



- Expert Verified, Online, **Free**.




CERTIFICATION TEST


- CertificationTest.net - Cheap & Quality Resources With Best Support

An engineer is implementing a Cisco ACI data center network that includes Cisco Nexus 2000 Series 10G fabric extenders. Which physical topology is supported?

A.




B.




C.

Nexus 9K Switches – ACI Mode



D.

Nexus 9K Switches – ACI Mode



Correct Answer: D

CiscoACI Highly Voted 4 years, 5 months ago

D correct

upvoted 8 times

designated Most Recent 8 months, 1 week ago

Correct answer is D

Note: Cisco does not support vPC connections between a FEX and 2 leaf switches, so each FEX must only be connected to 1 leaf switch.

Reference:

<https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/200232-Configuring-vPC-from-a-Host-to-Fabric-Ex.html>

upvoted 1 times

[Removed] 10 months ago

D is exact answer.

upvoted 1 times

eric0430 2 years ago

Any new exam takers? Is this still valid? I heard that Cisco Data Center certification is getting some revisions.

upvoted 1 times

student1956 2 years, 10 months ago

FEX (N2K Switchs) must be connected to one parent switch (N9K) .

D is correct

upvoted 1 times

  **aghibear** 2 years, 3 months ago

i failed the exam, all que

stions in the exam are different from the exam topics questions

upvoted 2 times

An ACI administrator notices a change in the behavior of the fabric. Which action must be taken to determine if a human intervention introduced the change?

- A. Inspect event records in the APIC UI to see all actions performed by users.
- B. Inspect /var/log/audit_messages on the APIC to see a record of all user actions.
- C. Inspect audit logs in the APIC UI to see all user events.
- D. Inspect the output of show command history in the APIC CLI.

Correct Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **instancepau** Highly Voted 👍 4 years, 2 months ago

should be audit logs.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/faults/guide/b_APIC_Faults_Errors/b_IFC_Faults_Errors_chapter_010.html

upvoted 13 times

🗳️ 👤 **[Removed]** Highly Voted 👍 4 years, 5 months ago

Answer should be Audit Logs to see what each and every user did on APIC

upvoted 9 times

🗳️ 👤 **designated** Most Recent 🕒 8 months, 1 week ago

C is correct

History > Audit Logs > Shows what each user did on APIC.

"Audit log—Holds objects that are records of user-initiated events such as logins and logouts (aaa:SessionLR) or configuration changes (aaa:ModLR) that are

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/faults/guide/b_APIC_Faults_Errors/b_IFC_Faults_Errors_chapter_010.html#cc

upvoted 1 times

🗳️ 👤 **[Removed]** 10 months ago

Selected Answer: C

C is right answer.

upvoted 1 times

🗳️ 👤 **mcsemcityp** 1 year, 11 months ago

Passed it today. All questions are in this pool, but you need to study all the comments as well.

upvoted 4 times

🗳️ 👤 **eric0430** 2 years ago

Any new exam takers? Is this still valid? I heard that Cisco Data Center certification is getting some revisions.

upvoted 1 times

🗳️ 👤 **7korn7** 2 years, 4 months ago

Selected Answer: C

C is the correct - checked on lab

upvoted 1 times

🗳️ 👤 **Supreme_123** 2 years, 7 months ago

Correct answer is C

upvoted 1 times

🗳️ 👤 **Gab99** 2 years, 8 months ago

Selected Answer: C

agree also seems likely



upvoted 1 times

🗳️ 👤 **patopizarro** 3 years, 2 months ago

Selected Answer: C

it should be Option C



upvoted 1 times

  **SS7640** 3 years, 4 months ago

Selected Answer: C

it should be Option C



upvoted 1 times

  **KSM03** 3 years, 7 months ago

Selected Answer: C

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/faults/guide/b_APIC_Faults_Errors/b_IFC_Faults_Errors_chapter_010.html

upvoted 3 times



  **KSM03** 3 years, 7 months ago

Please update answer or provide proof.

Answer should be:

C. Inspect audit logs in the APIC UI to see all user events.

upvoted 1 times

  **masal** 3 years, 9 months ago

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/faults/guide/b_APIC_Faults_Errors.pdf

Audit Logs is the answer, answer is C

upvoted 1 times

  **BigD69** 4 years ago

Anytime you see Human Intervention think AUDIT LOGS.

Each new event record MO is added to one of three separate event logs, depending on the cause of the event:

Audit log—Holds objects that are records of user-initiated events such as logins and logouts (aaa:SessionLR) or configuration changes (aaa:ModLR) that are required to be auditable.

Health score log—Holds records of changes in the health score (health:Record) of the system or components.

Event log—Holds records of other system-generated events (event:Record) such as link state transitions.

upvoted 3 times

  **mrpeet** 4 years ago

Correct answer is A. Audit log does not contain events as mentioned in C - events can be viewed in the Events viewer.

upvoted 1 times

  **thiyagas** 4 years, 5 months ago

Audit Logs

upvoted 3 times

An engineer is creating a configuration import policy that must terminate if the imported configuration is incompatible with the existing system. Which import mode achieves this result?

- A. merge
- B. atomic
- C. best effort
- D. replace

Correct Answer: B

Community vote distribution

B (100%)

  **hybersat** Highly Voted 4 years, 5 months ago

Atomic is correct..

Ref:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/basic-configuration/Cisco-APIC-Basic-Configuration-Guide-42x/Cisco-APIC-Basic-Configuration-Guide-42x_chapter_0100.html

importMode

Best-effort mode: each MO is applied individually, and errors only cause the invalid MOs to be skipped.

Note

If the object is not present on the controller, none of the children of the object get configured. Best-effort mode attempts to configure the children of the object.

Atomic mode: configuration is applied by whole shards. A single error causes whole shard to be rolled back to its original state.

upvoted 8 times

  **Nelsonnahun21** Most Recent 1 month, 1 week ago

Selected Answer: B



Atomic is correct

Atomic—ignores shards that contain objects that cannot be imported while proceeding with shards that can be imported. If the version of the incoming configuration is incompatible with the existing system, the import terminates.

Ref:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/aci-fundamentals/Cisco-ACI-Fundamentals-401/Cisco-ACI-Fundamentals-401_chapter_01011.html

upvoted 1 times

  **designated** 8 months, 1 week ago

B is correct:

Atomic—ignores shards that contain objects that cannot be imported while proceeding with shards that can be imported. If the version of the incoming configuration is incompatible with the existing system, the import terminates.

upvoted 1 times

  **[Removed]** 10 months ago

Selected Answer: B

B right answer.

upvoted 1 times

  **Mr_Certifiable** 1 year, 10 months ago

Selected Answer: B

Atomic—ignores shards that contain objects that cannot be imported while proceeding with shards that can be imported. If the version of the incoming configuration is incompatible with the existing system, the import terminates

Best-effort—ignores objects within a shard that cannot be imported. If the version of the incoming configuration is incompatible with the existing system, shards that are incompatible are not be imported while the import proceeds with those that can be imported.

Atomic Replace—overwrites existing configuration with imported configuration data. Any objects in the existing configuration that do not exist in the imported configuration are deleted.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/aci-fundamentals/Cisco-ACI-Fundamentals-401/Cisco-ACI-Fundamentals-401_chapter_01011.html

upvoted 2 times

  **eric0430** 2 years ago

Any new exam takers? Is this still valid? I heard that Cisco Data Center certification is getting some revisions.

upvoted 1 times

  **eddyedwards257** 2 years, 11 months ago

Atomic is correct : Merge and replace are Import Types not import Modes

upvoted 2 times

Which components must be configured for the BGP Route Reflector policy to take effect?

- A. spine fabric interface overrides and profiles
- B. access policies and profiles
- C. pod policy groups and profiles
- D. leaf fabric interface overrides and profiles

Correct Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **[Removed]** Highly Voted 👍 4 years, 5 months ago
Correct answer is C. Create Pod Policies .

<https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/aci-guide-configuring-fabric-bgp-route-reflectors.pdf>
upvoted 13 times

🗳️ 👤 **[Removed]** Highly Voted 👍 4 years, 5 months ago
Correct Answer :C .Pod policy and profile .
upvoted 9 times

🗳️ 👤 **[Removed]** Most Recent ⌚ 10 months ago
Selected Answer: C
The right answer is C.
upvoted 1 times

🗳️ 👤 **Chalmisco** 1 year, 5 months ago
Answer C - however from APIC version 3.0(1k) and later you would . Go to System > System Settings > BGP Route Reflector. See <https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/aci-guide-configuring-fabric-bgp-route-reflectors.pdf>
upvoted 2 times

🗳️ 👤 **Mr_Certifiable** 1 year, 10 months ago
Selected Answer: C
To configure the ACI Fabric BGP route reflectors, do the following:
1. Go to Fabric > Fabric Policies > Pod Policies > Policies > BGP Route Reflector default.
<https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/aci-guide-configuring-fabric-bgp-route-reflectors.pdf>
upvoted 2 times

🗳️ 👤 **GustavoF** 1 year, 11 months ago
C is correct.
upvoted 1 times

🗳️ 👤 **eric0430** 2 years ago
Any new exam takers? Is this still valid? I heard that Cisco Data Center certification is getting some revisions.
upvoted 1 times

🗳️ 👤 **student1956** 2 years, 10 months ago
Prior to APIC 3.0 , u must configure Pod Policy and Profile
APIC 3.0(1k) and Later , u should go to System > System settings > BGP Route Reflector
upvoted 4 times

🗳️ 👤 **patopizarro** 3 years, 2 months ago
Selected Answer: C
Correct answer is C

upvoted 3 times

🗲️ 👤 **ciscoaci2022** 3 years, 4 months ago

Correct answer is C, go to Fabric - Fabric Policy - Pod - Policy Group - Default

upvoted 3 times

🗲️ 👤 **SS7640** 3 years, 4 months ago

Selected Answer: C

Correct Answer should be Option C

upvoted 2 times

🗲️ 👤 **KSM03** 3 years, 7 months ago

Selected Answer: C

Correct answer: C

Reference: <https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/guide-c07-743150.html#2DistributeexternalrouteswithintheACIfabric>

upvoted 3 times

🗲️ 👤 **KSM03** 3 years, 7 months ago

Correct answer: C

Reference: <https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/guide-c07-743150.html#2DistributeexternalrouteswithintheACIfabric>

upvoted 1 times

🗲️ 👤 **thiyagas** 4 years, 5 months ago

Pod Policies

upvoted 6 times

Which type of policy configures the suppression of faults that are generated from a port being down?

- A. fault lifecycle assignment
- B. event lifecycle assignment
- C. fault severity assignment
- D. event severity assignment

Correct Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Irgond07** Highly Voted 4 years, 5 months ago

Options "C" is correct!

upvoted 5 times

🗳️ 👤 **designated** Most Recent 7 months, 2 weeks ago

C is correct, and the following will change the default fault:

... Policies -> Monitoring -> Common Policy -> Fault Severity Assignment Policies

upvoted 1 times

🗳️ 👤 **[Removed]** 10 months ago

Selected Answer: C

C is perfect answer.

upvoted 1 times

🗳️ 👤 **Mr_Certifiable** 1 year, 10 months ago

Selected Answer: C

Every APIC fault has a default severity. a fault might be considered more or less

severe than the default level. you might want to ignore a particular fault and squelch (suppress) it appearing in fault reports or status dashboards.

APIC two locations from which you can change the severity of a fault type:

- Directly from the Faults tab under a component in the APIC GUI
- In a monitoring policy

In the Faults tab of an APIC GUI , you can change the severity of a displayed fault or you can

suppress the auto-created fault squelch policy can be stored in the Fault

Severity Assignment Policies under one of the following:

- Tenants > common > Policies > Monitoring > default
- Fabric > Access Policies > Policies > Monitoring > default
- Fabric > Fabric Policies > Policies > Monitoring > default

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/faults/guide/b_APIC_Faults_Errors.pdf

upvoted 3 times

🗳️ 👤 **eric0430** 2 years ago

Any new exam takers? Is this still valid? I heard that Cisco Data Center certification is getting some revisions.

upvoted 1 times

Which type of profile needs to be created to deploy an access port policy group?

- A. attachable entity
- B. Pod
- C. module
- D. leaf interface

Correct Answer: D

Community vote distribution

D (75%)



A (25%)

 **Irgond07**  4 years, 5 months ago

Option "A" is the correct answer however Option "D" (interface profile) is also required. So, in this case there should be multiple choice question.
upvoted 11 times


 **n3tn3rd**  3 years, 7 months ago

Tricky Question! I'm not convinced AEP is correct as you can create an IPG without specifying an AEP! The question is to 'deploy' an IPG and I'd argue that's done by adding it to an interface selector, so my pick would be 'D'
upvoted 6 times

 **ashwind123**  3 months, 2 weeks ago

Selected Answer: A

The answer should be A as the question explicitly ask for access POLICY GROUP and under this we must attached AEP else wont work. D would make sense if the question about Leaf Interface profile which is mapping out the interface selector to respective access profiles.
upvoted 1 times

 **udo2020** 4 months, 2 weeks ago

Selected Answer: D

A leaf profile is the only possible solution. An attachable entity is no profile.
upvoted 1 times

 **designated** 7 months, 2 weeks ago


Selected Answer: D

AAEP is a profile while leaf interface is a Leaf Access Port policy group (tricky question cause it is asking about POLICY and not about PROFILE.
Fabric > Access Policies > Interfaces > Leaf Access Port > Create Leaf Access Port Policy Group
You will need to select an option for Attached Entity Profile (AEP) already created for this.
upvoted 1 times

 **Ceara** 1 year ago

Selected Answer: D

AEP is a policy, not a profile.
The interface profile needs to be created (Access Policies - Interfaces - Leaf Interfaces - Profile)
upvoted 2 times

 **SysAd** 1 year, 4 months ago

I also think its D because the question asked what is needed to deploy the policy group, it means he already have the policy group. AEP is part of the policy group, then you need the interface selector to deploy the policy group. If the question asked what is needed to create the policy group then it is AEP.
upvoted 1 times

 **msalamehi** 1 year, 4 months ago

Selected Answer: D

it is Leaf interface profile
upvoted 1 times

 **msalamehi** 1 year, 4 months ago

after thinking it is A, from Access port policy we link attached entity profile
upvoted 1 times

🗨️ 👤 **Mr_Certifiable** 1 year, 10 months ago

Selected Answer: A

n AAEP - Go to Fabric > Access Policies > Global Policies > Attachable Access Entity Profile

This is the "glue" that connects our domains (i.e., Physical, VMM and External routed Domains) and our Vlan Pool to Switches and Switch Interfaces.

<https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/aci-guide-configuring-fabric-access-policies.pdf>

upvoted 1 times

🗨️ 👤 **rss_01** 2 years, 6 months ago

Question seems to be related to a "deploy" action rather than to a create/configure one. I'd say D in this case

upvoted 3 times

🗨️ 👤 **nyanachi** 2 years, 8 months ago

When setting the access port policy group, only AEP is required, no leaf interface is required. A is correct, D is wrong.

upvoted 2 times

🗨️ 👤 **iulianm** 2 years, 10 months ago

I think D https://www.networklife.net/images/sheets/Networklife_CheatSheet_ACI_02_Fabric_access_policies.pdf

upvoted 1 times

🗨️ 👤 **student1956** 2 years, 10 months ago

The only Profile that should be assigned to Policy Group is attachable Entity Profile ,

I will go for answer A

upvoted 1 times

🗨️ 👤 **eddyedwards257** 2 years, 11 months ago

Think it is D: The question is what type of profile needs to be created , Interface Selectors are required in the Profile , nothing in the Policy Groups talks about "Profiles"

upvoted 3 times

🗨️ 👤 **rickyarchi** 3 years, 2 months ago

I think D is correct because you can create an interface profile without AEEP

upvoted 2 times

🗨️ 👤 **Kalpesh** 4 years, 4 months ago

I think D is correct, because its either Attachable Access entity or Attached Entity profile, not Attachable entity profile.

upvoted 1 times

🗨️ 👤 **RTL_dude** 4 years, 2 months ago

AEP and AAEP are the same thing, Cisco changed the terminology from a certain version.

upvoted 2 times

🗨️ 👤 **reminetwork** 4 years, 3 months ago

A is correct, you dont need leaf interface profile to create an Access Policy Port Group

<https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/aci-guide-configuring-fabric-access-policies.pdf>

upvoted 9 times

A situation causes a fault to be raised on the APIC. The ACI administrator does not want that fault to be raised because it is not directly relevant to the environment. Which action should the administrator take to prevent the fault from appearing?

- A. Under System -> Faults, right-click on the fault and select Acknowledge Fault so that acknowledged faults will immediately disappear.
- B. Create a stats threshold policy with both rising and falling thresholds defined so that the critical severity threshold matches the squelched threshold.
- C. Under System -> Faults, right-click on the fault and select Ignore Fault to create a fault severity assignment policy that hides the fault.
- D. Create a new global health score policy that ignores specific faults as identified by their unique fault code.

Correct Answer: C

Community vote distribution

C (100%)

🗲️ 👤 **designated** 7 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗲️ 👤 **zelya19** 1 year, 3 months ago

Selected Answer: C

Available options are:

Ignore Fault

Change Severity

Save as ...

Post ...

(greyed out) Share

(greyed out) Open in Object Store Browser

upvoted 2 times

🗲️ 👤 **Mr_Certifiable** 1 year, 10 months ago

Selected Answer: C

you might want to ignore a particular fault and squelch (suppress) it from appearing in fault reports or status dashboards. APIC provides two locations from which you can change the severity of a fault type:

Directly from the Faults tab under a component in the APIC GUI

In a monitoring policy

System > Faults or Tenant > name > name > Faults

System > Dashboard or Tenant > name > name > Dashboard.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/faults/guide/b_APIC_Faults_Errors.pdf

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/faults/guide/b_APIC_Faults_Errors/b_IFC_Faults_Errors_chapter_01.html

upvoted 1 times

🗲️ 👤 **pradeepmore** 4 years, 1 month ago

correct answer is C


upvoted 3 times

A RADIUS user resolves its role via the Cisco AV Pair. What object does the Cisco AV Pair resolve to?

- A. tenant
- B. security domain
- C. primary Cisco APIC
- D. managed object class

Correct Answer: D

Community vote distribution



nikomski Highly Voted 3 years, 11 months ago

Shouldn't this be B?

upvoted 17 times

prospio971 Most Recent 4 months, 2 weeks ago

Selected Answer: B

A managed object class represents ACI's object structure, but the AV Pair is used for authentication and authorization, which is tied to Security Domains.

upvoted 1 times

2eb1ea8 5 months ago

Selected Answer: B

Cisco AV Pair:

The Cisco AV (Attribute-Value) Pair is a RADIUS attribute used to assign roles and permissions to users in Cisco ACI. It typically includes information such as the user's role and the security domain they belong to.

Security Domain:

The security domain is a logical grouping of resources (e.g., tenants, EPGs, or other managed objects) that defines the scope of access for a user. When a RADIUS user authenticates, the Cisco AV Pair maps the user to a specific security domain, which determines what resources the user can access.

upvoted 2 times

udo2020 6 months ago

Selected Answer: D

I agree to D. The Cisco AV Pair (Attribute-Value Pair) is used in RADIUS to communicate attributes about a user. In the context of Cisco ACI, the AV Pair is used to assign roles and permissions to users.

When a RADIUS user authenticates, the Cisco AV Pair determines the user's role and permissions. These roles map to specific managed object classes within the Cisco ACI framework, allowing the user to access and interact with the corresponding objects in the APIC.

upvoted 1 times

[Removed] 7 months, 3 weeks ago

Selected Answer: D

Took some back and forth to understand what the ask was. The AV Pairs are carrying with them the privileges and roles(grouped privileges). The privileges are managed objects that you can't custom create, however you can regroup them in many ways to create a suitable access level.

The key to the answer is understanding what the whole string has. It does somewhat confuse a lot because of the word Security Domains on the beginning of the string, security domains are part of the structure as Tags or tenants etc.

<https://bestpath.io/cisco-aci-rbac/>

upvoted 2 times

zelya19 9 months, 4 weeks ago

Selected Answer: B

shell:domains =

SecurityDomainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,

SecurityDomainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2

upvoted 1 times

🗨️ 👤 **Redou2201** 1 year, 2 months ago

Selected Answer: D

based on the link in the comment I saw that: Roles and Privileges

A privilege controls access to a particular function within the system. The ACI fabric manages access privileges at the managed object (MO) level.
so for me it is D

upvoted 2 times

🗨️ 👤 **Mr_Certifiable** 1 year, 4 months ago

Selected Answer: D

AV Pair on the External Authentication Server

The Cisco APIC requires that an administrator configure a Cisco AV Pair on an external authentication server.

The Cisco AV pair specifies the APIC required RBAC roles and privileges for the user. The Cisco AV Pair format is the same for RADIUS, LDAP, or TACACS+.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/basic-configuration/Cisco-APIC-Basic-Configuration-Guide-42x/Cisco-APIC-Basic-Configuration-Guide-42x_chapter_011.html

For each of the defined roles in Cisco APIC, the APIC Roles and Privileges Matrix shows which managed object classes can be written and which can be read. <https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/apicroles/roles.html>

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/Security_config/b_Cisco_APIC_Security_Configuration_Guide/b_Cisco_APIC_Security_Guide_chapter_01000.html

upvoted 2 times

🗨️ 👤 **nyanachi** 2 years, 2 months ago

The ACI fabric manages access privileges at the managed object (MO) level.

Answer is D

upvoted 1 times

🗨️ 👤 **bizzar777** 2 years, 2 months ago

Selected Answer: D

managed object class is an object ; security domain is a tag

upvoted 1 times

🗨️ 👤 **Alphonza** 2 years, 7 months ago

Selected Answer: B

B is the answer Based on the following link

https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/basic-configuration/cisco-apic-basic-configuration-guide-51x/m_aaa.html

upvoted 2 times

🗨️ 👤 **Smoothery** 2 years, 8 months ago

Selected Answer: D

The ACI fabric manages access privileges at the managed object (MO) level. A privilege is an MO that enables or restricts access to a particular function within the system. For example, fabric-equipment is a privilege bit. This bit is set by the Application Policy Infrastructure Controller (APIC) on all objects that correspond to equipment in the physical fabric.

A role is a collection of privilege bits. For example, because an "admin" role is configured with privilege bits for "fabric-equipment" and "tenant-security," the "admin" role has access to all objects that correspond to equipment of the fabric and tenant security.

upvoted 1 times

🗨️ 👤 **Smoothery** 2 years, 8 months ago

Tricky. AV-Pairs are associated to a security domain/s. Then permissions assigned e.g. read and or write or both as you select the roles. The MO are assigned to various roles. The roles in turn are prescribed in the security domain. Therefore I am leaning more towards security domain for the AV as the SD is where you choose the MO. However with that said it would ideally involve both the SD and the MO as they are linked.

upvoted 3 times

🗨️ 👤 **ciscoaci2022** 2 years, 10 months ago

Should be "D"

A security domain is a tag associated with a certain subtree in the ACI MIT object hierarchy. For example, the default tenant "common" has a domain tag common. Similarly, the special domain tag all includes the entire MIT object tree. An administrator can assign custom domain tags to the MIT object hierarchy. For example, an administrator could assign the "solar" domain tag to the tenant named solar. Within the MIT, only certain objects

can be tagged as security domains. For example, a tenant can be tagged as a security domain but objects within a tenant cannot.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/b-Cisco-APIC-Security-Configuration-Guide-421/b-Cisco-APIC-Security-Configuration-Guide-421_chapter_011.html



upvoted 2 times

  **Annielover007** 2 years, 11 months ago

Selected Answer: B

B is correct

upvoted 2 times

  **jim13c** 3 years, 5 months ago

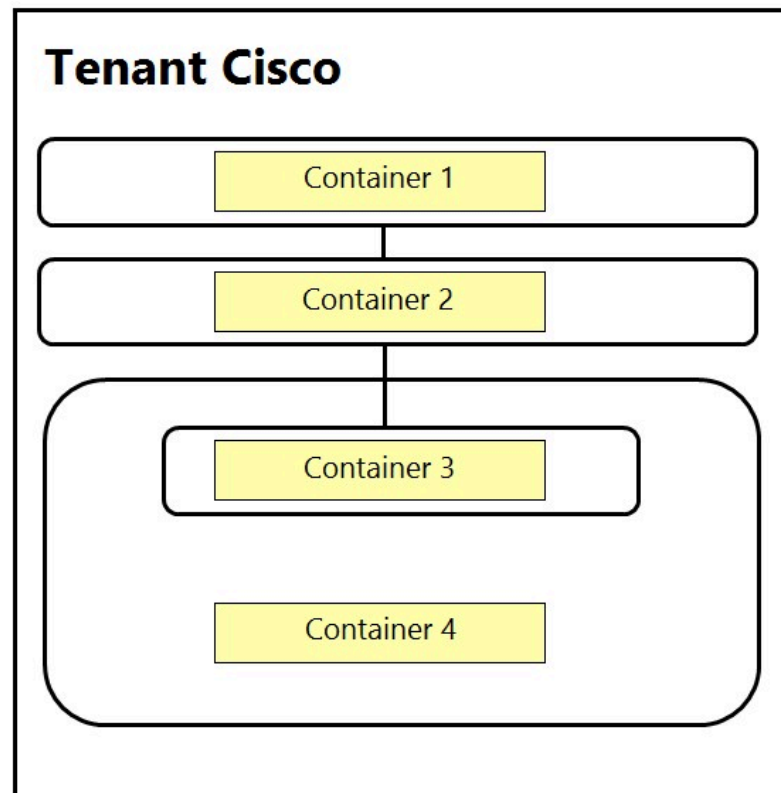
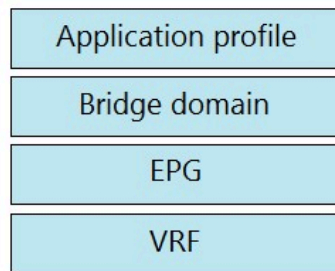
Agreed, should be B

upvoted 3 times

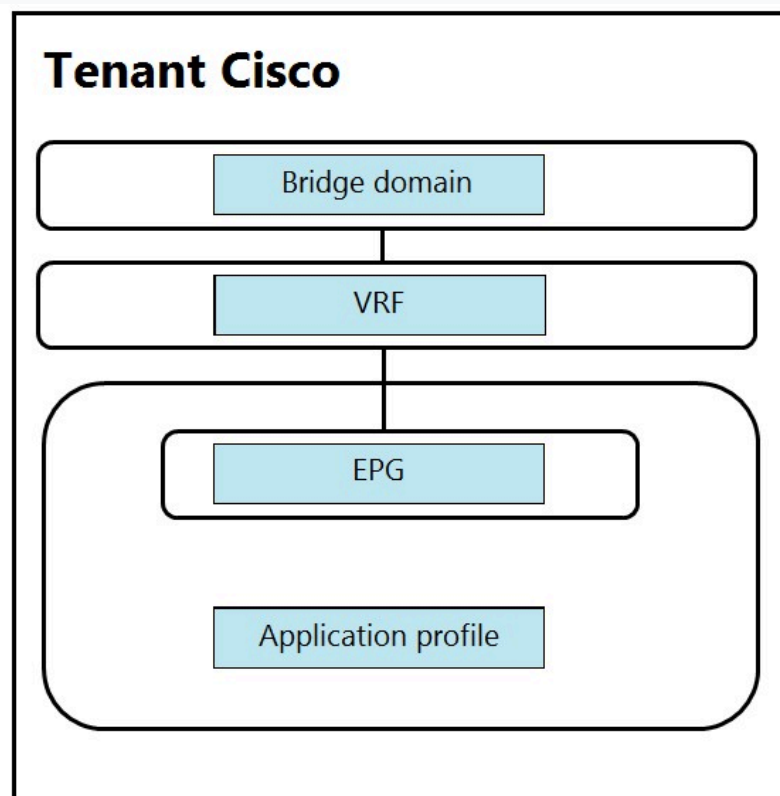
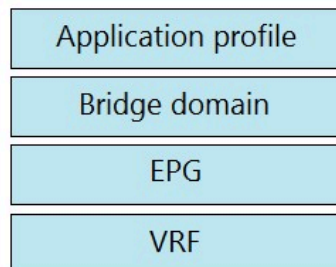
DRAG DROP -

An engineer is configuring a VRF for a tenant named Cisco. Drag and drop the child objects on the left onto the correct containers on the right for this configuration.

Select and Place:



Correct Answer:





[Removed] 3 years, 11 months ago

Correct Answer : VRF should come first in this hierarchy.

VRF-->BD->EPG->Application Profile .

upvoted 28 times

  **iulianm** 2 years, 4 months ago

From design it should be VRF - BD-EPG-AP AP can contain more EPG. In design you can see AP contain EPG

upvoted 3 times

  **[Removed]**  3 years, 11 months ago

The Answer is not entirely correct, Bridge Domain is a child object of a VRF

upvoted 10 times

  **[Removed]**  7 months, 3 weeks ago

I see commented here three years ago without offering an answer. But I am back, older and a little wiser.

VRF --- BD--- The lower bigger box is App profile---EPG.

The EPG is not configurable outside of an Application Profile,, it definitely remains the child of the AP. BD is clearly below the VRF . But if you answering without paying attention of the box nestings,, it would look like VRF-BD-EPG---App Profile, which is a bad representation of the truth.

upvoted 2 times

  **skulligno** 11 months, 3 weeks ago

Correct Answer :VRF->BD->EPG->Application Profile .

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/3-x/rest-api-config/b_Cisco_APIC_REST_API_Configuration_Guide_3x/b_Cisco_APIC_REST_API_Configuration_Guide_chapter_01011.html

upvoted 2 times

  **[Removed]** 1 year ago


According to the chapter "2.2 Describing Cisco ACI Policy Model Logical Constructs" image "ACI policy model logical constructs", the correct answer is "VRF > Bridge Domain > EPG > Application Profile" (source: official DCACI course)

upvoted 1 times

  **crooks_1988** 1 year, 3 months ago


I would also say: VRF -> BD ->App_Prof -> EPG

upvoted 2 times

  **tarachan291** 1 year, 4 months ago

VRF>BD>EPG>Application Profile

upvoted 2 times

  **boogycreek** 2 years, 2 months ago

Correct answer is application profile > VRF > BD > EPG

upvoted 2 times

  **Alphonza** 2 years, 7 months ago

VRF -> BD-> EPG-> App Profile Ref. Operating Cisco Application Centric Infrastructure pg 196

upvoted 6 times

  **StyleEast** 3 years, 1 month ago

Verified: VRF > BD > EPG > A.P.

upvoted 5 times

  **anasham** 3 years ago

BD refers to VRF

EPG refers to the BD and Configured within A.P



so VRF-BD->EPG-> app profile is correct down in the order the boxes are depicted.

upvoted 2 times

  **Ssay** 3 years, 1 month ago

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/Operating_ACI/guide/b_Cisco_Operating_ACI/b_Cisco_Operating_ACI_chapter_0111.html

upvoted 1 times

  **mr_siro** 3 years, 4 months ago

Yes, i think VRF -> BD -> EPG -> App_Prof

upvoted 2 times

  **Ssay** 3 years, 1 month ago

The correct answer is VRF -> BD -> App_Prof -> EPG

Explain in this link : https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/Operating_ACI/guide/b_Cisco_Operating_ACI/b_Cisco_Operating_ACI_chapter_0111.html

upvoted 5 times

Which feature dynamically assigns or modifies the EPG association of virtual machines based on their attributes?

- A. vzAny contracts
- B. standard contracts
- C. application EPGs
- D. uSeg EPGs

Correct Answer: D

🗨️ 👤 **Irgond07** Highly Voted 11 months ago

Options "D" is correct!

Microsegmented EPGs (uSeg) works based on attribute.

upvoted 10 times

🗨️ 👤 **prospio971** Most Recent 4 months, 2 weeks ago

Selected Answer: D

uSeg EPGs (micro-segmented Endpoint Groups) dynamically assign or modify the EPG association of virtual machines based on their attributes, such as VM name, operating system, or tags. This feature is particularly useful in environments where endpoint group membership needs to be flexible and automated, such as in virtualized or cloud-based infrastructures.

upvoted 1 times

Which feature allows firewall ACLs to be configured automatically when new endpoints are attached to an EPG?

- A. ARP gleaning
- B. dynamic endpoint attach
- C. hardware proxy
- D. network-stitching

Correct Answer: B

Community vote distribution

B (100%)

 **tarq** Highly Voted 3 years, 11 months ago

Dynamic Endpoint Attach is a feature that allows Firewall ACLs or ADC load balanced servers to be dynamically configured when new endpoints are attached to an EPG.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-734298.html>
upvoted 7 times

 **[Removed]** Most Recent 7 months, 3 weeks ago

Selected Answer: B

Provided answer is correct

upvoted 1 times

An engineer is implementing Cisco ACI at a large platform-as-a-service provider using APIC controllers, 9396PX leaf switches, and 9336PQ spine switches. The leaf switch ports are configured as IEEE 802.1p ports. Where does the traffic exit from the EPG in IEEE 802.1p mode in this configuration?

- A. from leaf ports tagged as VLAN 0
- B. from leaf ports untagged
- C. from leaf ports tagged as VLAN 4094
- D. from leaf ports tagged as VLAN 1

Correct Answer: A

  **Irgond07**  4 years, 5 months ago

A is correct option.

FYI...

Trunk (Tagged - classic IEEE 802.1q trunk)—Traffic for the EPG is sourced by the leaf switch with the specified VLAN tag. The leaf switch also expects to receive traffic tagged with that VLAN to be able to associate it with the EPG. Traffic received untagged is discarded.

- Access (Untagged)—Traffic for the EPG is sourced by the leaf as untagged. Traffic received by the leaf switch as untagged or with the tag specified during the static binding configuration is associated with the EPG.


- Access (802.1p)—If only one EPG is bound to that interface, the behavior is identical as in the untagged case. If other EPGs are associated with the same interface, traffic for the EPG is sourced with an IEEE 802.1q tag using VLAN 0 (IEEE 802.1p tag), or is sourced as untagged in the case of EX switches.

upvoted 6 times

  **alphatest** 3 years ago

With first generation switches, it can be UNTAGGED if there's only one EPG associated to the port, or VLAN0 when there are more. The question is ambiguous


upvoted 1 times

  **designated**  7 months, 2 weeks ago

Selected Answer: A

A is correct

upvoted 1 times

  **jmaroto** 3 years, 1 month ago

- Access (IEEE 802.1p) or native: With Cisco Nexus 9300-EX and later switches, this option is equivalent to the Access (untagged) option. This option exists because of first generation leaf switches. On Cisco Nexus 9300-EX or later switches, you can assign the native VLAN to a port either by using the Access (untagged) option or the Access (IEEE 802.1p) option. However, we recommend that you use the Access (untagged) option because the Access (IEEE 802.1p) option was implemented specifically to address

https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-application-centric-infrastructure-design-guide.html#_Toc6452851

The question reference first generation leaf and spine switches, this is the key

upvoted 1 times

  **MaxG** 3 years, 11 months ago

A is correct.

The VLAN 0 Priority Tagging feature enables 802.1Q Ethernet frames to be transmitted with the VLAN ID set to zero. These frames are called priority tagged frames. Setting the VLAN ID tag to zero allows the VLAN ID tag to be ignored and the Ethernet frame to be processed according to the priority configured in the 802.1P bits of the 802.1Q Ethernet frame header.

https://www.cisco.com/c/en/us/td/docs/switches/connectedgrid/cg-switch-sw-master/software/configuration/guide/vlan0/b_vlan_0.html
upvoted 4 times

How is an EPG extended outside of the ACI fabric?


- A. Create an external bridged network that is assigned to a leaf port.
- B. Create an external routed network that is assigned to an EPG.
- C. Enable unicast routing within an EPG.
- D. Statically assign a VLAN ID to a leaf port in an EPG.

Correct Answer: D

Reference:

<https://www.dcclessons.com/l2-external-network-with-aci>

Community vote distribution



tarq **Highly Voted** 3 years, 11 months ago

Extending the EPG out of the ACI Fabric: An EPG can be extended out of ACI fabric by statically assigning port to an EPG. As soon as leaf receives the traffic and determines the end point information, it assigns the traffic to Proper EPG by matching the VLAN ID on port.

upvoted 11 times

[Removed] **Most Recent** 7 months, 3 weeks ago

Selected Answer: A

L2 External Network with ACI

This topic will help you to understand how and what methods are used to extend the layer 2 network outside the ACI fabric.

There are various methods extend the layer 2 domain beyond ACI fabric.

Extending the EPG out of the ACI Fabric: An EPG can be extended out of ACI fabric by statically assigning port to an EPG. As soon as leaf receives the traffic and determines the end point information, it assigns the traffic to Proper EPG by matching the VLAN ID on port.

Extending the Bridge Domain out of the ACI Fabric: It is also possible to extend the bridge domain by creating the layer 2 outside connection (External Bridge network). By doing so, it extend the bridge domain to the outside network.

upvoted 1 times

[Removed] 7 months, 3 weeks ago

I am really struggling trying to find why answer A is not correct? Because creating an external bridged network does extend an EPG to the outside of ACI Fabric.

upvoted 1 times

Mr_Certifiable 1 year, 4 months ago

Selected Answer: D

There are three ways of extending the Layer 2 domain outside the ACI fabric:

1. Manually assign a port to a VLAN which in turn gets mapped into an EPG.

This action extends the EPG beyond the ACI fabric.

2. Create a Layer 2 connection to the outside network beyond the ACI fabric. This action allows a contract between an inside EPG and an outside EPG. This action extends the bridge domain.

3. Create a remote VTEP.

<https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/5eU6DfQV/LTRCRT-2611.pdf>

upvoted 2 times

Anantharajesh 1 year, 5 months ago

Isn't the answer is external bridge domain?

upvoted 1 times

[Removed] 7 months, 3 weeks ago

External bridged network is correct and better.

upvoted 1 times

  **Supreme_123** 2 years, 1 month ago

I think correct answer is D

upvoted 1 times

DRAG DROP -

Drag and drop the Cisco ACI filter entry options from the left onto the correct categories on the right indicating what are required or optional parameters.

Select and Place:

Name	Optional Parameters
ARP Flag	
Ether Type	
IP Protocol	
Source Port From	Required Parameters

Correct Answer:

Name	Optional Parameters
ARP Flag	
Ether Type	
IP Protocol	
Source Port From	Required Parameters

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/Operating_ACI/guide/b_Cisco_Operating_ACI/b_Cisco_Operating_ACI_chapter_01000.html

 **pepper_** Highly Voted 4 years, 3 months ago

Answer is right. https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/Operating_ACI/guide/b_Cisco_Operating_ACI/b_Cisco_Operating_ACI_chapter_01000.html
upvoted 7 times

 **designated** Most Recent 7 months, 2 weeks ago

Optional
ARP Flag
IP Protocol
Source Port From

Required
Name
Ethertype (even Unspecified)
upvoted 1 times

Where is the COOP database located?

- A. leaf
- B. spine
- C. APIC
- D. endpoint

Correct Answer: B

Reference:

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

  **pradeepmore** Highly Voted 4 years, 1 month ago

b spine

upvoted 6 times

  **designated** Most Recent 7 months, 2 weeks ago

Selected Answer: B

COOP database = Spine

upvoted 1 times

  **Pakawat** 3 years, 7 months ago

B is correct

upvoted 2 times

Which description regarding the initial APIC cluster discovery process is true?

- A. The APIC uses an internal IP address from a pool to communicate with the nodes.
- B. Every switch is assigned a unique AV by the APIC.
- C. The APIC discovers the IP address of the other APIC controllers by using Cisco Discovery Protocol.
- D. The ACI fabric is discovered starting with the spine switches.

Correct Answer: A

Community vote distribution

A (82%)

B (18%)

 **Alex321** Highly Voted 4 years, 5 months ago

Probably A is correct answer. D is wrong because APIC uses LLDP and not CDP during discovery process.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_010011.html:

"Each APIC in the Cisco ACI uses an internal private IP address to communicate with the ACI nodes and other APICs in the cluster. The APIC discovers the IP address of other APIC controllers in the cluster through the LLDP-based discovery process."

upvoted 26 times

 **Putaloka** Highly Voted 4 years, 5 months ago

Letter A is correct!

upvoted 11 times

 **designated** Most Recent 7 months, 2 weeks ago


Selected Answer: A

A is correct.

Address pool for TEP addresses: Enter the address pool for TEP addresses. The default IP address pool for TEP tunnel endpoint addresses is 10.0.0.0/16. This value is for the infrastructure VRF used for internal fabric communication. This subnet will not be exposed to your legacy network unless you are deploying the Cisco AVS/AVE. Regardless, the recommendation is to assign an unused subnet of size between and /16 and /21 subnet. The size of the subnet used will impact the scale of your pod. Most customers allocate an unused /16. This value cannot be changed once configured. Modifying this value requires a wipe of the fabric.

Reference: Cisco U


upvoted 2 times

 **imamus** 2 years, 2 months ago

Selected Answer: A

B is wrong because the AV is an object maintained by the APICs, to keep mappings between APIC IPs and IDs, not of the switches. The switches receive and advertise further AV pairs from one APIC to other APICs. Thus the correct one is A.

upvoted 2 times

 **imamus** 2 years, 2 months ago

B is wrong because the AV is an object maintained by the APICs, to keep mappings between APIC IPs and IDs, not of the switches. The switches receive and advertise further AV pairs from one APIC to other APICs. Thus the correct one is A.

upvoted 1 times

 **Robo2019** 2 years, 9 months ago

Answer is B,

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_010011.html#concept_A9C1C7BB9E8C4798AAD58CBA472B3BA5

>>When a switch reboots, the policy element (PE) on the leaf gets its AV from the APIC. The switch then advertises this AV to all of its neighbors and reports any discrepancies between its local AV and neighbors' AVs to all the APICs in its local AV.<<

upvoted 2 times

🗨️ 👤 **Robo2019** 2 years, 9 months ago

Selected Answer: B

Answer is B,

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_010011.html#concept_A9C1C7BB9E8C4798AAD58CBA472B3BA5

>>When a switch reboots, the policy element (PE) on the leaf gets its AV from the APIC. The switch then advertises this AV to all of its neighbors and reports any discrepancies between its local AV and neighbors' AVs to all the APICs in its local AV.<<

upvoted 2 times

🗨️ 👤 **patopizarro** 3 years, 2 months ago

Selected Answer: A

A is correct

upvoted 3 times

🗨️ 👤 **igebuntu** 3 years, 4 months ago

Agreed in A.

upvoted 1 times

🗨️ 👤 **Annielover007** 3 years, 4 months ago

Selected Answer: A

A is correct

upvoted 2 times

🗨️ 👤 **[Removed]** 3 years, 5 months ago

Selected Answer: A

A is correct

upvoted 2 times

🗨️ 👤 **mr_siro** 3 years, 10 months ago

A is correct

Explain:

The following describes the APIC cluster discovery process:

Each APIC in the Cisco ACI uses an internal private IP address to communicate with the ACI nodes and other APICs in the cluster. The APIC discovers the IP address of other APIC controllers in the cluster through the LLDP-based discovery process.

Link: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/m_provisioning.html#concept_8E5AD8328B9348D5B8F0FF1FF7D2EC96

upvoted 3 times

🗨️ 👤 **thiyagas** 4 years, 5 months ago

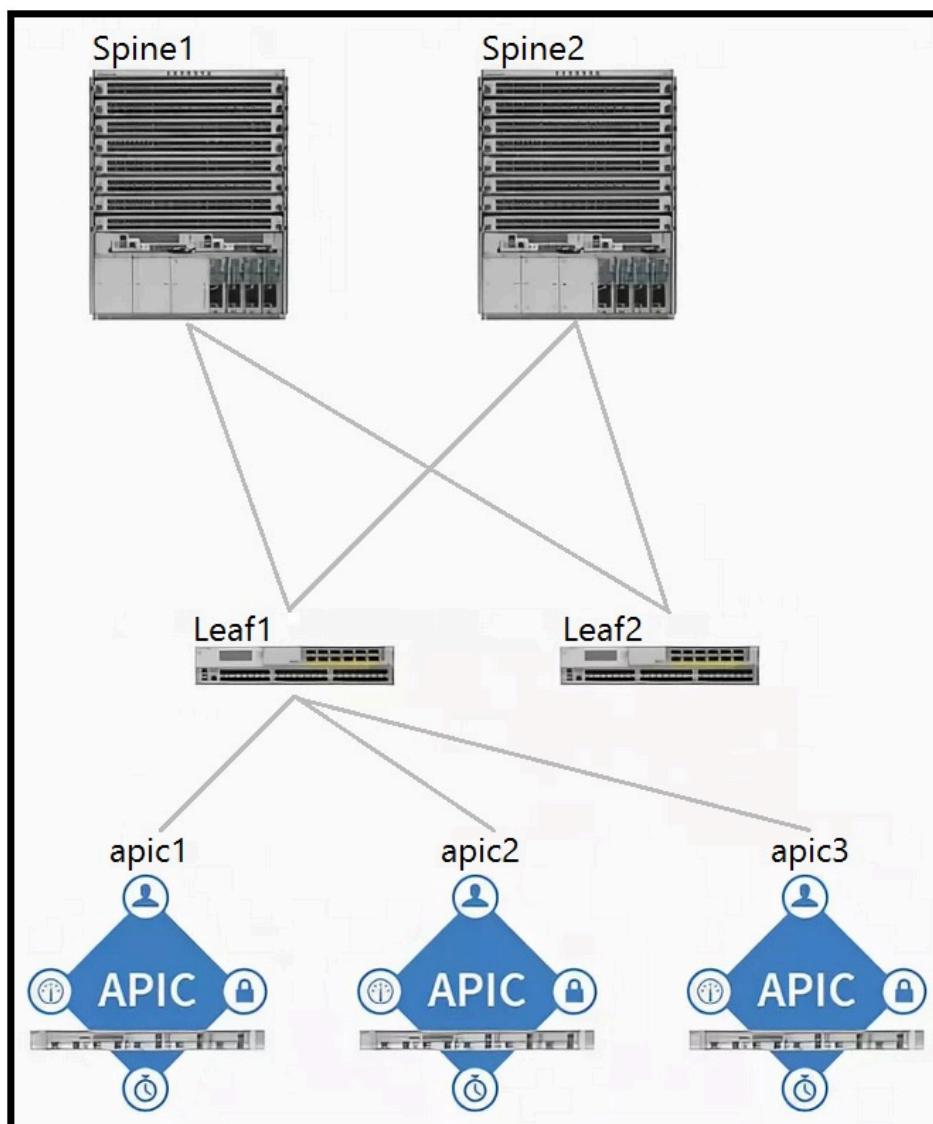
A is correct

upvoted 6 times

🗨️ 👤 **[Removed]** 4 years, 5 months ago

I concur.

upvoted 3 times



Refer to the exhibit. Which two components should be configured as route reflectors in the ACI fabric? (Choose two.)

- A. Spine1
- B. apic1
- C. Spine2
- D. Leaf1
- E. Leaf2
- F. apic2

Correct Answer: AC

pradeepmore Highly Voted 4 years, 1 month ago
a and c is correct , spine switches
upvoted 9 times

pradeepmore Highly Voted 4 years, 1 month ago
ans c is correct , spine switches
upvoted 5 times

designated Most Recent 7 months, 2 weeks ago
Selected Answer: AC
Spine Switches are always the route reflectors.
upvoted 1 times

When creating a subnet within a bridge domain, which configuration option is used to specify the network visibility of the subnet?

- A. limit IP learning to subnet
- B. scope
- C. gateway IP
- D. subnet control

Correct Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **thiyagas** Highly Voted 4 years, 5 months ago

B. Scope is correct answer
upvoted 16 times

🗳️ 👤 **jaboriel** Highly Voted 4 years, 5 months ago

shouldn't this be B (scope)?
upvoted 7 times

🗳️ 👤 **designated** Most Recent 7 months, 2 weeks ago

Selected Answer: B

B is correct.

Scope option has two options under Subnet creation:

- Advertised Externally
- Shared Between VRFs

Create Subnet:

Gateway IP: (address/mask)

Treat as virtual IP Address:

Make this IP address primary:

Scope:

- Advertised Externally
- Shared Between VRFs

Description

Subnet Control

- No Default SVI Gateway

Querier IP

IP Data-Plane Learning (disabled/enabled)

L3 Out for route profile (select value)

Policy Tag:

upvoted 1 times

🗳️ 👤 **zelya19** 1 year, 3 months ago

Selected Answer: B

Scope = The network visibility of the subnet.

upvoted 1 times

🗳️ 👤 **Chandera** 1 year, 10 months ago

We define subnet with BD Gateway IP .. example 192.168.1.1/24 ... So correct answer is C

upvoted 1 times

🗳️ 👤 **Gab99** 2 years, 8 months ago

Selected Answer: B

no brainer

upvoted 1 times

🗨️ 👤 **ciscoaci2022** 3 years, 4 months ago

B is correct, the Scope has three Options: Private to VRF, Advertised to Externally, and Shared between VRFs
upvoted 3 times

🗨️ 👤 **StyleEast** 3 years, 7 months ago

Answer is B.

upvoted 2 times

🗨️ 👤 **KSM03** 3 years, 7 months ago

Selected Answer: B

[https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/Operating_ACI/guide/b_Cisco_Operating_ACI/b_Cisco_Operating_ACI_chapter_0111.html#concept_247562B33DE34FC5ADBB5D516DA085EB)

[x/Operating_ACI/guide/b_Cisco_Operating_ACI/b_Cisco_Operating_ACI_chapter_0111.html#concept_247562B33DE34FC5ADBB5D516DA085EB](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/Operating_ACI/guide/b_Cisco_Operating_ACI/b_Cisco_Operating_ACI_chapter_0111.html#concept_247562B33DE34FC5ADBB5D516DA085EB)

Subnets—The network visibility of the subnet. The subnet is a portion of a network sharing a particular subnet address. The scope can be:

Shared Between VRFs—Defines subnets under an endpoint group, with the Shared option configured, to route leak to other tenants within the Fabric.

Advertise Externally—Defines subnets under a bridge domain, with the Public option configured, to share with Layer 3 outbound.

Private to VRF—Defines subnets under a bridge domain, with the Private option configured, to only be used in that tenant (will not be leaked). The default is Private.

upvoted 4 times

🗨️ 👤 **navinm** 4 years ago

The application profile is a set of requirements that an application instance has on the virtualizable fabric. The policy regulates connectivity and visibility among endpoints within the scope of the policy.

upvoted 1 times

🗨️ 👤 **[Removed]** 4 years, 5 months ago

Correct Answer: C .Gateway should be configured for subnet visibility of the respective bridge domain .

upvoted 1 times

🗨️ 👤 **[Removed]** 4 years, 5 months ago

Sorry for previous comments .B should be correct Answer for any network Visibility in ACI

upvoted 5 times

What does a bridge domain represent?

- A. Layer 3 cloud
- B. Layer 2 forwarding construct
- C. tenant
- D. physical domain

Correct Answer: B

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/L2_config/b_Cisco_APIC_Layer_2_Configuration_Guide/b_Cisco_APIC_Layer_2_Configuration_Guide_chapter_010.html

Community vote distribution

B (100%)

🗲️ 👤 **pradeepmore** Highly Voted 🍌 3 years, 7 months ago

B. Layer 2 forwarding construct

upvoted 6 times

🗲️ 👤 **zelya19** Most Recent 🕒 9 months, 4 weeks ago

Selected Answer: B

A bridge domain (fvBD) represents a Layer 2 forwarding construct within the fabric.

upvoted 1 times

🗲️ 👤 **eric0430** 1 year, 6 months ago

Any new exam takers? Is this still valid? I heard that Cisco Data Center certification is getting some revisions.

upvoted 1 times

Which table holds IP address, MAC address and VXLAN/VLAN information on a Cisco ACI leaf?

- A. endpoint
- B. adjacency
- C. RIB
- D. ARP

Correct Answer: A

Community vote distribution

A (100%)

  **pradeepmore** Highly Voted 3 years, 1 month ago

A. endpoint
upvoted 6 times



  **Mr_Certifiable** Most Recent 10 months, 2 weeks ago

Selected Answer: A

You can see the routes, ARP, CDP, LLDP, many other values under the Fabric > Inventory > Pod N > Leaf > Protocols, but it appears the endpoint tables and MAC tables are not visible in the GUI.

<https://community.cisco.com/t5/application-centric-infrastructure/display-local-endpoints-on-a-leaf-from-gui/td-p/3742617>

upvoted 2 times

  **Abrieg** 1 year, 4 months ago

Selected Answer: A

Cisco ACI replaced the MAC address table and ARP table with a single table called the endpoint table.

upvoted 2 times

Which two types of interfaces are supported on border leaf switches to connect to an external router? (Choose two.)

- A. subinterface with VXLAN tagging
- B. subinterface with 802.1Q tagging
- C. FEX host interface
- D. out of band interface
- E. Switch Virtual Interface

Correct Answer: BE

  **tarq** Highly Voted 4 years, 4 months ago

Out of given options, 2 interface types are mentioned, 3rd is L3 (routed) interface .

Border leaf switches can be configured with three types of interfaces to connect to an external router:

- Layer 3 (routed) interface
- Subinterface with IEEE 802.1Q tagging
- Switch Virtual Interface (SVI)

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737909.html>

upvoted 11 times

  **LV2022** Most Recent 5 months, 2 weeks ago

Selected Answer: BE

Border leaf switches can be configured with three types of interfaces to connect to an external router:

- Layer 3 (routed) interface
- Subinterface with IEEE 802.1Q tagging
- Switch Virtual Interface (SVI)

<https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-application-centric-infrastructure-design-guide.html>

upvoted 2 times

  **udo2020** 6 months ago

Selected Answer: BE

I think it should be B and E.

A can't be right, because VXLAN tagging is used internally within the ACI fabric, not for connections to external routers. Border leaf switches use VXLAN within the fabric but not for external interfaces.

upvoted 2 times

  **designated** 7 months, 2 weeks ago

Selected Answer: AB

A, B and E are correct:

Create L3 Out:

Interface Types:

> Layer3:

- Interface (default)
- Subinterface (than could be Encap: VLAN or VXLAN)
- SVI
- Floating SVI

> Layer 2:

- Port (default)
- Direct Port Channel

upvoted 1 times

  **eric0430** 2 years ago

Any new exam takers? Is this still valid? I heard that Cisco Data Center certification is getting some revisions.

upvoted 1 times

Which Cisco APIC configuration prevents a remote network that is not configured on the bridge domain from being learned by the fabric?

- A. enable Limit IP Learning to Subnet
- B. enable Unicast Routing
- C. enable IP Data-plane Learning
- D. enable ARP Flooding to BD

Correct Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **pradeepmore** Highly Voted 4 years, 1 month ago

A. enable Limit IP Learning to Subnet
upvoted 9 times

🗳️ 👤 **LV2022** Most Recent 5 months, 2 weeks ago

Selected Answer: A

Limit IP Learning to Subnet: This option is similar to a unicast reverse-forwarding-path check. If this option is selected, the fabric will not learn IP addresses from a subnet other than the one configured on the bridge domain.

<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2->

[x/L2_config/b_Cisco_APIC_Layer_2_Configuration_Guide/b_Cisco_APIC_Layer_2_Configuration_Guide_chapter_010.pdf](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/L2_config/b_Cisco_APIC_Layer_2_Configuration_Guide/b_Cisco_APIC_Layer_2_Configuration_Guide_chapter_010.pdf)

upvoted 1 times

🗳️ 👤 **designated** 7 months, 2 weeks ago

Selected Answer: A

A is correct:

Select the option "Limit Local Learning To BD/EPD Subnet(s):

Limits IP address learning to the bridge domain subnets only. Every BD can have multiples subnets associated with it. By default, all IPs are learned.

upvoted 1 times

🗳️ 👤 **zelya19** 1 year, 3 months ago

Selected Answer: A

Limit IP Learning to Subnet - This option does not prevent remote IP endpoint learning.

Enforce Subnet Check - Limits both local and remote endpoint learning to instances only when the source IP address belongs to a bridge-domain subnet in the VRF instance

Note: Neither the built-in prevention mechanism for second-generation leaf switches nor the Enforce Subnet Check feature is available on first-generation leaf switches. Instead, you can configure the Limit IP Learning To Subnet option and the Disable Remote EP Learn option on the border leaf. Refer to the section discussing each feature to learn the differences between the features.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

upvoted 2 times

🗳️ 👤 **Mr_Certifiable** 1 year, 10 months ago

Selected Answer: A

Fabric operations, we recommend configurations that cause ACI to only learn IP addresses, which are configured on a BD Subnet

<https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/aci-endpoint-learning.pdf>

upvoted 1 times

🗳️ 👤 **Necha_uk** 2 years ago

A work friend passed using these questions about six weeks ago, he said to keep an eye out for new ones.

upvoted 1 times

🗳️ 👤 **eric0430** 2 years ago

Any new exam takers? Is this still valid? I heard that Cisco Data Center certification is getting some revisions.
upvoted 2 times

An engineer needs to deploy a leaf access port policy group in ACI Fabric to support the following requirements:

- ⇒ Control the amount of application data flowing into the system
- ⇒ Allow the newly connected device to auto-negotiate link speed with the leaf switch

Which two ACI policies must be configured to achieve these requirements? (Choose two.)

- A. link level policy
- B. L2 interface policy
- C. slow drain policy
- D. ingress data plane policing policy
- E. ingress control plane policing policy

Correct Answer: AD


Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_Cisco_ACI_and_Forward_Error_Correction.html

[https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401/b_Cisco_APIC_Security_Guide_chapter_01110.html#:~:text=Use%20data%20plane%20policing%20\(DPP,dropping%20of%20packets%20occurs%20immediately](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401/b_Cisco_APIC_Security_Guide_chapter_01110.html#:~:text=Use%20data%20plane%20policing%20(DPP,dropping%20of%20packets%20occurs%20immediately)

Community vote distribution

AD (100%)

 **pradeepmore** Highly Voted 4 years, 1 month ago

- A. link level policy
 - D. ingress data plane policing policy
- upvoted 9 times

 **designated** Most Recent 7 months, 2 weeks ago

Selected Answer: AD

A and D are correct.

Fabric > Access Policies > Policies > Link Level Policy

A.Link Level Policy (auto negotiation - Off / On / On - Enforce)

Fabric > Access Policies > Policies > Data Plane Policing Policy

B. Ingress data plane policing

upvoted 1 times

 **Mr_Certifiable** 1 year, 10 months ago

Selected Answer: AD

Link Level Policies

You can configure link level policies, which are a type of access policy.

Data Plane Policing

Use data plane policing (DPP) to manage bandwidth consumption on ACI. DPP policies can apply to egress traffic, ingress traffic, or both.

<https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/security-configuration/cisco-apic-security-configuration-guide-release-52x/data-plane-policing-52x.html#:~:text=Use%20data%20plane%20policing%20%28DPP%29%20to%20manage%20bandwidth,monitors%20the%20data%20rates%20for%20a%20part>

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/basic-configuration/Cisco-APIC-Basic-Configuration-Guide-42x/Cisco-APIC-B-Guide-42x_chapter_0101.html#:~:text=Procedure%201%20On%20the%20menu%20bar%2C%20choose%20Fabric,configuration.%205%20Click%20Submit%20.%20...

upvoted 2 times

 **eric0430** 2 years ago

Any new exam takers? Is this still valid? I heard that Cisco Data Center certification is getting some revisions.

upvoted 1 times

A customer migrates a legacy environment to Cisco ACI. A Layer 2 trunk is configured to interconnect the two environments. The customer also builds ACI fabric in an application-centric mode. Which feature should be enabled in the bridge domain to reduce instability during the migration?

- A. Set Multi-Destination Flooding to Flood in BD.
- B. Enable Flood in Encapsulation.
- C. Set Multi-Destination Flooding to Flood in Encapsulation.
- D. Disable Endpoint Dataplane Learning

Correct Answer: A

Community vote distribution



danhu Highly Voted 4 years, 5 months ago

I think the answer is C.

IMPORTANT In a single BD with Multiple EPGs, each EPG is defined by a different VLAN encapsulation. IF those VLANs extend into you legacy non-ACI network, then (by default) these frames will leak from one VLAN to another - which is probably NOT what you want, which is why there is a Multi-Destination Flooding option to Flood in Encapsulation - to prevent leakage from one EPG to another - but again, this may break some protocols. But before setting your BD to Flood in Encapsulation re-read the IF statement above.

upvoted 17 times

Mohitrsh84 Highly Voted 4 years, 4 months ago

I go with A.

in application centric we can have multiple EPGs in 1 BD having different vlan tags. However on legacy it will be one Vlan/BD.

<https://community.cisco.com/t5/application-centric/aci-bridge-domain-setting-discussion/td-p/3182528>

Encapsulation = Encap VLAN = VLAN tagged on the wire. Assuming you have 1 VLAN : 1 EPG, and you have multiple EPGs in 1 BD, then multi-destination packets (which fall into this setting's category) are only flooded within the EPG. Setting this to "Flood in BD" will flood traffic through out the BD.

upvoted 11 times

Marioalfo2 2 years, 7 months ago

I agreee, the key part is that the deployment is on network-centric mode

upvoted 2 times

Lutech Most Recent 2 months, 2 weeks ago

Selected Answer: A

A is the right answer because flood in BD mean that the Fabric send the packets to al subnet configured in the BD, if we use Flood in Encapsulation the Fabric send the packets to all port in the same vlan encap.

upvoted 1 times

2eb1ea8 5 months ago

Selected Answer: C

During migration, the legacy environment and ACI fabric are interconnected via a Layer 2 trunk. If multi-destination traffic is not controlled, it could flood into the legacy environment, causing network instability. By limiting flooding to the ACI fabric's encapsulation, you ensure that the legacy environment is not overwhelmed with unnecessary traffic.

upvoted 2 times

prospio971 5 months, 1 week ago

Selected Answer: C

Flood in Encapsulation: This feature allows unknown unicast, broadcast, and multicast traffic to be flooded within the specified VLAN encapsulation. This behavior is essential during migrations to ensure that any unknown traffic from the legacy environment is properly handled without causing disruptions or instability.

upvoted 1 times

Edjane 7 months, 3 weeks ago

I think that you need to migration everything to ACI after you create the AP Centric, because. answer is A

upvoted 1 times

🗨️ 👤 **Rollizo** 8 months, 2 weeks ago

Selected Answer: A

This functionality was introduced in the Cisco APIC 3.1 releases primarily for service graph policy-based redirect (PBR) deployments, and it has been superseded by the ability to disable IP dataplane learning per-VRF instance (Cisco APIC release 4.0). We do not recommend disabling IP learning per bridge domain and it is not supported except when used with PBR.

Then it is no D

upvoted 1 times

🗨️ 👤 **mdriraa** 10 months, 3 weeks ago

Selected Answer: A

<https://community.cisco.com/t5/application-centric-infrastructure/aci-bridge-domain-setting-discussion/td-p/3182528>

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: A

<https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2020/pdf/BRKACI-3545.pdf>

- Flood in BD

Flood within the same BD
regardless of EPG or VLAN.

- Flood in Encapsulation

Flood within the same
access encap VLAN and BD
regardless of EPG.

upvoted 2 times

🗨️ 👤 **zelya19** 1 year, 3 months ago

Selected Answer: C

Recommended Solution

The flood in encapsulation option is used to limit flooding traffic inside the bridge domain to a single encapsulation. When EPG1/VLAN X and EPG2/VLAN Y share the same bridge domain and flood in encapsulation is enabled, the encapsulation flooding traffic does not reach the other EPG/VLAN.

<https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/6x/l2-configuration/cisco-apic-layer-2-networking-configuration-guide-60x/bridging-60x.html>

upvoted 2 times

🗨️ 👤 **msalamehi** 1 year, 5 months ago

I'll go with B

The feature lets you scope the flooding domain to the individual VLANs on which the traffic is received. This is roughly equivalent to scoping the flooding to the EPGs.

https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-application-centric-infrastructure-design-guide.html#_Toc129073543

upvoted 1 times

🗨️ 👤 **sailorsoul** 1 year, 5 months ago

Selected Answer: C

C as everybody says.

Network centric mode: 1 EPG --> 1 vlan.

Application centric mode : multiple vlan in 1 EPG.

upvoted 2 times

🗨️ 👤 **Hubertheman** 1 year, 6 months ago

Selected Answer: A

A. Set Multi-Destination Flooding to Flood in BD (Bridge Domain): Enabling multi-destination flooding in the bridge domain ensures that broadcast, unknown unicast, and multicast (BUM) traffic is flooded within the bridge domain. This setting can help maintain visibility of legacy systems during the migration phase, as it ensures that BUM traffic can still flow across the newly integrated ACI infrastructure. This option can be particularly useful during migrations to ensure that devices in the legacy network can still communicate with devices that have been moved to the ACI fabric.



upvoted 4 times

🗨️ 👤 **RUT** 1 year, 6 months ago

Selected Answer: B

The correct answer is B



upvoted 1 times

  **hebdeb** 1 year, 11 months ago

Selected Answer: C

Multi-Destination Flooding to Flood in Encapsulation


upvoted 1 times

  **hebdeb** 1 year, 11 months ago

Selected Answer: D

Application Centric means many VLAN in the same BD so Multi-Destination Flooding to Flood in Encapsulation

upvoted 2 times

  **eric0430** 2 years ago

Any new exam takers? Is this still valid? I heard that Cisco Data Center certification is getting some revisions.

upvoted 1 times



New ESXi hosts are procured in a data center compute expansion project. An engineer must update the configuration on the Cisco APIC controllers to support the addition of the new servers to the existing VMM domain. Which action should be taken to support this change?

- A. Create a range of internal VLANs in the associated VLAN pool.
- B. Set the encapsulation mode as VXLAN.
- C. Enable infrastructure VLAN in the associated AEP.
- D. Map the leaf interface selector to the AEP that is associated with the VMM domain.

Correct Answer: D

Community vote distribution

D (100%)

  **manet** Highly Voted 4 years ago

I think that D is correct
upvoted 6 times

  **2eb1ea8** Most Recent 5 months ago

Selected Answer: A

When adding new ESXi hosts to a VMM domain, the ACI fabric needs additional VLANs to handle the traffic for the new hosts. These VLANs are used internally by the ACI fabric for communication between the leaf switches and the ESXi hosts. By creating a range of internal VLANs in the VLAN pool associated with the VMM domain, you ensure that the fabric has enough VLANs to support the new hosts.
upvoted 1 times

  **designated** 7 months, 2 weeks ago



Selected Answer: D

D is correct since we are just adding a new ESXi host
upvoted 1 times

  **Mr_Certifiable** 1 year, 10 months ago

Selected Answer: D

<https://aci-lab.ciscolive.com/lab/pod9/acivmm/add-hosts-dvs>
upvoted 2 times

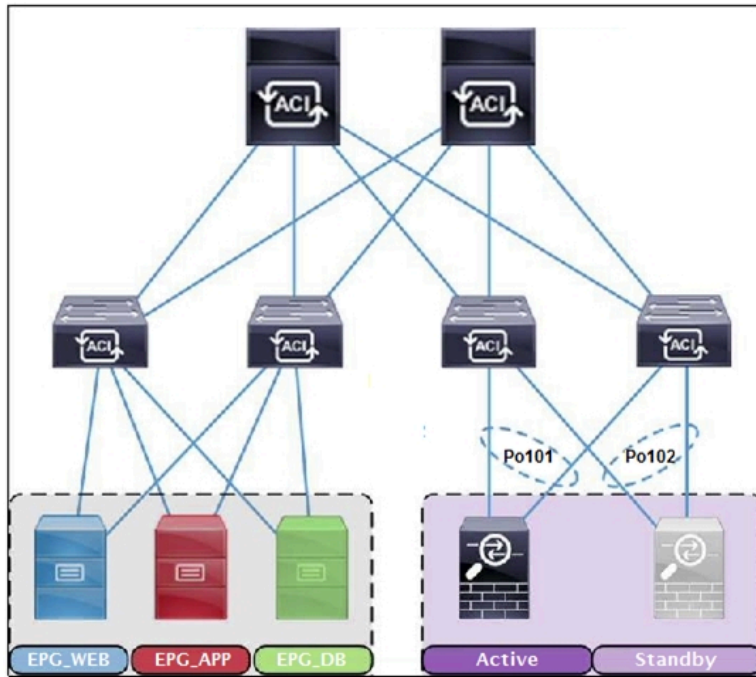
  **eric0430** 2 years ago

Any new exam takers? Is this still valid? I heard that Cisco Data Center certification is getting some revisions.
upvoted 1 times

  **igebuntu** 3 years, 4 months ago

More correctly: Apply an Interface Policy Group - which has the AAEP with the existing VMM domain configured - to the Interface Selectors.
upvoted 4 times

DRAG DROP -



Refer to the exhibit. A Cisco ACI fabric is newly deployed, and the security team requires more visibility of all inter EPG traffic flows. All traffic in a VRF must be forwarded to an existing firewall pair. During failover, the standby firewall must continue to use the same IP and MAC as the primary firewall. Drag and drop the steps from the left into the implementation order on the right to configure the service graph that meets the requirements. (Not all steps are used.)

Select and Place:

Answer Area

Apply a service graph template and select vzAny EPG as the consumer and provider.	Step 1
Select a redirect policy with the Layer 3 destination.	Step 2
Create a Layer 4 to Layer 7 service graph template.	Step 3
Select a redirect policy with enabled anycast and the Layer 3 destination.	Step 4
Select the same cluster interface under Consumer Connector and Provider Connector.	Step 5
Create a service bridge domain and a Layer 4 to Layer 7 device with one cluster interface.	Step 6
Select the existing contract with custom IP EtherType filter.	

Correct Answer:

Answer Area

Apply a service graph template and select vzAny EPG as the consumer and provider.

Create a service bridge domain and a Layer 4 to Layer 7 device with one cluster interface.

Select a redirect policy with the Layer 3 destination.

Create a Layer 4 to Layer 7 service graph template.

Create a Layer 4 to Layer 7 service graph template.

Select a redirect policy with enabled anycast and the Layer 3 destination.

Select a redirect policy with enabled anycast and the Layer 3 destination.

Select the existing contract with custom IP EtherType filter.

Select the same cluster interface under Consumer Connector and Provider Connector.

Select the same cluster interface under Consumer Connector and Provider Connector.


Create a service bridge domain and a Layer 4 to Layer 7 device with one cluster interface.

Apply a service graph template and select vzAny EPG as the consumer and provider.

Select the existing contract with custom IP EtherType filter.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/3-x/L4-L7_services_deployment/3_2_1/b_L4L7_Deploy_321/b_L4L7_Deploy_321_chapter_01001.html

 **[Removed]**  2 years, 11 months ago

To configure Service Graph in managed or unmanaged mode , Configuration steps should be as follows :

- 1.Create a service bridge domain and a Layer 4 to Layer 7 device with on cluster interface.
- 2.Create a Layer 4 to Layer 7 service graph template .
- 3.Select a redirect policy with enabled any cast and the Layer 3 destination .
- 4.Apply a service a graph template and select vzAny EPG as the consumer and provider
- 5.Select the existing contract with customer IP Ether Type filter .
- 6.Select the same cluster interface under Consumer Connector and Provider connector .

upvoted 9 times

 **onix** 1 year, 11 months ago


https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/cli/nx/cfg/b_APIC_NXOS_CLI_User_Guide/b_APIC_NXOS_CLI_User_Guide_chapter_011001.pdf

Anycast services are not supported with the following features and options:

- Two firewalls in an Active/Standby relationship (in this scenario, the Anycast service is active in only one pod and all traffic is sent using the active service)

So 3 should be: Select a redirect policy with the Layer 3 destination.

upvoted 4 times

 **nikomski** 2 years, 11 months ago

- 1.Create a service bridge domain and a Layer 4 to Layer 7 device with on cluster interface.
- 2.Create a Layer 4 to Layer 7 service graph template .
- 3.Apply a service a graph template and select vzAny EPG as the consumer and provider
- 4.Select the existing contract with customer IP Ether Type filter .
- 5.Select a redirect policy with enabled any cast and the Layer 3 destination .
- 6.Select the same cluster interface under Consumer Connector and Provider connector .

upvoted 13 times

 **Jey10** 2 years, 7 months ago

It is OK except 5 => it should be without anycast

upvoted 6 times

  **Said75** Highly Voted 7 months ago

Correct Answer and verified on my lab :



1. Create a service bridge domain and a Layer 4 to Layer 7 device with on cluster interface.
2. Create a Layer 4 to Layer 7 service graph template .
3. Apply a service a graph template and select vzAny EPG as the consumer and provider
4. Select the existing contract with customer IP Ether Type filter .
5. Select a redirect policy with the Layer 3 destination .
6. Select the same cluster interface under Consumer Connector and Provider connector .

upvoted 7 times

  **2eb1ea8** Most Recent 5 months ago

1. Create a service bridge domain and a Layer 4 to Layer 7 device with on cluster interface.
2. Create a Layer 4 to Layer 7 service graph template.
3. Apply a service a graph template and select vzAny EPG as the consumer and provider.
4. Select the existing contract with customer IP EtherType filter.
5. Select a redirect policy with enabled anycast and the layer 3 destination (Enabling anycast ensures that the standby firewall can use the same IP and MAC address as the primary firewall during failover. The layer 3 destination specifies the firewall's IP address).
6. Select the same cluster interface under Consumer Connector and Provider connector.

upvoted 1 times

  **korthab** 1 year, 4 months ago

I think this is the correct answer based on the steps i watched on labminutes.com.

1. Create a service bridge domain and a Layer 4 to Layer 7 device with on cluster interface.
2. Create a Layer 4 to Layer 7 service graph template.
3. Apply a service a graph template and select vzAny EPG as the consumer and provider.
4. Select the existing contract with customer IP EtherType filter.
5. Select a redirect policy with Layer 3 destination.
6. Select the same cluster interface under Consumer Connector and Provider connector.

Labminutes LINK:

https://www.labminutes.com/dc0032_aci_service_graph_pbr_fw_1

Anycast LINK:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/cli/nx/cfg/b_APIC_NXOS_CLI_User_Guide/b_APIC_NXOS_CLI_User_Guide_chapter_011001.pdf

Quote:

" Anycast services are not supported with the following features and options:

- Two firewalls in an Active/Standby relationship (in this scenario, the Anycast service is active in only one pod and all traffic is sent using the active service) "

upvoted 5 times

  **ciscoaci2022** 1 year, 10 months ago

The correct answer should be:

1. Create a service bridge domain and a Layer 4 to Layer 7 device with on cluster interface.
2. Create a Layer 4 to Layer 7 service graph template .
3. Apply a service a graph template and select vzAny EPG as the consumer and provider
4. Select the existing contract with customer IP Ether Type filter .
5. Select a redirect policy with the Layer 3 destination .
6. Select the same cluster interface under Consumer Connector and Provider connector

since the Anycast services are not supported with the following features and options:

- Two firewalls in an Active/Standby relationship (in this scenario, the Anycast service is active in only one pod and all traffic is sent using the active service)

So 3 should be: Select a redirect policy with the Layer 3 destination.

upvoted 4 times

🗨️ 👤 **muhnator** 1 year, 10 months ago

1. Create a service bridge domain and a Layer 4 to Layer 7 device with on cluster interface.
2. Create a Layer 4 to Layer 7 service graph template .
3. Apply a service a graph template and select vzAny EPG as the consumer and provider
4. Select the existing contract with customer IP Ether Type filter .
5. Select a redirect policy with the Layer 3 destination .
6. Select the same cluster interface under Consumer Connector and Provider connector .

upvoted 4 times

🗨️ 👤 **nabilzay** 2 years, 10 months ago

I think nikmski's answer is right based on this doc:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/L4-L7_Services_Deployment/guide/b_L4L7_Deploy_ver201/b_L4L7_Deploy_ver201_chapter_010100.html#id_27316

However not sure if anycast is needed, I'd say no

upvoted 2 times

🗨️ 👤 **Kalpesh** 2 years, 10 months ago

I think anycast is also not needed as it's a active/standby setup not Active/Active.

upvoted 3 times

🗨️ 👤 **thiyagas** 2 years, 11 months ago

not sure if this answer is correct... any comment.?

upvoted 2 times

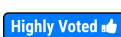
An engineer is extending an EPG out of the ACI fabric using static path binding. Which statement about the endpoints is true?

- A. Endpoints must connect directly to the ACI leaf port.
- B. External endpoints are in a different bridge domain than the endpoints in the fabric.
- C. Endpoint learning encompasses the MAC address only.
- D. External endpoints are in the same EPG as the directly attached endpoints.


Correct Answer: D

Community vote distribution

D (100%)

  **[Removed]**  2 years, 5 months ago

Correct Answer : D .Extending the EPG should be statically binding with same EPG in ACI.
upvoted 6 times


  **hybersat**  2 years, 5 months ago

Should this not be D?
upvoted 5 times


  **Supreme_123**  7 months, 3 weeks ago


Correct answer D
upvoted 1 times


  **SherifAbdelMoteleb** 1 year, 1 month ago



 D is correct
upvoted 2 times

  **patopizarro** 1 year, 2 months ago

 D is correct
upvoted 2 times

  **KSM03** 1 year, 7 months ago

 The user can extend an EPG beyond an ACI leaf by statically assigning a leaf port (along with a VLAN ID) to an EPG. After doing so, all the traffic received on this leaf port with the configured VLAN ID will be mapped to the EPG and the configured policy for this EPG will be enforced. The endpoints need not be directly connected to the ACI leaf port. They can be behind a layer 2 network as long as the VLAN associated with the EPG is enabled within the layer 2 network that connects the remote endpoint to the ACI fabric.
upvoted 5 times

  **mr_siro** 1 year, 10 months ago


D is correct
upvoted 1 times

  **thomyohan** 2 years ago

sorry D is right !
upvoted 2 times

  **thomyohan** 2 years ago

The user can extend an EPG beyond an ACI leaf by statically assigning a leaf port (along with a VLAN ID) to an EPG. After doing so, all the traffic received on this leaf port with the configured VLAN ID will be mapped to the EPG and the configured policy for this EPG will be enforced. The endpoints need not be directly connected to the ACI leaf port. They can be behind a layer 2 network as long as the VLAN associated with the EPG is enabled within the layer 2 network that connects the remote endpoint to the ACI fabric.
(c) is the right answer
upvoted 1 times

  **Skappie** 2 years, 1 month ago

I think answer should be D.

Question is specific: engineer is extending an EPG out of the ACI fabric. This eliminates A.

C is also incorrect, if EPG is mapped to a BD with a subnet, endpoints will be learn with IPs. Not enough information for C.

D is correct. epg static binding to leaf interface <-> external switch <-> endpoint connected to external switch = those external endpoints will be in the same EPG as directly connected endpoints

upvoted 5 times

  **jithin1234** 2 years, 3 months ago



why not a? in option D, there is external endpoint

upvoted 1 times

  **Skappie** 2 years, 1 month ago

Question is specific: engineer is extending an EPG out of the ACI fabric. This eliminates A.

upvoted 1 times

  **nabilzay** 2 years, 4 months ago

It should be D

upvoted 4 times

  **thiyagas** 2 years, 5 months ago

Correct Answer is D

upvoted 4 times

Which setting prevents the learning of Endpoint IP addresses whose subnet does not match the bridge domain subnet?

- A. "Limit IP learning to network" setting within the bridge domain.
- B. "Limit IP learning to subnet" setting within the EPG.
- C. "Limit IP learning to network" setting within the EPG.
- D. "Limit IP learning to subnet" setting within the bridge domain.

Correct Answer: D

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/L2_config/b_Cisco_APIC_Layer_2_Configuration_Guide/b_Cisco_APIC_Layer_2_Configuration_Guide_chapter_010.html

Community vote distribution

D (100%)

🗲️ 👤 **Ksinghb** Highly Voted 👍 3 years, 10 months ago

"Limit IP learning to subnet" setting within the bridge domain.
upvoted 6 times

🗲️ 👤 **designated** Most Recent 🕒 7 months, 2 weeks ago

Selected Answer: D

"Limit Local IP Learning to BD/EPG Subnet(s)" checkbox in 6.0 > correct!
upvoted 1 times

🗲️ 👤 **zelya19** 1 year, 3 months ago

Selected Answer: D

"Limit Local IP Learning to BD/EPG Subnet(s)" checkbox in 6.0
upvoted 2 times

🗲️ 👤 **Huberttheman** 1 year, 11 months ago

Selected Answer: D

D. "Limit IP learning to subnet" setting within the bridge domain.
upvoted 1 times

Which endpoint learning operation is completed on the egress leaf switch when traffic is received from an L3Out?

- A. The source MAC and IP address of the traffic is learned as a local endpoint.
- B. The source MAC address of the traffic is learned as a remote endpoint.
- C. No source MAC or IP address of the traffic is learned as a remote endpoint.
- D. The source IP address of the traffic is learned as a remote endpoint.

Correct Answer: C

Community vote distribution

C (60%)

B (40%)

TCoder Highly Voted 4 years, 4 months ago

C is the right answer. For the Leaf connected to L3out, Source MAC is learnt only while other leafs do not learn source MAC or IP when they receive traffic from L3out.

upvoted 15 times

anasham Highly Voted 3 years, 6 months ago

Correct answer is C.

ACI uses concept like traditional networking for L3 out connectivity.

IT uses RIB and ARP table, not local endpoint table.

IT doesn't learn IP address from data plane; it learns the subnet via control plane using routing protocol

It learns MAC address of next-hop(router) for this subnet from data plane.

As the mac address is the MAC of the router (not that of the endpoint) and it doesn't learn endpoint IP address(/32 or /128), leaf doesn't learn MAC or IP of the endpoint when it receives from L3out

upvoted 10 times

ashwind123 Most Recent 3 months, 2 weeks ago

Selected Answer: C

C is right answer.

Local endpoint learning with an incoming packet from L3Out to Cisco ACI:

Only the source MAC address is learned as a local endpoint. The source IP address is not learned as a local endpoint.

Remote endpoint learning with an incoming packet from L3Out to Cisco ACI:

No source MAC or IP address is learned as a new remote endpoint by a packet.*

upvoted 2 times

designated 7 months, 2 weeks ago

Selected Answer: C

C is correct.

L3Out is a way to communicate ACI infrastructure using endpoint learning with the traditional network which uses flood, arp, mac, IP and all of them as DESTINATION and not SOURCE as ACI.

upvoted 1 times

zelya19 1 year, 3 months ago

Selected Answer: C

No source MAC or IP address is learned as a new remote endpoint by a packet.*

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html#Silenthostsconsiderations>

upvoted 1 times

sailorsoul 1 year, 5 months ago

Selected Answer: C

C is correct. Only source MAC is learned.

upvoted 1 times

🗳️ 👤 **Huberttheman** 1 year, 6 months ago

Selected Answer: C

Routing Over Learning: For traffic from L3Outs like the Internet, ACI relies more on routing protocols and policy-based forwarding rather than endpoint learning. The fabric will use its routing table to forward traffic to the appropriate destinations, based on the routes learned through dynamic routing protocols or static routing configurations.

upvoted 1 times

🗳️ 👤 **Macc10** 1 year, 8 months ago

Selected Answer: B

Cisco ACI uses a behavior similar to that in traditional networks for L3Out connectivity. The Cisco ACI L3Out domain learns the MAC address only from the data plane. IP addresses are not learned from the data plane in an L3Out domain; instead, Cisco ACI uses ARP to resolve next-hop IP and MAC relationships to reach the prefixes behind external routers.

upvoted 1 times

🗳️ 👤 **Mr_Certifiable** 1 year, 10 months ago

Selected Answer: B

assume local EP

Local endpoint learning with an incoming packet from L3Out to Cisco ACI:

Only the source MAC address is learned as a local endpoint. The source IP address is not learned as a local endpoint.

Remote endpoint learning with an incoming packet from L3Out to Cisco ACI:

No source MAC or IP address is learned as a new remote endpoint by a packet

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

upvoted 1 times

🗳️ 👤 **asd248402** 2 years ago

Selected Answer: C

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html#L3Outendpointlearningconsiderations>

see figure 7 wich explain it very well

upvoted 1 times

🗳️ 👤 **Mr_Certifiable** 2 years ago

C is my vote - based on snip [The Cisco ACI L3Out domain learns the MAC address only from the data plane. IP addresses are not learned from the data plane in an L3Out domain; instead, Cisco ACI uses ARP to resolve next-hop IP and MAC relationships to reach the prefixes behind external routers.]

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

upvoted 1 times

🗳️ 👤 **mdsuresh** 2 years, 9 months ago

B is the right answer. For C, without source mac and IP, how the return traffic will flow?

upvoted 1 times

🗳️ 👤 **alphatest** 3 years ago

Selected Answer: C

In this document

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html#L3Outendpointlearningconsiderations>

it is made clear that no remote endpoint is learned by the egress leaf (Leaf 1 in figure 7). The stale endpoint issue is actually due to this fact.

upvoted 2 times

🗳️ 👤 **netguru** 3 years, 1 month ago

B is definitely not correct. As per Cisco White paper-

Cisco ACI learns a MAC or IP address as a remote endpoint when a packet comes into a Cisco ACI leaf switch from another leaf switch through a

spine switch.

So remote endpoints are those which are connected to other leaf. I think correct answer is A.

upvoted 1 times

  **patopizarro** 3 years, 2 months ago



Selected Answer: B

B is the right answer.

The Cisco ACI L3Out domain learns the MAC address only from the data plane. IP addresses are not learned from the data plane in an L3Out domain; instead, Cisco ACI uses ARP to resolve next-hop IP and MAC relationships to reach the prefixes behind external routers.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

upvoted 2 times



  **wcorrea** 3 years, 11 months ago

I think D is right. The question say on Egress switch... Egress is LEAF 2 in fig5.

"In figure 5..... only IP address 192.168.1.1 is learned as a remote endpoint on LEAF2."

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

upvoted 1 times

  **Ruslans** 3 years, 11 months ago

ocal endpoint learning with an incoming packet from L3Out to Cisco ACI:

C, Only the source MAC address is learned as a local endpoint. The source IP address is not learned as a local endpoint.

upvoted 1 times

```
<fvTenant name="ACILab">
  <fvCtx name="pvn1"/>
  <fvBD name="bd1">
    <fvRsCtx tnFvCtxName="pvn1"/>
    <fvSubnet ip="10.1.100.1/24"/>
  </fvBD>
</fvTenant>
```

Refer to the exhibit. Which two objects are created as a result of the configuration? (Choose two.)

- A. application profile
- B. attachable AEP
- C. bridge domain
- D. endpoint group
- E. VRF

Correct Answer: CE

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/rest_cfg/2_1_x/b_Cisco_APIC_REST_API_Configuration_Guide/b_Cisco_APIC_REST_API_Configuration_Guide_chapter_01110.html


 **designated** 7 months, 2 weeks ago

Selected Answer: CE

fvCtx = VRF

fvBD = Bridge Domain

upvoted 1 times

 **Jey10** 4 years, 1 month ago

And also subnet

upvoted 1 times

 **hybersat** 4 years, 5 months ago

Correct, here is an example from the link in the answer..

Configure the tenant, VRF, and bridge domain.

This example configures tenant t1 with VRF v1 and bridge domain bd1. The tenant, VRF, and BD are not yet deployed.

Example:

```
<fvTenant name="t1">
  <fvCtx name="v1"/>
  <fvBD name="bd1">
    <fvRsCtx tnFvCtxName="v1"/>
    <fvSubnet ip="44.44.44.1/24" scope="public"/>
    <fvRsBDToOut tnL3extOutName="l3out1"/>
  </fvBD>
</fvTenant>
```

upvoted 4 times

What must be enabled in the bridge domain to have the endpoint table learn the IP addresses of endpoints?

- A. L2 unknown unicast: flood
- B. GARP based detection
- C. unicast routing
- D. subnet scope

Correct Answer: C

Reference:

<https://hsvglobalschool.in/dhkycw/cisco-aci-bridge-domain.html>

Community vote distribution

C (100%)

🗲️ 👤 **designated** 7 months, 2 weeks ago

Selected Answer: C

IP = Unicast Routing

upvoted 1 times

🗲️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: C

Unicast Routing is correct

upvoted 1 times

🗲️ 👤 **DSAM9** 3 years, 4 months ago

C. Unicast Correct

upvoted 3 times

An engineer is extending EPG connectivity to an external network. The external network houses the Layer 3 gateway and other end hosts. Which ACI bridge domain configuration should be used?

- A. Forwarding: Custom L2 Unknown Unicast: Hardware Proxy L3 Unknown Multicast Flooding: Flood Multi Destination Flooding: Flood in BD ARP Flooding: Enabled
- B. Forwarding: Custom L2 Unknown Unicast: Flood L3 Unknown Multicast Flooding: Flood Multi Destination Flooding: Flood in BD ARP Flooding: Enabled
- C. Forwarding: Custom L2 Unknown Unicast: Hardware Proxy L3 Unknown Multicast Flooding: Flood Multi Destination Flooding: Flood in BD ARP Flooding: Disabled
- D. Forwarding: Custom L2 Unknown Unicast: Flood L3 Unknown Multicast Flooding: Flood Multi Destination Flooding: Flood in BD ARP Flooding: Disabled

Correct Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **[Removed]** Highly Voted 4 years, 5 months ago

Correct Answer : B.

<https://aci-lab.ciscolive.com/lab/pod9/tenants/create-l2-bd>

upvoted 16 times

🗳️ 👤 **nabilzay** Highly Voted 4 years, 4 months ago

Should be B, you need ARP flood to be able to communicate correctly with the external L2 domain

upvoted 7 times

🗳️ 👤 **designated** Most Recent 7 months, 2 weeks ago

Selected Answer: B

Due to this specific configurations are required to achieve the goal of connecting the two separate domains. There are two unique characteristics that have to be configured in a bridge domain so that it will work properly.

ARP Flooding: This parameters forces the ACI fabric to flood to all ports in the same Layer2 domain any ARP requests that are received. In this way any broadcast ARP requests that arrive from an external Layer 2 are forwarded to all endpoints in the same layer 2 domain or to the exterior from inside the ACI fabric

Unknown unicast forwarding: This parameter also changes the behavior of the ACI fabric to flood any unknown unicast entries to the ports of the same layer 2 domain.

upvoted 1 times

🗳️ 👤 **sailorsoul** 1 year, 5 months ago

Selected Answer: B

B is correct.

upvoted 1 times

🗳️ 👤 **Gab99** 2 years, 8 months ago

Selected Answer: B

i feel B also

upvoted 1 times

🗳️ 👤 **patopizarro** 3 years, 2 months ago

Selected Answer: B

Correct Answer : B. Flood in BD ARP Flooding: Enabled

upvoted 5 times

🗳️ 👤 **polle** 3 years, 7 months ago

Should be B

upvoted 2 times

🗳️ 👤 **thiyagas** 4 years, 5 months ago

shouldn't this be B..?

upvoted 6 times

An engineer configured a bridge domain with the hardware-proxy option for Layer 2 unknown unicast traffic. Which statement is true about this configuration?

- A. The leaf switch drops the Layer 2 unknown unicast packet if it is unable to find the MAC address in the local forwarding tables.
- B. The Layer 2 unknown hardware proxy lacks support of the topology change notification.
- C. The leaf switch forwards the Layers 2 unknown unicast packets to all other leaf switches if it is unable to find the MAC address in its local forwarding tables.
- D. The spine switch drops the Layer 2 unknown unicast packet if it is unable to find the MAC address in the proxy database.

Correct Answer: D

Community vote distribution

D (100%)

🗲️ 👤 **[Removed]** Highly Voted 👍 2 years, 5 months ago

Correct Answer : D. Spine will take the decision to forward the traffic for unknown unicast if Spine proxy is being selected .
upvoted 18 times

🗲️ 👤 **KSM03** Highly Voted 👍 1 year, 7 months ago

Selected Answer: D

When using hardware-proxy, you should consider enabling unicast routing and defining a subnet on the bridge domain. This is because with hardware-proxy on, if a MAC address has been aged out in the spine switch-proxy, traffic destined to this MAC address is dropped.
Reference: <https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-application-centric-infrastructure-design-guide.html#Usinghardwareproxytoreduceflooding>
upvoted 6 times

🗲️ 👤 **Gab99** Most Recent 🕒 8 months, 3 weeks ago

Selected Answer: D

of course D no question
upvoted 1 times

🗲️ 👤 **mr_siro** 1 year, 10 months ago

D is correct
upvoted 3 times

🗲️ 👤 **jim13c** 1 year, 11 months ago

D is correct. <https://community.cisco.com/t5/application-centric/unknown-unicast-in-aci-fabric/td-p/2569161>
upvoted 2 times

🗲️ 👤 **thiyagas** 2 years, 5 months ago

Correct answer is D
upvoted 5 times

An engineer configured Layer 2 extension from the ACI fabric and changed the Layer 2 unknown unicast policy from Flood to Hardware Proxy. How does this change affect the flooding of the L2 unknown unicast traffic?

- A. It is forwarded to one of the spines to perform as a spine proxy.
- B. It is flooded within the whole fabric.
- C. It is dropped by the leaf when the destination endpoint is not present in the endpoint table.
- D. It is forwarded to one of the APICs to perform as a proxy.

Correct Answer: A

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/L2_config/b_Cisco_APIC_Layer_2_Configuration_Guide/b_Cisco_APIC_Layer_2_Configuration_Guide_chapter_010.html

  **designated** 7 months, 2 weeks ago

A is correct.

Hardware proxy for Layer 2 unknown unicast traffic is the default option. If the destination MAC is not in the ingress leaf endpoint table, the packet is sent to the spine proxy. This forwarding behavior uses the COOP database on spine switches to forward unknown unicast traffic to the destination leaf without relying on flood-and-learn behavior, as long as the MAC address is known to the COOP database on spine switch.

If MAC address of an endpoint is unknown to spine COOP database (for example, silent host), the spine drops the packet.

upvoted 2 times

  **marceaubueno** 1 year, 6 months ago

A is correct. Hardware proxy instructs to the leaf to send the unknown unicast packet to the spine which will check the mac address of the endpoint in its coop table.

upvoted 3 times

Create Subnet

Gateway IP:
address/mask

Treat as virtual IP address: ☐

Make this IP address primary: ☐

Scope: ☒ Private to VRF
☐ Advertised Externally
☐ Shared between VRFs

Description:

Subnet Control: ☐ No Default SVI Gateway
☐ Querier IP

L3 Out for Route Profile:

Route Profile:

ND RA Prefix policy:

Cancel Submit

When the subnet is configured on a bridge domain, on which physical devices is the gateway IP address configured?

- A. all leaf switches and all spine nodes
- B. only spine switches where the bridge domain of the tenant is present
- C. only leaf switches where the bridge domain of the tenant is present
- D. all border leaf nodes where the bridge domain of the tenant is present

Correct Answer: C

Reference:

<http://www.netdesignarena.com/index.php/2016/06/16/aci-tenant-building-blocks-forwarding-logic/>

nikomski Highly Voted 4 years, 5 months ago

C is correct

upvoted 12 times

[Removed] Highly Voted 4 years, 5 months ago

Correct Answer: Gateway will be available only the leaf switch where bridge domain is present in a specific tenant .

upvoted 10 times

designated Most Recent 7 months, 2 weeks ago

Selected Answer: C

Complete answer:

Also, just like a VLAN and its SVI on a normal switch, the ACI bridge domain provides a default gateway and its subnet for endpoints. ACI bridge domain and the gateway will be pervasively deployed on any leaf switches with EPGs associated to the BD. Hence, a default gateway for all endpoints will be right in front of them on the first leaf they will see without having to forward a packet to another leaf that has the default gateway for the endpoint. This gateway is called a pervasive gateway, or anycast gateway. One bridge domain can have multiple pervasive gateways just like a VLAN SVI with a secondary IP on a normal switch.

upvoted 1 times

designated 7 months, 2 weeks ago

Selected Answer: C

C is correct

ACI bridge domain and the gateway will be pervasively deployed on any leaf switches with EPGs associated to the BD.

upvoted 1 times

Which method does the Cisco ACI fabric use to load-balance multidestination traffic?

- A. forwarding tag trees
- B. PIM routing
- C. spanning trees
- D. shortest-path trees

Correct Answer: A

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_010010.html

Community vote distribution

A (100%)

 **mr_siro** Highly Voted 3 years, 10 months ago

A is corect

The predefined topologies based on which ACI forwards multi-destination traffic are called forwarding tag (FTag) trees. Each FTag tree does not necessarily use all fabric uplinks. That is why ACI creates multiple FTag trees and load balances multi-destination traffic across them. All switches in a fabric understand based on the FTag bits in the GIPo address how to forward the traffic they receive further along the specified FTag tree. Four bits are used to identify FTag IDs; ACI fabrics support up to 12 FTag trees.

upvoted 6 times

 **designated** Most Recent 7 months, 2 weeks ago

Selected Answer: A

A is correct

The ACI fabric uses Forwarding Tag (FTAG) trees to load balance multi-destination traffic. All multi-destination traffic is forwarded in the form of encapsulated IP multicast traffic within the fabric. The ingress leaf assigns an FTAG to the traffic when forwarding it to the spine. The FTAG is assigned in the packet as part of the destination multicast address. In the fabric, the traffic is forwarded along the specified FTAG tree. Spine and any intermediate leaf switches forward traffic based on the FTAG ID. One forwarding tree is built per FTAG ID. Between any two nodes, only one link forwards per FTAG. Because of the use of multiple FTAGs, parallel links can be used with each FTAG choosing a different link for forwarding. The larger the number of FTAG trees in the fabric means the better the load balancing potential is. The ACI fabric supports up to 12 FTAGs.

upvoted 1 times

 **Mr_Certifiable** 1 year, 10 months ago

Selected Answer: A

The ACI fabric uses Forwarding Tag (FTAG) trees to load balance multi-destination traffic. All multi-destination traffic is forwarded in the form of encapsulated IP multicast traffic within the fabric. The ingress leaf assigns an FTAG to the traffic when forwarding it to the spine. The FTAG is assigned in the packet as part of the destination multicast address. In the fabric, the traffic is forwarded along the specified FTAG tree. Spine and any intermediate leaf switches forward traffic based on the FTAG ID

<https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/aci-fundamentals/cisco-aci-fundamentals-52x/fundamentals-52x.html>

upvoted 1 times

 **sourabh1000** 3 years, 10 months ago

yeah A is right

upvoted 1 times

What happens to the traffic flow when the Cisco ACI fabric has a stale endpoint entry for the destination endpoint?

- A. The leaf switch does not learn the source endpoint through data plane learning.
- B. The leaf switch drops the traffic that is destined to the endpoint.
- C. The leaf switch floods the traffic to the endpoint throughout the fabric.
- D. The leaf switch sends the traffic to the wrong destination leaf.

Correct Answer: B

Reference:

<https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKACI-2641.pdf>

Community vote distribution

D (100%)

  **Carol254**  4 years, 4 months ago


I think the answer is D.

Because of this stale remote endpoint, any traffic from LEAF1 toward IP2 will fail, because LEAF1 sends packets to the wrong leaf.

This stale remote endpoint on LEAF1 needs to be manually cleared to resume communication. The command syntax to manually clear a particular remote IP endpoint is shown here:

```
LEAF1# clear system internal epm endpoint key vrf <vrf-name> ip <ip-address>
```

upvoted 15 times

  **ashwind123**  3 months, 2 weeks ago

Selected Answer: D

I've seen couple of times wrong or stale entry on ACI which results traffic send to wrong destination leaf resulting in traffic blackholing.

upvoted 1 times

  **designated** 7 months, 2 weeks ago

Selected Answer: D

D is correct:

In Cisco ACI, when a stale endpoint entry exists for a destination endpoint, the fabric incorrectly assumes the stale information is valid. This results in the traffic being forwarded to the wrong destination leaf (where the endpoint was previously located). This can cause communication issues until the endpoint table is updated with the correct information through control plane updates or data plane learning.

upvoted 1 times

  **Mr_Certifiable** 2 years ago



as stated - context would indicate the switch learns it is - stale - therefore it is dropped.

upvoted 1 times

  **Supreme_123** 2 years, 7 months ago

I think is D

upvoted 1 times

  **Gab99** 2 years, 8 months ago

Selected Answer: D

stale means stuck not there endpoint packet lost

upvoted 1 times

  **mdsuresh** 3 years, 2 months ago

If the traffic sourced behind L3out to an endpoint which moved to another leaf in ACI fabric, then answer is B.



If traffic sourced from ACI fabric, to the endpoint which was in ACI fabric itself and now it is moved to L3out, then the Answer is D.

upvoted 1 times

  **muhnator** 3 years, 4 months ago



answer is D

upvoted 1 times

  **Ruslans** 3 years, 11 months ago

D, Because of this stale remote endpoint, any traffic from LEAF1 toward IP2 will fail, because LEAF1 sends packets to the wrong leaf.

upvoted 1 times

  **danhvu** 4 years, 5 months ago

I think the answer is D

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

upvoted 4 times

Which action sets Layer 2 loop migration in an ACI Fabric with a Layer 2 Out configured?

- A. Enable MCP on the ACI fabric.
- B. Disable STP in the external network.
- C. Disable STP on the ACI fabric.
- D. Enable STP on the ACI fabric.

Correct Answer: A

🗨️ 👤 **designated** 7 months, 2 weeks ago

Selected Answer: A

A is correct

MisCabling Protocol (MCP) detects loops from external sources (such as misbehaving servers and external networking equipment running STP) and will err-disable the interface on which Cisco ACI receives its own packet.

The best practice is to enable this option (potentially also with "Enable MCP PDU per VLAN") on leaf node ports that are connected to external Layer 2 networks that may introduce loops.

upvoted 1 times

🗨️ 👤 **Marinheiro** 1 year, 1 month ago

It's loop mitigation, not "migration".

upvoted 2 times

🗨️ 👤 **villain_jack** 3 years, 3 months ago

A is correct

https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-aci-best-practices-quick-summary.html#_Toc97911120

upvoted 4 times


An engineer is implementing a connection that represents an external bridged network. Which two configurations are used? (Choose two.)

- A. Layer 2 remote fabric
- B. Layer 2 outside
- C. Layers 2 internal
- D. Static path binding
- E. VXLAN outside

Correct Answer: BD

Community vote distribution



BD (100%)

  **thiyagas**  4 years, 5 months ago

correct answer is BD

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c07-732033.html#_Toc395143568

upvoted 25 times

  **hybersat** 4 years, 5 months ago

I agree, here is the snip from the docs above:

- Extend the bridge domain out of the ACI fabric - Another option to extend the layer 2 domain is to create a layer 2 outside connection (or external bridged network, as called in the APIC GUI) for a given bridge domain. It effectively extends the bridge domain to the outside network.




The following sections explain these three options in greater detail.

Extend the EPG Out of the ACI Fabric

The user can extend an EPG beyond an ACI leaf by statically assigning a leaf port (along with a VLAN ID) to an EPG. After doing so, all the traffic received on this leaf port with the configured VLAN ID will be mapped to the EPG and the configured policy for this EPG will be enforced. The endpoints need not be directly connected to the ACI leaf port. They can be behind a layer 2 network as long as the VLAN associated with the EPG is enabled within the layer 2 network that connects the remote endpoint to the ACI fabric.

To statically assign port to an EPG, go to menu Tenant>Application Profiles>EPG>Static Binding (Paths).



upvoted 7 times

  **designated**  7 months, 2 weeks ago

Selected Answer: BD

B and D are correct



upvoted 1 times

  **7korn7** 2 years, 4 months ago

Selected Answer: BD

B and d is the correct answer

upvoted 1 times

  **Gab99** 2 years, 8 months ago

Selected Answer: BD

B D seems more logical

upvoted 1 times

  **korthab** 2 years, 11 months ago

Selected Answer: BD

I think the answer is B, D


upvoted 1 times

  **patopizarro** 3 years, 2 months ago

Selected Answer: BD



i think answer is BD

upvoted 1 times

  **ciscoaci2022** 3 years, 3 months ago

Agree, BD is correct

upvoted 1 times

  **Dontshoot0095** 3 years, 12 months ago

I Agree - BD

upvoted 3 times

Which two actions extend a Layer 2 domain beyond the ACI fabric? (Choose two.)

- A. extending the routed domain out of the ACI fabric
- B. creating a single homed Layer 3 Out
- C. creating an external physical network
- D. extending the bridge domain out of the ACI fabric
- E. extending the EPG out of the ACI fabric

Correct Answer: DE

Community vote distribution

DE (100%)

  **[Removed]**  4 years, 5 months ago

Answer should be D, E

<https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2015/pdf/LABSDN-2331-LG.pdf>

upvoted 24 times

  **thiyagas**  4 years, 5 months ago

correct answer is DE

upvoted 11 times

  **designated**  7 months, 2 weeks ago

Selected Answer: DE

D and E are correct

L3Out and EPG extension



upvoted 1 times

  **sailorsoul** 1 year, 5 months ago

Selected Answer: DE

100% DE.

upvoted 1 times

  **Gab99** 2 years, 8 months ago

Selected Answer: DE

most likely

upvoted 1 times

  **SherifAbdelMoteleb** 3 years, 1 month ago

Selected Answer: DE

Answer should be D, E

upvoted 1 times

  **Brute_Wane** 3 years, 1 month ago

Selected Answer: DE

Correct answer should be D and E.

upvoted 1 times

  **Leviatan_93** 3 years, 11 months ago

I agree with us, the correct answer is D,E

upvoted 4 times

When Cisco ACI connects to an outside Layers 2 network, where does the ACI fabric flood the STP BPDU frame?

- A. within the bridge domain
- B. within the APIC
- C. within the access encap VLAN
- D. between all the spine and leaf switches

Correct Answer: C

Community vote distribution

C (89%)

11%

 **[Removed]**  4 years, 5 months ago

Correct Answer: C.

STP BPDUs from Legacy switches are flooded within the EPG, not the BD. This is a change from most all flooded traffic in ACI. Most of the time when we talk about traffic being flooded inside of ACI, the flooding is occurring at the BD level.

upvoted 13 times



 **[Removed]**  4 years, 5 months ago

Answer is C

The ACI fabric is an IP-based fabric that implements an integrated overlay, allowing any subnet to be placed anywhere in the fabric and supports a fabric-wide mobility domain for virtualized workloads. STP is not required within the ACI fabric and leaf. The spine and APIC don't run STP instances.

When connecting to an outside layer 2 network, the ACI fabric floods the STP BPDU frame within the boundary of the EPG. External switches are expected to break any potential loop upon receiving the flooded BPDU from the ACI fabric. Figure 69 depicts this process.

upvoted 8 times

 **designated**  7 months, 2 weeks ago

Selected Answer: C

For instance, if EPG1, port 1/1, is configured to match VLAN 5 from a switch, another port of that switch for that same Layer 2 domain can be connected only to EPG1 using the same encapsulation of VLAN 5. Otherwise, the external switch would receive the BPDU for VLAN 5 tagged with a different VLAN number.


***** Cisco ACI floods BPDUs only between the ports in the bridge domain that have the same encapsulation*****

upvoted 1 times

 **[Removed]** 1 year, 1 month ago

I come back three years later and the answer is still un changed and it is C

upvoted 3 times

 **kamel86** 1 year, 8 months ago

Selected Answer: C

Among the Layer 2 frames that require multidestination forwarding, Cisco ACI handles spanning tree BPDUs in a slightly different way than other frames because to avoid loops and to preserve the access encapsulation VLAN information associated to the BPDU (within the bridge domain), this traffic is assigned the VXLAN VNID that identifies the access encapsulation VLAN (instead of the bridge domain VNID) and flooded to all ports of the bridge domain that carry the same access encapsulation (regardless of the EPG).

upvoted 1 times

 **Mr_Certifiable** 2 years, 1 month ago

References:

The same VLAN number can be used by one EPG on one leaf and by another EPG on a different leaf.

If the two EPGs are in the same bridge domain, they share the same flood domain VLAN for BPDUs and they share the broadcast domain.

(Page 323 - Deploying ACI the complete guide)

BPDUs are flooded within the fabric encap of an EPG (allocated based on domain/vlanpool).

In order for BPDUs to be flooded properly, all interfaces within the EPG that are connected to external bridges MUST reside in the same physical or L2 external domain

and vlan encapsulation. (BRKACI -3101- Page 76 - Common mistakes that cause loops)

<https://learningnetwork.cisco.com/s/question/0D53i00000KsrWnCAJ/aci-flood-domain-for-bpdu-different-physical-domains>

upvoted 1 times

🗲️ 👤 **jecq** 2 years, 5 months ago

Sorry, the answer is C.

upvoted 1 times

🗲️ 👤 **jecq** 2 years, 6 months ago

Selected Answer: A

The correct answer is A. The answer C is incorrect, because you can have the same vlan encaps even in another Tenants.

upvoted 1 times

🗲️ 👤 **korthab** 2 years, 10 months ago

Selected Answer: C

Based on the text below i think the answer is C:

Among the Layer 2 frames that require multidestination forwarding, Cisco ACI handles spanning tree BPDUs in a slightly different way than other frames because to avoid loops and to preserve the access encapsulation VLAN information associated to the BPDU (within the bridge domain), this traffic is assigned the VXLAN VNID that identifies the access encapsulation VLAN (instead of the bridge domain VNID) and flooded to all ports of the bridge domain that carry the same access encapsulation (regardless of the EPG). This behavior also applies more in general to Layer 2 flooding when using the feature called "Flood in Encapsulation". In this document, we refer to this specific encapsulation as the FD_VLAN VXLAN encapsulation or FD_VLAN VNID, or FD VNID for simplicity. The FD_VLAN fabric encapsulation (or FD_VLAN VNID or FD VNID) is different from the bridge domain VNID.

Source:

<https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-application-centric-infrastructure-design-guide.html>

upvoted 3 times

🗲️ 👤 **jmaroto** 2 years, 12 months ago

Selected Answer: C

I think C is the good one. The Cisco ACI fabric does not run Spanning Tree Protocol natively, but it can forward BPDUs within the EPGs.

The flooding scope for BPDUs is different from the flooding scope for data traffic. The unknown unicast traffic and broadcast traffic are flooded within the bridge domain. Spanning Tree Protocol BPDUs are flooded within a specific VLAN encapsulation (also known as FD_VLAN), and in many cases, though not necessarily, an EPG corresponds to a VLAN.

upvoted 3 times

🗲️ 👤 **Brute_Wane** 3 years, 1 month ago

Selected Answer: C

<https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-application-centric-infrastructure-design-guide.html>

upvoted 1 times

🗲️ 👤 **ciscoaci2022** 3 years, 3 months ago

correct answer is C

upvoted 1 times

🗲️ 👤 **rhonen** 3 years, 9 months ago

Each BPDU frame is flooded within each access encaps VLAN ID. No configuration is required for the BPDU flooding. The external switches are in charge of breaking any potential loops.

upvoted 2 times

🗲️ 👤 **Dontshoot0095** 3 years, 12 months ago

C is correct

upvoted 2 times

🗲️ 👤 **mrpeet** 4 years ago

Correct answer is A.

From referenced documentation:

"The ACI leaf floods the BPDU frame within the EPG by using the VXLAN network identifier (VNID) assigned for the EPG when it encapsulates the BPDU in VXLAN format. The flooding scope of the BPDU is different than the one for data traffic. The unknown unicast traffic and broadcast traffic are flooded within the bridge domain. On the outside layer 2 network, STP instances are aligned with the VLAN boundary. To keep it consistent, the ACI fabric maintains the STP boundary by flooding the BPDU within the scope of the EPG."

Answer C mentions the encapsulated VLAN where the documentation mentions VXLAN so this is completely different. With these setups, typically VLAN=EPG=BD so answer A is the best option.

upvoted 4 times

🗨️ 👤 **RTL_dude** 4 years, 2 months ago

Indeed, C.

Correct link: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c07-732033.html#_Toc395143573

upvoted 3 times

🗨️ 👤 **thiyagas** 4 years, 5 months ago

correct answer is C

upvoted 6 times

On which two interface types should a user configure storm control to protect against broadcast traffic? (Choose two.)

- A. APIC facing interfaces
- B. port channel on a single leaf switch
- C. all interfaces on the leaf switches in the fabric
- D. endpoint-facing trunk interface
- E. fabric uplink interfaces on the leaf switches

Correct Answer: BD

Community vote distribution

BD (100%)

🗳️ 👤 **designated** 7 months, 2 weeks ago

Selected Answer: BD

B and D are correct

A traffic storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. You can use traffic storm control policies to prevent disruptions on Layer 2 ports by broadcast, unknown multicast, or unknown unicast traffic storms on physical interfaces. By default, storm control is not enabled in the ACI fabric. ACI bridge domain (BD) Layer 2 unknown unicast flooding is enabled by default within the BD but can be disabled by an administrator. In that case, a storm control policy only applies to broadcast and unknown multicast traffic. If Layer 2 unknown unicast flooding is enabled in a BD, then a storm control policy applies to Layer 2 unknown unicast flooding in addition to broadcast and unknown multicast traffic. Traffic storm control (also called traffic suppression)

upvoted 1 times

🗳️ 👤 **patopizarro** 3 years, 2 months ago

Selected Answer: BD

B and D.

Typically, a fabric administrator configures storm control in fabric access policies on the following interfaces:

A regular trunk interface.

A direct port channel on a single leaf switch.

A virtual port channel (a port channel on two leaf switches).

upvoted 4 times

🗳️ 👤 **muhnator** 3 years, 4 months ago

B and D are correct

upvoted 1 times

🗳️ 👤 **manet** 3 years, 12 months ago

[https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/L2_config/b_Cisco_APIC_Layer_2_Configuration_Guide/b_Cisco_APIC_Layer_2_Configuration_Guide_chapter_01010.html)

[x/L2_config/b_Cisco_APIC_Layer_2_Configuration_Guide/b_Cisco_APIC_Layer_2_Configuration_Guide_chapter_01010.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/L2_config/b_Cisco_APIC_Layer_2_Configuration_Guide/b_Cisco_APIC_Layer_2_Configuration_Guide_chapter_01010.html)

upvoted 2 times

Which two dynamic routing protocols are supported when using Cisco ACI to connect to an external Layer 3 network? (Choose two.)

- A. iBGP
- B. VXLAN
- C. IS-IS
- D. RIPv2
- E. eBGP

Correct Answer: AE

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/ACI_Best_Practices/b_ACI_Best_Practices/b_ACI_Best_Practices_chapter_010010.html

Community vote distribution

AE (100%)

 **designated** 7 months, 2 weeks ago

Selected Answer: AE

A and E are correct.

The supported dynamic routing protocols are BGP, OSPF, or EIGRP.

A routed Layer 3 connection uses a set of protocols that determine the path that data follows in order to travel across multiple networks from its source to its destination. Cisco ACI routed connections perform IP forwarding according to the protocol selected, such as BGP, OSPF, or EIGRP.
upvoted 2 times

 **[Removed]** 1 year, 1 month ago

Not sure how old this question is,, answer is A, E.


Now there is OSPF, Static Routes and Eigrp for L3outs
upvoted 1 times

 **Lorygru** 2 years, 1 month ago

Selected Answer: AE

AE is correct

upvoted 1 times

 **7korn7** 2 years, 4 months ago

Selected Answer: AE

The answer is correct

upvoted 1 times

 **Supreme_123** 2 years, 8 months ago

I think, it's correct A,E

upvoted 1 times

What must be configured to redistribute externally learned OSPF routes within the ACI fabric?

- A. Route Control Profile
- B. BGP Route Reflector
- C. BGP Inter-leak Route Map
- D. PIM Sparse Mode

Correct Answer: B

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/ACI_Best_Practices/b_ACI_Best_Practices/b_ACI_Best_Practices_chapter_010010.html

Community vote distribution

B (50%)

A (50%)


 **targ** Highly Voted 4 years, 4 months ago

Correct Answer B

Once the border leaf learns the external routes, it can then redistribute the external routes of a given VRF instance to an MP-BGP address family (VPNv4 or VPNv6). MP-BGP maintains a separate BGP routing table for each VRF instance. Within MP-BGP, the border leaf switch advertises routes to a spine switch, which is a BGP route reflector. The routes are then propagated to all the leaf switches where the VRF instances are instantiated.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/guide-c07-743150.html>

upvoted 12 times

 **hybersat** Highly Voted 4 years, 5 months ago

Shouldn't this be A?

We are not talking transit routing here...

upvoted 5 times

 **prospio971** Most Recent 4 months, 2 weeks ago

Selected Answer: A

To redistribute externally learned OSPF routes within the Cisco ACI fabric, a Route Control Profile must be configured. This profile allows you to control which routes are imported or exported between the ACI fabric and external routing domains.

BGP Route Reflector: This is used for internal BGP route reflection within the ACI fabric and is unrelated to OSPF redistribution.

upvoted 1 times

 **designated** 7 months, 2 weeks ago

Selected Answer: B

B is correct..


WITHIN ACI FABRIC is Route Reflector

upvoted 1 times

 **Tavlon** 1 year ago

Its actually under Route Control Enforcement , profile so its simply A

upvoted 1 times

 **zelya19** 1 year, 3 months ago

Selected Answer: B

In Cisco ACI, a border leaf node on which Layer 3 Out (L3Outs) are deployed redistributes L3Out routes to the BGP IPv4/IPv6 address family and then to the MP-BGP VPNv4/VPNv6 address family along with the VRF information so that L3Out routes are distributed from a border leaf node to other leaf nodes through the spine nodes. Interleak redistribution in the Cisco ACI fabric refers to this redistribution of L3Out routes to the BGP IPv4/IPv6 address family. By default, interleak happens for all L3Out routes, such as routes learned through dynamic routing protocols, static routes, and directly-connected subnets of L3Out interfaces, except for routes learned through BGP. Routes learned through BGP are already in the BGP IPv4/IPv6 table and are ready to be exported to MP-BGP VPNv4/VPNv6 without interleak.

<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/5-x/l3-configuration/cisco-apic-layer-3-networking-configuration-guide-50x/route-control-50x.html>

upvoted 1 times

  **Huberttheman** 1 year, 5 months ago

Selected Answer: A

To redistribute externally learned OSPF routes within the Cisco ACI fabric, you need to configure a Route Control Profile. This profile is used to control the import and export of routes between different routing protocols or instances within the ACI fabric, including the redistribution of routes from external sources such as OSPF into the ACI fabric's routing protocol

upvoted 1 times

  **Mr_Certifiable** 2 years ago

B - Distribute external routes within the ACI fabric

ACI uses Multi-Protocol BGP (MP-BGP) with VPNv4 in the ACI infra VRF (overlay-1 VRF) to distribute external routes from a border leaf to other leaf switches. Similar to other configurations/components in the ACI infra VRF such as ISIS between each switch, this configuration is also automated in the background. The only two configurations that users need to perform are as follows:

- Select the BGP AS number.
- This is the AS number to represent the entire ACI fabric. It is used for infra MP-BGP between leaf and spines, and for BGP in user L3Outs to establish BGP peers with external devices.
- Select spine switches as BGP Route Reflectors.
- Each leaf switch will be a BGP client for the selected route-reflector spine switches.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/guide-c07-743150.html>

upvoted 1 times

  **Narbledeath** 2 years, 2 months ago

Its B

2. Distribute external routes within the ACI fabric

ACI uses Multi-Protocol BGP (MP-BGP) with VPNv4 in the ACI infra VRF (overlay-1 VRF) to distribute external routes from a border leaf to other leaf switches. Similar to other configurations/components in the ACI infra VRF such as ISIS between each switch, this configuration is also automated in the background. The only two configurations that users need to perform are as follows:

- Select the BGP AS number.
- This is the AS number to represent the entire ACI fabric. It is used for infra MP-BGP between leaf and spines, and for BGP in user L3Outs to establish BGP peers with external devices.
- Select spine switches as BGP Route Reflectors.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/guide-c07-743150.html>

upvoted 1 times

  **[Removed]** 4 years, 5 months ago

Correct Answer:B.


upvoted 3 times

Regarding the MTU value of MP-BGP EVPN control plane packets in Cisco ACI, which statement about communication between spine nodes in different sites is true?

- A. By default, spine nodes generate 9000-bytes packets to exchange endpoints routing information. As a result, the Inter-Site network should be able to carry 9000-bytes packets.
- B. By default, spine nodes generate 1500-bytes packets to exchange endpoints routing information. As a result, the Inter-Site network should be able to carry 1800-bytes packets.
- C. By default, spine nodes generate 1500-bytes packets to exchange endpoints routing information. As a result, the Inter-Site network should be able to carry 1500-bytes packets.
- D. By default, spine nodes generate 9000-bytes packets to exchange endpoints routing information. As a result, the Inter-Site network should be able to carry 9100-bytes packets.

Correct Answer: D

Community vote distribution



therufus Highly Voted 4 years, 3 months ago

Correct answer is A

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

Excerpt:

MTU of MP-BGP control-plane communication between spine nodes in different sites: By default, the spine nodes generate 9000-byte packets for exchanging endpoint routing information. If that default value is not modified, the ISN must support an MTU size of at least 9000 bytes
upvoted 13 times

Jey10 4 years, 1 month ago

It is D from the link you give => "Maximum MTU of the frames generated by the endpoints connected to the fabric: If the endpoints are configured to support jumbo frames (9000 bytes), then the ISN should be configured with at least a 9100-byte MTU value. If the endpoints are instead configured with the default 1500-byte value, then the ISN MTU size can be reduced to 1600 bytes"
upvoted 2 times

anasham 3 years, 6 months ago

@Jey10; what you have written is about Data Plane MTU. The question is about Control plane MTU. So the answer is A as therufus said
upvoted 6 times

RTL_dude 4 years, 2 months ago

I agree, the question is about exchanging endpoint routing information (i.e. control plane traffic) *between* spines in different sites, not local endpoint to spine traffic for which the 9100 byte MTU applies to.
upvoted 3 times

veld2 3 years, 3 months ago

MTU of MP-BGP control-plane communication between spine nodes in different sites: By default, the spine nodes generate 9000-byte packets for exchanging endpoint routing information. If that default value is not modified, the ISN must support an MTU size of at least 9000 bytes, otherwise the exchange of control-plane information across sites would not succeed (despite being able to establish MP-BGP adjacencies).

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html>
upvoted 1 times

ashwind123 Most Recent 3 months, 1 week ago

Selected Answer: A

Its A as it explicitly talk about control plane MTU and not data plane MTU
upvoted 1 times

designated 7 months, 2 weeks ago

Selected Answer: A

A is correct..

There is not overhead from VXLAN for control plane traffic.

Go to System > Control Plane MTU and you will see that the value that could be set is between 576 and 9000.

And we can also set this option to 1500 for remote-leaf architecture.



<https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/6x/system-management-configuration/cisco-apic-system-management-configuration-guide-60x/basic-operations-60x.html>

upvoted 2 times

  **Lalag** 11 months, 3 weeks ago

A is the right answer as per the link <https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html>

upvoted 1 times

  **zelya19** 1 year, 3 months ago

Selected Answer: A

No mentioning of supporting jumbo frames within the fabric -> A.

upvoted 1 times

  **959836c** 1 year, 6 months ago

Selected Answer: D

Answer is D. While it's true the question is about the control plane and a mtu of 9000 is suitable. For the data plane you need to configure 9100 for mtu. In a real world environment you need to account for the data plane and go with the highest value. It could be A, but it's a technically either way.

upvoted 1 times

  **Macc10** 1 year, 8 months ago

Selected Answer: A

Difficult question but I am swaying towards A rather than B due to the fact it mentions control plane within the question...

upvoted 1 times

  **kamel86** 1 year, 8 months ago

Selected Answer: A

From ACI Multi-site white paper:

The MTU of MP-BGP control-plane communication between spine nodes in different sites: By default, the spine nodes generate 9000-byte packets for exchanging endpoint routing information. If that default value is not modified, the ISN must support an MTU size of at least 9000 bytes; otherwise, the MP-BGP exchange of control-plane information across sites would not succeed.

upvoted 1 times

  **Chandaria** 1 year, 10 months ago

Here the ask is about control plane so answer is A.

If it was for data plane then answer would be D.

upvoted 1 times

  **Mr_Certifiable** 2 years, 1 month ago

IPN requirements for a Remote Leaf solution are as follows:

MTU: The solution must support an end-to-end MTU that is at least 100 bytes higher than that of the endpoint source traffic. Assuming that 1500 bytes has been configured for data plane MTU, Remote Leaf can be deployed using a minimum MTU of 1600 bytes. An IPN MTU this low, however, necessitates that ACI administrators lower the ACI fabricwide control plane MTU, which is 9000 bytes by default.

<https://www.ciscopress.com/articles/article.asp?p=3150964&seqNum=3>

upvoted 1 times

  **Rododendron2** 2 years, 2 months ago

Selected Answer: D

The minimum MTU value to configure in the ISN depends on two factors:

- The maximum MTU of the frames generated by the endpoints connected to the fabric: If the endpoints are configured to support jumbo frames (9000 bytes), then the ISN should be configured with at least a 9050-byte MTU value. If the endpoints are instead configured with the default 1500-byte value, then the ISN MTU size can be reduced to 1550 bytes.

upvoted 1 times

  **bizzar777** 2 years, 8 months ago

Selected Answer: D

Maximum transmission unit (MTU) of Multiprotocol Border Gateway Protocol (MP-BGP) Ethernet Virtual Private Network (EVPN) control plane communication between spine nodes in different sites - By default, the spine nodes generate 9000-byte packets to exchange endpoint routing

information. If that default value is not modified, the Inter Site Network (ISN) must support an MTU size of at least 9100 bytes. In order to tune the default value, modify the corresponding system settings in each APIC domain.

upvoted 2 times

🗨️ 👤 **mdsuresh** 2 years, 9 months ago

D is correct!

upvoted 1 times

🗨️ 👤 **korthab** 2 years, 10 months ago

This question has two answers on the internet.

But if you want to support endpoint traffic of a MTU of 9000, it is required to increase the MTU to at least 9100 on the ISN.

So what is this question about? Something that is minimal required for the MP-BGP protocol to work? Or an recommended setup looking further than only the communication between Spines.

I would like to pick MTU 9100 for the ISN, but not sure if it is a trick question.

LINK 1 = MTU 9100

<https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/214270-configure-aci-multi-site-deployment.html>

LINK 2 = MTU 9000

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

upvoted 2 times

🗨️ 👤 **eddyedwards257** 2 years, 11 months ago

Maximum MTU of the frames generated by the endpoints connected to the fabric: If the endpoints are configured to support jumbo frames (9000 bytes), then the ISN should be configured with at least a 9100-byte MTU value. If the endpoints are instead configured with the default 1500-byte value, then the ISN MTU size can be reduced to 1600 bytes.

- MTU of MP-BGP control-plane communication between spine nodes in different sites: By default, the spine nodes generate 9000-byte packets for exchanging endpoint routing information. If that default value is not modified, the ISN must support an MTU size of at least 9000 bytes, otherwise the exchange of control-plane information across sites would not succeed (despite being able to establish MP-BGP adjacencies). The default value can be tuned by modifying the corresponding system settings in each APIC domain

upvoted 1 times

🗨️ 👤 **leetingo** 3 years, 4 months ago

the ISN mtu value needed depends on whether you need jumbo frame support for endpoint or not. In this question it does not say anything about endpoint jumbo frame support, in this case 9000 is enough for ISN control plane packets

upvoted 1 times

🗨️ 👤 **Player** 3 years, 6 months ago

A is correct. We are asked exactly about the control plane traffic between sites, not about the size of frames of the EPs connected to the fabric. Link provided by therufus is clearly describing that.

upvoted 1 times

The screenshot shows the Cisco ACI GUI configuration for a connection. On the left is a navigation tree with folders like 'L4-L7', 'Service Parameters', 'Service Graph Templates', and 'Prod_to_Trans'. The main area displays two tables:

Terminate Nodes:	
Name	Provider/Consumer
T1	Consumer
T2	Provider

Connections:				
Name	Connected Nodes	Direct Connect	Unicast Route	Adjacency Type
C1	N1, T1	<input type="checkbox"/>	<input checked="" type="checkbox"/>	
C2	N1, T2	False		

At the bottom right of the connections table are 'Update' and 'Cancel' buttons. The 'Adjacency Type' dropdown for the second connection (C2) is highlighted with a red box.

Refer to the exhibit. Which Adjacency Type value should be set when the client endpoint and the service node interface are in a different subnet?

- A. Routed
- B. Unicast
- C. L3Out
- D. L3

Correct Answer: D

prospio971 5 months, 1 week ago

Selected Answer: D

When the client endpoint and the service node interface are in different subnets, Layer 3 routing is required for communication between them. Setting the Adjacency Type to L3 ensures that the ACI fabric performs the necessary Layer 3 forwarding between the different subnets.

upvoted 1 times

crooks_1988 9 months, 3 weeks ago

D. L3 or L2 available for configuration only

upvoted 2 times

Which endpoint learning operation is completed on the ingress leaf switch when traffic is received from a Layer 3 Out?

- A. The source MAC address of the traffic is learned as a local endpoint.
- B. The source MAC address of the traffic is learned as a remote endpoint.
- C. The source IP address of the traffic is learned as a remote endpoint.
- D. The source IP address of the traffic is learned as a local endpoint.

Correct Answer: A

Community vote distribution

A (100%)

  **TCoder**  4 years, 4 months ago

It is A. The source MAC address is learnt as local endpoint.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.pdf>

Refer to Table 6 which says,

Local endpoint learning with an incoming packet from L3Out
to Cisco ACI:

Only the source MAC address is learned as a local endpoint.

The source IP address is not learned as a local endpoint.

upvoted 19 times

  **Kalpesh**  4 years, 4 months ago

Isn't it A? As per below document, in endpoint learning behavior for traffic coming in from L3 out, only source mac address will be learned as local endpoint and of course that will be on ingress leaf for traffic coming in

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

upvoted 7 times

  **designated**  7 months, 2 weeks ago



Selected Answer: C

C is correct..

When traffic is received on an ACI leaf switch from a Layer 3 Out (L3Out) connection, the source IP address of the traffic is learned as a remote endpoint. This is because the endpoint is external to the ACI fabric and communicates with the fabric through the L3Out. The ACI fabric tracks such endpoints as remote entities since they are not directly connected to the fabric as local endpoints.

- Local Endpoint Learning: Refers to endpoints connected directly to the ACI leaf switch on an access port or through a Layer 2 domain.
- Remote Endpoint Learning: Refers to endpoints external to the ACI fabric, typically reachable via an L3Out or through VXLAN tunnels from another leaf.



upvoted 2 times

  **emoemo** 1 year, 10 months ago

I think B is correct.

A is a description of "dest MAC", I thought.



upvoted 1 times

  **emoemo** 1 year, 10 months ago

sorry, This question is about "ingress leaf".

So, The correct answer is C ???

upvoted 3 times

  **hebdeb** 1 year, 11 months ago

Selected Answer: A

It is A

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

739989.html#L3Outendpointlearningconsiderations

upvoted 1 times

  **asd248402** 2 years ago

Selected Answer: A

Scenario 1

Local endpoint learning with an incoming packet from L3Out to Cisco ACI:

Only the source MAC address is learned as a local endpoint. The source IP address is not learned as a local endpoint.

upvoted 1 times

  **Mr_Certifiable** 2 years ago

A

Scenario 1

Local endpoint learning with an incoming packet from L3Out to Cisco ACI:

Only the source MAC address is learned as a local endpoint. The source IP address is not learned as a local endpoint.



Scenario 2

Remote endpoint learning with an incoming packet from L3Out to Cisco ACI:

No source MAC or IP address is learned as a new remote endpoint by a packet.*

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

upvoted 1 times

  **7korn7** 2 years, 2 months ago

Selected Answer: A

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

Local endpoint learning with an incoming packet from L3Out to Cisco ACI:


Only the source MAC address is learned as a local endpoint. The source IP address is not learned as a local endpoint.

upvoted 1 times

  **korthab** 2 years, 10 months ago

This question is not stating where the L3Out is connected and therefore it is not possible to determine it is a local or remote endpoint.

upvoted 1 times

  **eddyedwards257** 2 years, 11 months ago

Is it not B ? The source MAC on this leaf is learnt as a remote endpoint given the L3 out from which the end host is being may be on another leaf . It does not say it's the same leaf as what the L3 configured on ?

upvoted 1 times

  **Dspam** 3 years, 10 months ago

Local endpoint learning with an incoming packet from L3Out to Cisco ACI:

Only the source MAC address is learned as a local endpoint. The source IP address is not learned as a local endpoint.

upvoted 1 times

An engineer must connect Cisco ACI fabric using Layer 2 with external third-party switches. The third-party switches are configured using 802.1s protocol. Which two constructs are required to complete the task? (Choose two.)

- A. spanning tree policy for mapping MST Instances to VLANs
- B. MCP policy with PDU per VLAN enabled
- C. MCP instance policy with administrative slate disabled
- D. dedicated EPG for native VLAN
- E. static binding of native VLAN in all existing EPGs

Correct Answer: AD

Community vote distribution


AD (100%)

 **thiyagas** Highly Voted 3 years, 11 months ago

i think the answer is A and D

1. Create a special MST EPG and map it to all ports facing non-ACI switches that run MST. This ensures that ACI does not drop MST BPDUs.
2. Navigate to Fabric > Access Policies > Policies > Switch > Spanning-Tree > default and create MST region policies that include the MST region names, MST instance IDs, revision IDs, and relevant VLAN encapsulations. This ensures that ACI knows which EPGs to flush when it receives a TCN and also out of which ports it should forward MST BPDUs.

upvoted 21 times

 **nikomski** 3 years, 11 months ago

I agree

upvoted 2 times

 **Mohitkrsh84** 3 years, 10 months ago


A & D is correct.

MST (IEEE 802.1s), BPDUs do not carry a VLAN tag, and the BPDUs are sent over the native VLAN. Typically, the native VLAN is not used to carry data traffic, and the native VLAN may not be configured for data traffic on the Cisco ACI fabric. As a result, to help ensure that MST BPDUs are flooded to the desired ports, the user must create an EPG (an MST EPG) for VLAN 1 as native VLAN to carry the BPDUs. This EPG connects to the external switches that run MST.

In addition, the administrator must configure the mapping of MST instances to VLANs to define which MAC address table must be flushed when a Topology Change Notification (TCN) occurs. When a TCN event occurs on the external Layer 2 network, this TCN reaches the leafs to which it connects via the MST EPG, and flushes the local endpoint information associated with these VLANs on these leafs; as result, these entries are removed from the spine-proxy mapping database.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737909.html>

upvoted 8 times

 **sailorsoul** Most Recent 11 months, 1 week ago

Selected Answer: AD

dedicated EPG for native vlan when running MST

upvoted 1 times

 **Gab99** 2 years, 2 months ago

Selected Answer: AD


A and D makes sense

upvoted 3 times

 **rsm123th** 2 years, 6 months ago

B and E from <https://community.cisco.com/t5/data-center-and-cloud-documents/spanning-tree-mst-switches-interaction-with-aci/ta-p/3146184>

upvoted 1 times

 **thomyohan** 3 years, 6 months ago

Answer is A and D

[To accept traffic for any VLAN, the VLAN needs to be provisioned, either by a statically assigned port and VLAN to an EPG, or by the EPG being provisioned dynamically by the APIC when there is integration between the APIC and Virtual Machine Manager (VMM), (such as vCenter or Microsoft SCVMM). As a result, to ensure MST BPDU is flooded to the desired ports, the user needs to create an EPG to carry the BPDU]



This answer is not apt - ----> E. static binding of (native VLAN)* in all existing EPGs
upvoted 3 times

  **[Removed]** 3 years, 11 months ago

Answer is D and E

Additional configuration is required in order for Multiple Spanning Tree (MST) BPDU to be flooded properly. The BPDU frame for Per-VLAN Spanning Tree (PVST) and Rapid Per-VLAN Spanning Tree (RPVST) have a VLAN tag. The ACI leaf can identify which EPG the BPDU needs to be flooded based on the VLAN tag in the frame. However, for MST (802.1s), BPDU frames don't carry a VLAN tag and they are sent over the native VLAN. Typically, the native VLAN is not used to carry data traffic and the native VLAN may not be configured for data traffic on the ACI fabric. By default there is no native VLAN enabled on the ACI fabric. To accept traffic for any VLAN, the VLAN needs to be provisioned, either by a statically assigned port and VLAN to an EPG, or by the EPG being provisioned dynamically by the APIC when there is integration between the APIC and Virtual Machine Manager (VMM), (such as vCenter or Microsoft SCVMM). As a result, to ensure MST BPDU is flooded to the desired ports, the user needs to create an EPG to carry the BPDU. As shown in Figure 71, the mode needs to be "native" given that the BPDU frame is untagged.

upvoted 4 times

  **hybersat** 3 years, 11 months ago

I agree DE, der this for further details: <https://community.cisco.com/t5/data-center-documents/spanning-tree-mst-switches-interaction-with-aci/ta-p/3146184>

upvoted 3 times

Create L3Out



1. Identity

2. Nodes And Interfaces

3. Protocols

4. External EPG

Nodes and Interfaces

The L3Out configuration consists of node profiles and interface profiles. An L3Out can span across multiple nodes in the fabric. All nodes used by the L3Out can be included in a single node profile and is required for nodes that are part of a VPC pair. Interface profiles can include multiple interfaces. When configuring dual stack interfaces a separate interface profile is required for the IPv4 and IPv6 configuration, that is automatically taken care of by this wizard.

Use Defaults: ☒

Interface Types

Layer 3: **Routed** Routed Sub SVI Floating SVILayer 2: **Port** Direct Port Channel

Nodes

Node ID F1P1L1 (Node - 1001)	Router ID 10.1.7.1	Loopback Address 10.1.1.1 <small>Leave empty to not configure any Loopback</small>	<input type="checkbox"/> + Hide Interfaces
Interface Select a port	IP Address <input type="text"/> <small>address/mask</small>	MTU (bytes) inherit	<input type="checkbox"/> +

Previous

Cancel

Next

Refer to the exhibit. An engineer must configure an L3Out peering with the backbone network. The L3Out must forward unicast and multicast traffic over the link.

Which two methods should be used to configure L3Out to meet these requirements? (Choose two.)

- A. Layer 3 routed port
- B. VPC with SVI
- C. port channel with SVI
- D. Layer 3 routed subinterface
- E. Layer 3 floating SVI

Correct Answer: AD

Community vote distribution

AD (100%)

Orlando1 Highly Voted 4 years, 3 months ago

Should be A and D. "PIM is supported on Layer 3 Out routed interfaces and routed subinterfaces including Layer 3 port-channel interfaces. PIM is not supported on Layer 3 Out SVI interfaces."

<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2->

x/L3_config/b_Cisco_APIC_Layer_3_Configuration_Guide/b_Cisco_APIC_Layer_3_Configuration_Guide_chapter_01111.html#id_21570

upvoted 12 times

nabilzay Highly Voted 4 years, 4 months ago

I think it is A and D as PIM is not supported on L3out SVIs

upvoted 5 times

designated Most Recent 7 months, 2 weeks ago

Selected Answer: AD

Wow.. this is too much..

PIM is supported on Layer 3 Out routed interfaces and routed subinterfaces including Layer 3 port-channel interfaces. PIM is not supported on Layer 3 Out SVI interfaces.

It was hard to find this.

upvoted 1 times

hebdeb 1 year, 11 months ago

Selected Answer: AD

Should be A and D

upvoted 2 times

🗲️ 👤 **7korn7** 2 years, 2 months ago

Selected Answer: AD

should be A and D

upvoted 1 times

🗲️ 👤 **Gab99** 2 years, 8 months ago

Selected Answer: AD

PIM not supported on SVI

upvoted 2 times

🗲️ 👤 **Nickname456** 2 years, 9 months ago

A, D Layer 3 Out ports and sub-interfaces are supported while external SVIs are not supported. Since external SVIs are not supported, PIM cannot be enabled in L3-VPC.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/L3-configuration/Cisco-APIC-Layer-3-Networking-Configuration-Guide-401/Cisco-APIC-Layer-3-Networking-Configuration-Guide-401_chapter_010010.html

upvoted 1 times

🗲️ 👤 **danhvu** 4 years, 5 months ago

I think A is used Routed port instead of the Routed interface

upvoted 1 times

🗲️ 👤 **nikomski** 4 years, 5 months ago

Why A is not correct?

upvoted 1 times

Create vCenter Domain

Virtual Switch Name:

Virtual Switch: VMware vSphere Distributed Switch Cisco AVS Cisco AVE

Associated Attachable Entity Profile:

Delimiter:

Enable Tag Collection: ☐

Access Mode: Read Only Mode Read Write Mode

Endpoint Retention Time (seconds):

VLAN Pool:

Security Domains:

Name	Description
------	-------------

Refer to the exhibit. An engineer is integrating a VMware vCenter with Cisco ACI VMM domain configuration. ACI creates port-group names with the format of

"Tenant | Application | EPG". Which configuration option is used to generate port groups with names formatted as "Tenant=Application=EPG"?

- A. enable tag collection
- B. security domains
- C. delimiter
- D. virtual switch name

Correct Answer: C

Community vote distribution

C (100%)

[Removed] Highly Voted 4 years, 5 months ago

Answer is C: Delimiter

Step 8 Configuring the delimiter during VMM domain creation, perform the following actions:

On the menu bar, choose VM NETWORKING > Inventory

In the Navigation pane, right-click VMware and click Create vCenter Domain.

In the Create vCenter Domain dialog box, enter a Name.

Optional: In the Delimiter field, enter one of the following: |, ~, !, @, ^, +, or =. If you do not enter a symbol, the system default | delimiter will appear in the VMware PortGroup name.

upvoted 17 times

designated Most Recent 7 months, 2 weeks ago

Selected Answer: C

C is correct

(Optional) In the Delimiter field, enter one of the following: |, ~, !, @, ^, +, or =.

<https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/virtualization-guide/cisco-aci-virtualization-guide-52x/m-cisco-aci-with-vmware-vds.html>

If you do not enter a symbol, the system default | delimiter will appear in the policy.

upvoted 1 times

asd248402 2 years ago

Selected Answer: C

Answer is C: Delimiter

upvoted 1 times

🗨️ 👤 **asd248402** 2 years ago

Selected Answer: C

delimiter make more sense

upvoted 1 times

🗨️ 👤 **Gab99** 2 years, 8 months ago

Selected Answer: C

I vote again

upvoted 1 times

🗨️ 👤 **Gab99** 2 years, 8 months ago

Selected Answer: C

makes sense only

upvoted 1 times

🗨️ 👤 **korthab** 2 years, 10 months ago

Selected Answer: C

Answer C is correct.

upvoted 1 times

🗨️ 👤 **anasham** 3 years, 11 months ago

By default, ACI names the distributed port group using the format
tenant|application|epg. The character separating the parameters is called a delimiter
character. In recent versions of ACI code, the naming can be customized using the
Custom EPG Name field. from the 00-620 ACI guide

upvoted 2 times

🗨️ 👤 **luisma84** 4 years, 2 months ago

Ans is C

upvoted 2 times

🗨️ 👤 **Carol254** 4 years, 4 months ago

Answer is C.

upvoted 1 times

🗨️ 👤 **thiyagas** 4 years, 5 months ago

Answer is C

upvoted 1 times

Create vCenter Domain

vCenter Credentials:

Profile name	Username	Description
Lab-VCenter	admin	

vCenter:

Name	IP	Type	Stats Collection
Lab-VCenter	vcenter.aci.lab	vCenter	Disabled

Port Channel Mode: select a value

- Static Channel – Mode On
- LACP Active
- LACP Passive
- MAC Pinning+
- MAC Pinning-Physical-NIC-load

Cancel Submit

Refer to the exhibit. An engineer is implementing Cisco ACI VMware vCenter integration for a blade server that lacks support of bonding. Which port channel mode results in "route based on originating virtual port" on the VMware VDS?

- A. Static Channel – Mode On
- B. MAC Pinning-Physical-NIC-load
- C. LACP Passive
- D. MAC Pinning+
- E. LACP Active

Correct Answer: D

Community vote distribution

D (100%)

designated 7 months, 2 weeks ago

Selected Answer: D

D is correct:

Portchannel mode VDS

- LACP Enabled:

> LACP Active/Passive - Route based on IP hash (downlink port group)

- LACP Disabled

> MAC Pinning - • Route based on originating virtual port

> MACPinning-Physical-NIC-Load - Route based on physical NIC load

> Static Channel ON - Route based on IP hash

upvoted 1 times

jecq 2 years, 6 months ago

Selected Answer: D



<https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/6x/virtualization/cisco-aci-virtualization-guide-60x/ACI-Virtualization-Guide-60x-aci-with-vmware-vds.pdf>

upvoted 3 times

When configuring Cisco ACI VMM domain integration with VMware vCenter, which object is created in vCenter?

- A. datacenter
- B. VMware vSphere Standard vSwitch
- C. VMware vSphere Distributed Switch
- D. cluster

Correct Answer: C

  **Narbleddeath** 8 months, 2 weeks ago

C.

Creating a vCenter Domain Profile Using the Basic GUI

Before You Begin

Before you create a VMM domain profile, you must establish connectivity to external network using in-band management network on the APIC.

Procedure

Step 1 Login to the Basic Mode in the APIC GUI.

Step 2 On the menu bar, choose VM NETWORKING > Inventory.

Step 3 In the Navigation pane, right-click VMware and click Create vCenter Domain.

Step 4 In the Create vCenter Domain dialog box, in the Virtual Switch Name field, enter a Name.

Step 5 In the Virtual Switch field, verify that VMware vSphere Distributed Switch is selected

https://www.cisco.com/c/dam/m/en_us/solutions/data-center-virtualization/application-centric-infrastructure/aci-virtualization-guide-chapter.pdf
upvoted 2 times

DRAG DROP -

Drag and drop the Cisco ACI Layer 4 to Layer 7 service insertion terms on the left to the correct descriptions on the right.

Select and Place:

concrete interfaces	ensures reachability between L3 domains
service graph	rendered with local resources that are available in the fabric
device cluster	contains an active-standby pair of firewalls or load balancers
VRF stitching	encapsulations programmed based on their association with logical interfaces

Correct Answer:

concrete interfaces	VRF stitching
service graph	service graph
device cluster	device cluster
VRF stitching	concrete interfaces

 **Huberttheman** 11 months ago

concrete interfaces: encapsulations programmed based on their association with logical interfaces

service graph: rendered with local resources that are available in the fabric

device cluster: contains an active-standby pair of firewalls or load balancers

VRF stitching: ensures reachability between L3 domains

upvoted 2 times

An engineer has set the VMM resolution immediacy to pre-provision in a Cisco ACI environment. No Cisco Discovery Protocol neighborhood has been formed between the hypervisors and the ACI fabric leaf nodes. How does this affect the download policies to the leaf switches?

- A. No policies are downloaded because LLDP is the only supported discovery protocol.
- B. Policies are downloaded when the hypervisor host is connected to the VMM VDS.
- C. Policies are downloaded to the ACI leaf switch regardless of Cisco Discovery Protocol neighborhood.
- D. No policies are downloaded because there is no discovery protocol neighborhood.

Correct Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **MaxG** Highly Voted 3 years, 11 months ago

Definitely C

When using pre-provision immediacy, policy is downloaded to ACI leaf switch regardless of CDP/LLDP neighborhood. Even without a hypervisor host connected to the VMM switch.

upvoted 9 times

🗳️ 👤 **JJCR86** Highly Voted 3 years, 11 months ago

Correct answer is C.

Pre-provision—Specifies that a policy (for example, VLAN, VXLAN binding, contracts, or filters) is downloaded to a leaf switch even before a VM controller is attached to the virtual switch (for example, VMware VDS). This pre-provisions the configuration on the switch.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_01011.html

upvoted 7 times

🗳️ 👤 **udo2020** Most Recent 6 months, 3 weeks ago

Selected Answer: C

I think it is C. But there is one more thing to know:

Only leaf switches with access policies referencing an AAEP associated with a VMM domain download policies as a result of EPG pre-provisioning.

upvoted 1 times

🗳️ 👤 **designated** 7 months, 2 weeks ago

Selected Answer: C

C is correct

> Pre-provision: VLAN will be deployed on all leaf interfaces under the AAEP associated to the VMM domain regardless of VM controller of hypervisor status.

> Immediate: VLAN will be deployed on leaf interfaces only when hypervisors are detected through LLDP or CDP. This information has to be bi-direction.

> On Demand: VLAN will be deployed on leaf interfaces only when hypervisors are detected as mentioned in Immediate mode and when at least one VM is associated to the corresponding port group. Both conditions need to be met for the VLAN to be deployed on leaf interfaces in On-Demand mode.

upvoted 1 times

🗳️ 👤 **jecq** 2 years, 6 months ago

I think that the correct answer is D, because in question say "No Cisco Discovery Protocol neighborhood has been formed".

upvoted 1 times

🗳️ 👤 **jecq** 2 years, 6 months ago

Sorry, the correct answer is C, I got confused with the "immediate" option.

upvoted 1 times

🗨️ 👤 **korthab** 2 years, 10 months ago

Selected Answer: C

The answer is C:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_01011.html

Pre-provision—Specifies that a policy (for example, VLAN, VXLAN binding, contracts, or filters) is downloaded to a leaf switch even before a VM controller is attached to the virtual switch (for example, VMware VDS). This pre-provisions the configuration on the switch.

..
..
..

When using pre-provision immediacy, policy is downloaded to ACI leaf switch regardless of CDP/LLDP neighborship. Even without a hypervisor host connected to the VMM switch.

upvoted 1 times

🗨️ 👤 **erikhara** 4 years, 4 months ago

It should be . B

When using pre-provision immediacy, policy is downloaded to ACI leaf switch regardless of CDP/LLDP neighborship. Even without a hypervisor host connected to the VMM switch.

upvoted 3 times

🗨️ 👤 **erikhara** 4 years, 4 months ago

I think its . A

When using pre-provision immediacy, policy is downloaded to ACI leaf switch regardless of CDP/LLDP neighborship. Even without a hypervisor host connected to the VMM switch.

upvoted 1 times

🗨️ 👤 **hybersat** 4 years, 5 months ago

This must be D. Both CDP and LLDP are supported, but the immediate resolution is selected that does require CDP/LLDP. Only pre-provisioned does not require CDP/LLDP.

here is a snip from the docs (https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_01011.html)

Resolution Immediacy

Immediate—Specifies that EPG policies (including contracts and filters) are downloaded to the associated leaf switch software upon ESXi host attachment to a DVS. LLDP or OpFlex permissions are used to resolve the VM controller to leaf node attachments.

The policy will be downloaded to leaf when you add host to the VMM switch. CDP/LLDP neighborship from host to leaf is required.

upvoted 1 times

🗨️ 👤 **hybersat** 4 years, 5 months ago

Shoot... I was wrong.. it is C..

Sorry..

upvoted 8 times

In the context of VMM, which protocol between ACI leaf and compute hosts ensures that the policies are pushed to the leaf switches for immediate and on demand resolution immediacy?

- A. VXLAN
- B. LLDP
- C. ISIS
- D. STP

Correct Answer: B

Community vote distribution

B (100%)

🗲️ 👤 **[Removed]** Highly Voted 3 years, 11 months ago

Correct Answer: B .

Immediate—Specifies that EPG policies (including contracts and filters) are downloaded to the associated leaf switch software upon ESXi host attachment to a DVS. LLDP or OpFlex permissions are used to resolve the VM controller to leaf node attachments.

The policy will be downloaded to leaf when you add host to the VMM switch. CDP/LLDP neighborship from host to leaf is required.

On Demand—Specifies that a policy (for example, VLAN, VXLAN bindings, contracts, or filters) is pushed to the leaf node only when an ESXi host is attached to a DVS and a VM is placed in the port group (EPG).

The policy will be downloaded to leaf when host is added to VMM switch and virtual machine needs to be placed into port group (EPG). CDP/LLDP neighborship from host to leaf is required.

With both immediate and on demand, if host and leaf lose LLDP/CDP neighborship the policies are removed.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_01011.html

upvoted 16 times

🗲️ 👤 **sailorsoul** Most Recent 11 months, 1 week ago

Selected Answer: B

No doubt it is B

upvoted 1 times

🗲️ 👤 **Redou2201** 1 year, 4 months ago

Selected Answer: B

it is LLDP

upvoted 1 times

🗲️ 👤 **Gab99** 2 years, 2 months ago

Selected Answer: B

makes more sense

upvoted 1 times

🗲️ 👤 **Nickname456** 2 years, 3 months ago

Correct should be B

upvoted 1 times

🗲️ 👤 **mdsuresh** 2 years, 10 months ago



B is correct

upvoted 1 times

🗲️ 👤 **thomyohan** 3 years, 6 months ago

Correct Answer : B

upvoted 2 times

  **thiyagas** 3 years, 11 months ago

Protocol between Leaf and Compute used is lldp correct.? not sure if the vxlan is the answer

upvoted 2 times

Which two components are essential parts of a Cisco ACI Virtual Machine Manager (VMM) domain policy configuration? (Choose two.)

- A. Layer 3 outside interface association
- B. EPG static port binding
- C. VMM domain profile
- D. EPG association
- E. IP address pool association

Correct Answer: *CD*

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_01011.html

  **designated** 7 months, 2 weeks ago

Selected Answer: *CD*

VMM Main Components:

- > Virtual Machine Manager Domain Profile (C)
 - > Credential
 - > Controller
 - > EPG Association (D)
 - > Attachable Entity Profile Association (AEP)
 - > VLAN Pool Association
- upvoted 2 times

  **Marinheiro** 1 year, 1 month ago

That's correct, C and D.

upvoted 1 times

Domain - F1-VCSAB1_VCD

Policy Operational Associated EPGs

General VSwitch Policy Faults History

Properties

Port Channel Policy: F1-VCSAB1_VCD_lacpLag

LLDP Policy: F1-VCSAB1_VCD_lldplfPo

CDP Policy: select an option

MTU Policy: select an option

STP Policy: select an option

Firewall Policy: select an option

NetFlow Exporter Policy: select an option

Enhanced Lag Policy

Name	Mode	Load Balancing Mode	Number of Links
	LACP Active	Source and Destination IP Address	2


Update Cancel

Refer to the exhibit. An engineer configures the Cisco ACI fabric for VMM integration with ESXi servers that are to be connected to the ACI leaves. The server team requires the network switches to initiate the LACP negotiation as opposed to the servers. The LAG group consists of two 10 Gigabit Ethernet links. The server team also wants to evenly distribute traffic across all available links. Which two enhanced LAG policies meet these requirements? (Choose two.)

- A. LACP Mode: LACP Standby
- B. LB Mode: Destination IP Address and TCP/UDP Port
- C. LB Mode: Source and Destination MAC Address
- D. LB Mode: Source IP Address and TCP/UDP Port
- E. LACP Mode: LACP Active

Correct Answer: BD

Community vote distribution



Rododendron2 Highly Voted 2 years, 2 months ago

Selected Answer: CE

CE.

Wrong comments here, unclear why. 2 part to setup LACP & LB algorithm.

LACP needs to be active, unclear why omitted here in discussion by someone, it is enhanced LAG policy, so part of question. Cannot be source IP (TEP is not changing), cannot be dst port - same for common services, cannot be dest IP - same destination VM. So fits source & dst mac.

upvoted 5 times

designated Most Recent 7 months, 2 weeks ago

Selected Answer: BE

It should be

B and E..

We need to select a mode and this is LACP Active and we can select just one load balancing hash when we are creating a Port Channel Policy for VDS and since the vCenter has a lot of IPs inside it, the best choice it would be IP destination. Also, the source IP or MAC will be the same.

upvoted 2 times

mdrira 10 months, 4 weeks ago

Selected Answer: BD

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/aci_virtual_edge/configuration/2-x/Cisco-ACI-Virtual-Edge-Configuration-Guide-202/Cisco-ACI-Virtual-Edge-Configuration-Guide-202_chapter_0100.html

upvoted 1 times

🗨️ 👤 **netguy85** 11 months, 1 week ago

Selected Answer: DE

i think it is D and E

upvoted 1 times

🗨️ 👤 **Marinheiro** 1 year, 1 month ago

Selected Answer: BD

BD,

Cisco Application Centric Infrastructure Virtual Edge, VXLAN mode traffic always uses the source IP address as the TEP IP address. To ensure proper load balancing, we recommend the algorithm Source and Destination TCP/UDP Port.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/aci_virtual_edge/configuration/2-x/Cisco-ACI-Virtual-Edge-Configuration-Guide-202/Cisco-ACI-Virtual-Edge-Configuration-Guide-202_chapter_0100.html

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: AB

<https://aci-lab.ciscolive.com/lab/pod10/acivmm/vmm-build>

LACP Standby will give chance to the Leaf switches to initiate LACP BPDUs, Loadbalancing will be more variable to many destinations (More random flows).

upvoted 1 times

🗨️ 👤 **zelya19** 1 year, 3 months ago

Selected Answer: AB

This is a vSwitch policy and they want network switches to initiate LACP negotiations -> LACP Standby. Traffic is flowing from a limited number of VMs to the rest of the network -> LB Mode: Destination IP/Port. LB Mode can be different in the opposite direction!

upvoted 3 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

This is the correct answer.

The configs are will be configured on the VDS and face back to the ACI Leaves.

<https://aci-lab.ciscolive.com/lab/pod10/acivmm/vmm-build>

upvoted 1 times

🗨️ 👤 **Redou2201** 1 year, 8 months ago

for me it is C and E based on this document:

<https://community.cisco.com/t5/data-center-and-cloud-knowledge-base/vmm-integration-with-enhanced-lacp/ta-p/4522997>

upvoted 2 times

🗨️ 👤 **Mr_Certifiable** 2 years, 1 month ago

possibly CE

consider vmotion - The LAG group consists of two 10 Gigabit Ethernet links

Cisco ACI considers the frequent move of an IP address from one MAC address to the other and potentially between ports as a misconfiguration. Features such as rogue endpoint control may quarantine the endpoints and raise a fault.

<https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-application-centric-infrastructure-design-guide.html>

upvoted 2 times

🗨️ 👤 **NSF2** 2 years, 4 months ago

In my opinion, it must be D and E.

Because usually source port is ephemeral port which changes all the time, as result deifferent hashes that put traffic across multiple links.


Secondly LACP must be ACTIVE.

upvoted 2 times

🗨️ 👤 **Gab99** 2 years, 8 months ago

Selected Answer: BD

complicated but I agree
upvoted 3 times

  **7korn7** 2 years, 9 months ago

Selected Answer: BD

The answer is B and D.

Explanation:

LACP Mode: LACP Standby - In LACP available only two modes - Active and Passive

LACP Mode: Active its standard LACP mode

LB Mode: Source and Destination MAC Address its a standard LACP balancing mode

And question says: Which two enhanced LAG policies meet these requirements? (Choose two.)

LB Mode: Destination IP Address and TCP/UDP Port

LB Mode: Source IP Address and TCP/UDP Port

This two modes are LACP enhanced balancing modes

upvoted 4 times

  **Drunken_Commie** 3 years, 1 month ago

everyone forgets that we're configuring the "remote" side of connection (DVS on servers). ergo, it has to be LACP Passive (here called "standby"), and best-from-provided load balancing decision would be destination ip + tcp/udp

upvoted 2 times

  **muhnator** 3 years, 4 months ago

We dont know if its a "normal VDS" or an "AVE". So maybe they want to point to the AVE due to

(https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/aci_virtual_edge/configuration/2-x/Cisco-ACI-Virtual-Edge-Configuration-Guide-202/Cisco-ACI-Virtual-Edge-Configuration-Guide-202_chapter_0100.html)

so because "For Cisco ACI Virtual Edge VXLAN mode, it is mandatory to use a load-balancing algorithm having a UDP port. We recommend the algorithm Source and Destination TCP/UDP Port."

B and E should be correct



upvoted 2 times

  **duracell** 3 years, 8 months ago

I think the answers are not complete. --> answer should be LACP active / Src/Dst TCP/UDP

Possible answers: SRC/DST MAC, SRC/DST IP, SRC/DST TCP/UDP and LACP active or passive

upvoted 1 times

  **duracell** 3 years, 8 months ago

Cisco Application Centric Infrastructure Virtual Edge, VXLAN mode traffic always uses the source IP address as the TEP IP address. To ensure proper load balancing, we recommend the algorithm Source and Destination TCP/UDP Port.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/aci_virtual_edge/configuration/2-x/Cisco-ACI-Virtual-Edge-Configuration-Guide-202/Cisco-ACI-Virtual-Edge-Configuration-Guide-202_chapter_0100.html

upvoted 2 times

  **apot** 3 years, 9 months ago

I think C and E.

I think that B and D are not correct as there are no such options

upvoted 2 times

  **mvfpeof05l** 3 years, 11 months ago

there is no mode of lacp standby in the APIC settings. LACP Mode: LACP Standby - is wrong!

In the Mode field, choose one of the following options appropriate to your setup:

Static Channel - Mode On

LACP Active

LACP Passive

MAC Pinning

MAC Pinning-Physical-NIC-load

but the answer options clearly state - LACP Standby.

upvoted 1 times

Which tenant is used when configuring in-band management IP addresses for Cisco APICs, leaf nodes, and spine nodes?

- A. default
- B. infra
- C. common
- D. mgmt

Correct Answer: D

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_Configuring_Static_Management_Access.html#concept_CFF63FEBE947424291B0F10E6F23DA7D

Community vote distribution

D (100%)

🗳️ **designated** 7 months, 2 weeks ago

Selected Answer: D

Management IP addresses = mgmt
upvoted 1 times

🗳️ **Marinheiro** 1 year, 1 month ago

Selected Answer: D

D

ACI fabric nodes have two options for management connectivity; out-of-band (OOB), which governs the dedicated physical management port on the back of the device, or in-band (INB), which is provisioned using a specific EPG/BD/VRF in the management tenant with a degree of configurable parameters.

<https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/218028-troubleshoot-aci-management-and-core-ser.html>

upvoted 1 times

🗳️ **duracell** 3 years, 8 months ago

tenant mgmt is used for in- and out-of-band-mgmt
upvoted 2 times

🗳️ **mr_siro** 3 years, 10 months ago

D is correct

guide config in-band using apic gui

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/basic-configuration/Cisco-APIC-Basic-Configuration-Guide-411/Cisco-APIC-Basic-Configuration-Guide-411_chapter_0100.html#task_730EB317BDE94145BAF91092F88B5257

upvoted 4 times

🗳️ **wcorrea** 3 years, 10 months ago

B is correct. APIC connects to the Out of band (OOB) management net via dedicated mgmt0.

APIC also connect inband to the leafs via tenant infra vrf overlay-1

upvoted 2 times

What represents the unique identifier of an ACI object?

- A. universal resource identifier (URI)
- B. application programming interface
- C. management information tree
- D. distinguished name

Correct Answer: D

Reference:

<https://www.slideshare.net/CiscoDevNet/introduction-to-aci-apis>

Community vote distribution

D (100%)

🗳️ 👤 **Marinheiro** 7 months, 3 weeks ago

Selected Answer: D

D

The ACI object model is represented in the Management Information Tree (MIT) that is an object oriented data base. Every branch represents a functional area, and every node is a managed object - has a CLASS and a globally unique Distinguished Name formed by a parent name and the relative name.

<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/policy-model-guide/b-Cisco-ACI-Policy-Model-Guide.html>

upvoted 2 times

🗳️ 👤 **Rododendron2** 1 year, 8 months ago

D. Every MO in the system can be identified by a unique distinguished name (DN)

upvoted 1 times

🗳️ 👤 **Narbledeath** 1 year, 8 months ago

Correct, D

<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/policy-model-guide/b-Cisco-ACI-Policy-Model-Guide.html>

upvoted 2 times

Which new construct must a user create when configuring in-band management?

- A. VLAN pool
- B. management contract
- C. management tenant
- D. bridge domain

Correct Answer: B

Community vote distribution

B (55%)

A (45%)

 **nikomski**  4 years, 5 months ago

Shouldn't this be A - VLAN pool?


upvoted 8 times

 **memelas** 4 years, 4 months ago

A is correct, since Tenant and Bridge Domain are present by default on the GUI.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/basic-configuration/Cisco-APIC-Basic-Configuration-Guide-411/Cisco-APIC-Basic-Configuration-Guide-411_chapter_0100.html#task_730EB317BDE94145BAF91092F88B5257

upvoted 13 times

 **nabilzay** 4 years, 4 months ago

Agree, should be A as for the rest of the constructs a default one should already exist that you can use

upvoted 6 times

 **manet**  4 years ago

Shouldn't be B (management contract), since VLAN Pool is not created. Is it assigned only one VLAN ID when in-band mgmt EPG is created.

upvoted 5 times

 **JJCR86** 3 years, 11 months ago

Correct answer should be B. Both a contract and VLAN pool should be created however the questions specifies "construct" and the only logical construct among VLAN pool and contracts is the latter.

upvoted 4 times

 **udo2020**  6 months, 3 weeks ago

Selected Answer: B

B is the only possible solution.

upvoted 1 times

 **designated** 7 months, 2 weeks ago

Selected Answer: B

B is correct since contract is a logical construct.

Cisco ACI Logical Constructs:

Tenant

VRF

Bridge Domain

EPG

Application Profile

Contract

Reference: Cisco U

upvoted 2 times

 **Lalag** 11 months, 3 weeks ago

A should be the right answer, the contract required for external communication of In-band,

upvoted 1 times

🗨️ 👤 **zelya19** 1 year, 3 months ago

Selected Answer: B

A contract is a policy construct used to define communication between EPGs.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-743951.html>

upvoted 1 times

🗨️ 👤 **Huberttheman** 1 year, 5 months ago

Selected Answer: A

In Cisco ACI, the typical components involved in configuring in-band management include a management tenant, management VRF, bridge domain for management, VLAN pool, and the assignment of IP addresses to the in-band interfaces of the fabric devices (APICs, leaf nodes, and spine nodes).

A "management contract" is more commonly associated with defining the communication rules or policies for managing the ACI fabric, but it is not directly related to the creation of the necessary constructs for in-band management

upvoted 1 times

🗨️ 👤 **959836c** 1 year, 6 months ago

Selected Answer: A

It's physical domain then vlan pool

upvoted 1 times

🗨️ 👤 **Macc10** 1 year, 7 months ago

Selected Answer: B

This is not D, Correct answer is B.

Bridge domain "inb" is pre-defined within the management tenant.

As per Red Nectar Article - "I'll assign the default gateway IP address to the pre-defined inb Bridge Domain in the mgmt Tenant".

<https://rednectar.net/2016/12/22/configuring-in-band-management-for-the-apic-on-cisco-aci-part-1-via-an-epg/>

I can confirm this is correct within own lab environment.

upvoted 4 times

🗨️ 👤 **LevisRock** 1 year, 9 months ago

Selected Answer: B

should be B

upvoted 2 times

🗨️ 👤 **rss_01** 2 years, 6 months ago

Vlan Pool is not a construct ,but still required.So i'd say B

upvoted 2 times

🗨️ 👤 **Smoothery** 3 years, 2 months ago

Selected Answer: A

inb (Bridge domain pre existing). No Vlan pool and ideally one of the first steps required. MGMT tenant exists. Management contract is the last "to do". Construct literally means build blocks. Hence nothing can be built with thr VLANS.

upvoted 3 times

🗨️ 👤 **mr_siro** 3 years, 10 months ago

D is correct,

pls refer link: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/basic-configuration/Cisco-APIC-Basic-Configuration-Guide-411/Cisco-APIC-Basic-Configuration-Guide-411_chapter_0100.html#task_730EB317BDE94145BAF91092F88B5257

upvoted 1 times

🗨️ 👤 **joe1989** 3 years, 9 months ago

in the same link you have provided, watch closely it does Not mention "Create Bridge Domain", it only says expand the bridge domain(which must mean it is pre-existing). please correct me if my observation is incorrect...

upvoted 2 times

🗨️ 👤 **igebuntu** 3 years, 4 months ago

Built-in BD name is "inb" in the mgmt tenant. No need to create.

upvoted 4 times

🗨️ 👤 **mvfpeof05l** 3 years, 11 months ago

Why is it not a bridge domain, it also needs to be created?

upvoted 1 times

What must be configured to allow SNMP traffic on the APIC controller?

- A. out-of-band management interface
- B. contract under tenant mgmt
- C. SNMP relay policy
- D. out-of-band bridge domain

Correct Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **designated** 7 months, 2 weeks ago

Selected Answer: B

B is correct

upvoted 2 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: B

Answer is B

upvoted 1 times

🗳️ 👤 **igebuntu** 3 years, 4 months ago

As a matter of fact the correct answer would be A,B - but since OOB interface is configured during the APIC initial setup, it may be less relevant answer here.

upvoted 4 times

Which type of port is used for in-band management within ACI fabric?

- A. spine switch port
- B. APIC console port
- C. leaf access port
- D. management port

Correct Answer: C

  **MaxG**  11 months ago

C is correct.

<https://howtoaci.com/2019/06/15/configuring-in-band-management-connectivity-in-aci/>

upvoted 5 times

  **Ravindrakedar**  11 months, 1 week ago

Answer should be B: APIC Console port

upvoted 1 times

  **sourabh1000** 10 months ago

no, its inband magmt.. not oob so answer c is correct

upvoted 1 times

Edit Stats Threshold

transmit B2B credit cumulative

Normal Value: 300

Threshold Direction: Both Rising Falling

Rising Thresholds to Config: ☒ ☐

- ☒ Critical
- ☐ Major
- ☐ Minor
- ☒ Warning

Rising

	Set	Reset
Critical	500	
Major		
Minor		
Warning	400	300

Refer to the exhibit. A client reports that the ACI domain connectivity to the fiber channel storage is experiencing a B2B credit oversubscription. The environment has a SYSLOG server for state collection messages. Which value should be chosen to clear the critical fault?

- A. 300
- B. 410
- C. 350
- D. 510

Correct Answer: B

Community vote distribution

A (100%)

Dspam Highly Voted 3 years, 10 months ago

critical value range is 500 - X

warning range is 400 - 300 .

options 300 and 350 cannot be used as it will overlap into the warning range.

510 is beyond the 500 range of critical logical deduction .. on 410 - 500 will create a unique value range that will uniquely trigger critical

upvoted 11 times

designated Most Recent 7 months, 2 weeks ago



Selected Answer: B

It should some value between 499 and 401 and there is only 410.

Above of 500 is Critical

Below of 400 is Warning

upvoted 2 times



  **Rollizo** 8 months, 1 week ago

Selected Answer: B

410 is the only value. Critical has to reset to a value less than 500.

Cannot be 300 or 350 because it is the same as warning


upvoted 2 times

  **mdrira** 10 months, 3 weeks ago

Selected Answer: A

Since the normal value is 300, clearing fault should be 300

upvoted 1 times

  **Ravindrakedar** 3 years, 11 months ago

Not sure how it comes as 410, can anyone explain same here

upvoted 1 times

Which statement about ACI syslog is true?

- A. Notifications for different scopes of syslog objects can be sent only to one destination.
- B. Syslog messages are sent to the destination through the spine.
- C. All syslog messages are sent to the destination through APIC.
- D. Switches send syslog messages directly to the destinations.

Correct Answer: D

Community vote distribution



D (100%)

  **[Removed]**  3 years, 11 months ago

Correct Answer:D .Please refer below URL in page:66

<https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2016/pdf/BRKACI-2303.pdf>

upvoted 15 times

  **Carol254** 3 years, 10 months ago

Syslog

BRKACI-2303 66

- MOs with associated faults or stats have a scope

- Notifications for different scopes can be sent to different destinations

- Faults, Event Records and Audit

Records can be dispatched using syslog, as well as callhome & SNMP traps

- Switches send syslog message directly to the destinations, APIC is not involved in forwarding switch syslog messages.

upvoted 14 times

  **thiyagas**  3 years, 11 months ago

Correct Answer is D

upvoted 9 times

  **msalamehi**  11 months ago

correct answer is D

A doesn't make sense

upvoted 1 times

  **asd248402** 1 year, 6 months ago

Selected Answer: D

switch send logs directly

upvoted 1 times

  **Rododendron2** 1 year, 8 months ago

Selected Answer: D

Obviously - only thing that makes sense -Switches send syslog messages directly to the destinations.

upvoted 1 times

  **Gab99** 2 years, 2 months ago

Selected Answer: D

D is clear based explanation here

upvoted 1 times

A data center administrator is upgrading an ACI fabric. There are 3 APIC controllers in the fabric and all the servers are dual-homed to pairs of leaf switches configured in VPC mode. How should the fabric be upgraded to minimize possible traffic impact during the upgrade?

- A. 1. Create two maintenance groups for the APIC controllers: VPC left and VPC right. 2. Upgrade the first group of controllers. 3. Upgrade the second group of controllers. 4. Upgrade the leaf switches.
- B. 1. Create two maintenance groups for APIC controllers: VPC left and VPC right. 2. Upgrade the leaf switches. 3. Upgrade the first group of controllers. 4. Upgrade the second group of controllers.
- C. 1. Create two maintenance groups for the leaf switches: VPC left and VPC right. 2. Upgrade the APIC controllers. 3. Upgrade the first group of leaf switches. 4. Upgrade the second group of leaf switches.
- D. 1. Create two maintenance groups for the leaf switches: VPC left and VPC right. 2. Upgrade the first group of switches. 3. Upgrade the second group of switches. 4. Upgrade the APIC controllers.

Correct Answer: C

Community vote distribution

C (100%)

[Removed] 4 years, 5 months ago

Correct Answer : C .APIC controller should be upgraded first and then all switches with even and odd Group .
upvoted 25 times

mrpeet 4 years ago

Correct answer is C

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/apic-installation-upgrade-downgrade/Cisco-APIC-Installation-Upgrade-Downgrade-Guide/m_upgrading_and_downgrading_the_apic_controller_and_switch_software.html#Cisco_Reference.dita_1b95beff-a430-4236-9258-cb38cc3c7d7f
upvoted 7 times

designated 7 months, 2 weeks ago

Selected Answer: C

C is correct

APIC > One side VPC (left) > other side VPC (right)
upvoted 1 times

Jhony_Bless 10 months, 1 week ago

Selected Answer: C

Apics first, the the 2 groups of switches.
upvoted 1 times

Macc10 1 year, 8 months ago

C is the Answer - Controller needs to be upgraded first prior 2 switches.
upvoted 1 times

asd248402 2 years ago

C as apic are the first that needs to be upgraded
upvoted 1 times

Gab99 2 years, 8 months ago

Selected Answer: C

C of course!
upvoted 2 times

korthab 2 years, 11 months ago

Selected Answer: C

The correct answer is C. Recommended by Cisco TAC and the Cisco Consultant that was helping us upgrade ACI to upgrade the APIC first and than the Fabric with Minimum of two groups VPC1 and VPC2.
upvoted 2 times

jmaroto 3 years ago

Selected Answer: C

C is correct answer

upvoted 1 times

🗉 👤 **mdsuresh** 3 years, 3 months ago

C is correct answer

upvoted 1 times

🗉 👤 **mdsuresh** 3 years, 4 months ago

Correct Answer is C

upvoted 1 times

🗉 👤 **Leviatan_93** 3 years, 11 months ago

Correct answer C

upvoted 4 times

Which protocol does ACI use to securely save the configuration in a remote location?

- A. SCP
- B. HTTPS
- C. TFTP
- D. FTP

Correct Answer: A

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_Using_Import_Export_to_Recover_Config_States.html

  **thomyohan** Highly Voted 4 years ago

Answer : SCP (Both FTP and SCP are supported....but the keyword "secure")

upvoted 6 times

  **designated** Most Recent 7 months, 2 weeks ago

Selected Answer: A

We have SFTP now then this could be the second answer.

FTP is WRONG.

upvoted 1 times

  **designated** 7 months, 2 weeks ago

Selected Answer: A

SCP > secure transfer

upvoted 1 times

  **knorraple** 4 years ago

keyword is securly

upvoted 2 times

  **manet** 4 years ago

I think that also D (FTP) should be correct answer

upvoted 1 times

Which two protocols support accessing backup files on a remote location from the APIC? (Choose two.)

- A. TFTP
- B. FTP
- C. SFTP
- D. SMB
- E. HTTPS

Correct Answer: BC

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/troubleshooting/b_APIC_Troubleshooting/b_APIC_Troubleshooting_appendix_010011.html

  **designated** 7 months, 2 weeks ago

Selected Answer: BC

It is correct.. B and C.

Protocols supported by Remote Location: FTP, SCP, SFTP (default)
upvoted 1 times

  **imamus** 2 years, 1 month ago

I guess its those because the remote location allowed protocols in the apic are SCP, FTP, SFTP

<https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/aci-guide-creating-backup-for-apic-cluster.pdf>

upvoted 3 times

Which attribute should be configured for each user to enable RADIUS for external authentication in Cisco ACI?

- A. cisco-security domain
- B. cisco-auth-features
- C. cisco-aci-role
- D. cisco-av-pair

Correct Answer: D

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/Security_config/b_Cisco_APIC_Security_Configuration_Guide/b_Cisco_APIC_Security_Guide_chapter_01011.html

 **Rododendron2** 7 months ago

attribute - for each user ? what is that supposed to mean ?

upvoted 2 times

 **Mr_Certifiable** 7 months, 1 week ago

The Cisco APIC requires that an administrator configure a Cisco AV Pair on an external authentication server. The Cisco AV pair specifies the APIC required RBAC roles and privileges for the user. The Cisco AV Pair format is the same for RADIUS, LDAP, or TACACS+

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/basic-configuration/Cisco-APIC-Basic-Configuration-Guide-42x/Cisco-APIC-Basic-Configuration-Guide-42x_chapter_011.html

upvoted 2 times

```
aaa authentication login fallback
  realm radius
  group radius-1

aaa authentication login console
  realm radius
  group radius-1

aaa authentication login default
  realm radius
  group radius-1



aaa banner 'WELCOME TO ACI'
aaa group radius radius-1
  server 10.1.1.1 priority 0
  server 10.2.2.2 priority 1

aaa user default-role-no-login
```

Refer to the exhibit. Which action should be taken to ensure authentication if the RADIUS servers are unavailable?

- A. Adjust the priority of server 10.1.1.1 to 1.
- B. Assign the user to the default role.
- C. Set the default login realm to LDAP.
- D. Set the fallback login to local.

Correct Answer: D

  **prospio971** 7 months, 1 week ago

Selected Answer: D

To ensure authentication in case of RADIUS server unavailability, you should implement a fallback authentication mechanism. This involves configuring a local authentication source on the Cisco ACI devices.

upvoted 1 times

A network engineer demonstrates Cisco ACI to a customer. One of the test cases is to validate a disaster recovery event by resetting the ACI fabric to factory and then restoring the fabric to the state it was in before the event. Which setting must be enabled on ACI to export all configuration parameters that are necessary to meet these requirements?

- A. enabled AES encryption
- B. generated a tech-support file
- C. encrypted export destination
- D. enabled JSON format export

Correct Answer: A



  **nabilzay** Highly Voted 4 years, 4 months ago

Answer is A, without AES encrypt enabled password info is not backed up
upvoted 13 times

  **designated** Most Recent 7 months, 2 weeks ago

Selected Answer: A

Admin > Import/Export > Export Policies > Configuration > Modify Global AES Encryption Settings > Enabled!
upvoted 2 times

  **Dspam** 3 years, 10 months ago

Always enable AES encryption when performing fabric backup configuration exports. Doing so will assure that all the secure properties of the configuration will be successfully imported when restoring the fabric.
upvoted 3 times

  **mrpeet** 4 years ago



Answer is A. AES encryption may be optional for backup but is required for the use case of restoring a factory-fresh install to a fully working environment.
upvoted 2 times

  **Mohitkrsh84** 4 years, 4 months ago

isn't A and D both correct. A is optional though.

<https://unofficialaciguide.com/2017/10/21/creating-a-backup-for-your-apic-cluster/>

upvoted 1 times

  **MrSaint** 3 years, 6 months ago

D states that "enabled JSON format export" but it also is an option, since you can enable XML format to export as well. About A, yes AES is optional BUT if you dont enable it, the fabric wont export credentials like VMM or service pack for managed L4-L7 Services.
upvoted 2 times

An engineer wants to filter the System Faults page and view only the active faults that are present in the Cisco ACI fabric. Which two lifecycle stages must be selected for filtering? (Choose two.)

- A. Raised
- B. Retaining
- C. Soaking, Clearing
- D. Raised, Clearing
- E. Soaking

Correct Answer: AE

Community vote distribution

AE (100%)

BigD69 Highly Voted 3 years ago

A and E. Raised and Soaking.

Soaking = initial state of a fault, when problem condition first detected.

Raised = fault not resolved within soaking interval.

All other options are fault resolved for one reason or another.

upvoted 10 times

nikomski Highly Voted 3 years, 5 months ago

Shouldn't this one be A & E?

upvoted 7 times

Mohitkrsh84 3 years, 4 months ago

Yes, it should be A & E.

upvoted 9 times

mrpeet 3 years ago

No. Soaking means APIC isn't sure whether it's a fault or not.

upvoted 1 times

Dameel Most Recent 8 months, 3 weeks ago

Selected Answer: AE

Clearing is not used to see active faults

upvoted 1 times

Necha_uk 12 months ago

Selected Answer: AE

A and E. Raised and Soaking.

Soaking = initial state of a fault, when problem condition first detected.

Raised = fault not resolved within soaking interval.

upvoted 2 times

muhnator 2 years, 4 months ago

Selected Answer: AE

A and E is correct

upvoted 3 times

Xixmk 2 years, 8 months ago

A and E.

D - Raised-Clearing means that fault trigger is no longer active, fault remains in that state until timer expires or trigger activates again, then it goes directly to Raised state, without previously going through Soaking. A, E are correct since both represent a fault with active trigger.

upvoted 3 times

An engineer must limit management access to the Cisco ACI fabric that originates from a single subnet where the NOC operates. Access should be limited to SSH and HTTPS only. Where should the policy be configured on the Cisco APIC to meet the requirements?

- A. policy in the management tenant
- B. ACL on the console interface
- C. ACL on the management interface of the APIC
- D. policy on the management VLAN

Correct Answer: A

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/cli/nx/cfg/b_APIC_NXOS_CLI_User_Guide/b_APIC_NXOS_CLI_User_Guide_chapter_01000.html

Community vote distribution

A (100%)

🗨️ 👤 **Marinheiro** 7 months, 3 weeks ago

Selected Answer: A

This document left me a little confused:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/cli/nx/cfg/b_APIC_NXOS_CLI_User_Guide/b_APIC_NXOS_CLI_User_Guide_chapter_01000.pdf

However, I tested it in the lab and I'm going with option A.

upvoted 1 times

In the context of ACI Multi-Site, when is the information of an endpoint (MAC/IP) that belongs to site 1 advertised to site 2 using the EVPN control plane?

- A. Endpoint information is not exchanged across sites unless COOP protocol is used.
- B. Endpoint information is not exchanged across sites unless a policy is configured to allow communication across sites.
- C. Endpoint information is exchanged across sites as soon as the endpoint is discovered in one site.
- D. Endpoint information is exchanged across sites when the endpoints are discovered in both sites.

Correct Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **thiyagas** Highly Voted 3 years, 11 months ago

correct answer is B

The endpoint information is stored in the local COOP database. Spine nodes in site 1 know about locally connected endpoints, and the same is true for spine nodes in site 2. Note that at this point no information is exchanged across sites for EP1 and EP2 EPGs, because there is no policy in place yet indicating a need for those endpoints to communicate.

An intersite policy is defined in Cisco ACI Multi-Site Orchestrator and is then pushed and rendered in the two sites.

upvoted 17 times

🗳️ 👤 **thomyohan** Highly Voted 3 years, 6 months ago

A COOP notification is generated inside each fabric from the leaf nodes on which EP1 and EP2 are discovered and sent to the local spine nodes.

The endpoint information is stored in the local COOP database. Spine nodes in site 1 know about locally connected endpoints, and the same is true for spine nodes in site 2. Note that at this point no information is exchanged across sites for EP1 and EP2 EPGs, because there is no policy in place yet indicating a need for those endpoints to communicate. Correct Answer : B

upvoted 11 times

🗳️ 👤 **sailorsoul** Most Recent 11 months ago

Selected Answer: B

voting B as everybody explain.

upvoted 1 times

🗳️ 👤 **hebdeb** 1 year, 5 months ago

Selected Answer: B

It is B

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html#CiscoACIMultiSiteoverlaycontrolplane>

upvoted 2 times

🗳️ 👤 **Mr_Certifiable** 1 year, 7 months ago

B

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html#CiscoACIMultiSiteoverlaydataplane>

The endpoint information is stored in the local COOP database. Spine nodes in site 1 know about locally connected endpoints, and the same is true for spine nodes in site 2. Note that at this point no information is exchanged across sites for EP1 and EP2 EPGs, because there is no policy in place yet indicating a need for those endpoints to communicate.

upvoted 2 times

🗳️ 👤 **Gab99** 2 years, 2 months ago

Selected Answer: B

B as they explained

upvoted 1 times

🗳️ 👤 **Xixmk** 3 years, 2 months ago

No Endpoint information exchange takes place unless a contract is applied or EPG is being stretched across sites
upvoted 2 times

🗨️ 👤 **[Removed]** 3 years, 11 months ago

Correct Answer: B .Because COOP only works in local fabric within the site.
upvoted 5 times

🗨️ 👤 **hybersat** 3 years, 11 months ago

I think B is correct.

Coop are always used on spines, but not between multi sites. If traffic is allowed one spine in on site will forward the BUM packet to the remote site spines using anycast via EVPN and it will lookup like normal there. If there is a host it will return the answer back via EVPN and it's put in the local coop db. That is at least how I read the docs :-)
upvoted 4 times

Which statement regarding ACI Multi-Pod and TEP pool is true?

- A. The IP addresses used in the IPN network can overlap TEP pool of the APIC.
- B. A different TEP pool must be assigned to each Pod.
- C. The Pod1 TEP pool must be split and a portion of the TEP pool allocated to each Pod.
- D. The same TEP pool is used in all Pods.

Correct Answer: B

Reference:

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739714.html>

 **thomyohan** Highly Voted 1 year ago

In a Cisco ACI Multi-Pod deployment, each Pod is assigned a separate and not overlapping TEP pool Answer : B
upvoted 8 times

Which two statements regarding ACI Multi-Site are true? (Choose two.)

- A. The Multi-Site orchestrator must be directly attached to one ACI leaf.
- B. Routers in the Inter-Site network must run OSPF, DHCP relay, and MP-BGP.
- C. ACI Multi-Site is a solution that supports a dedicated APIC cluster per site.
- D. ACI Multi-Site is a solution that allows one APIC cluster to manage multiple ACI sites.
- E. The Inter-Site network routers should run OSPF to establish peering with the spines.

Correct Answer: CE

Community vote distribution

CE (100%)

[Removed] Highly Voted 3 years, 11 months ago

Correct Answer :CE .Cisco ACI multi-sites need separate APIC cluster in each site .
upvoted 13 times

thiyagas Highly Voted 3 years, 11 months ago

Correct Answer is CE
upvoted 9 times

sailorsoul Most Recent 11 months, 1 week ago

Selected Answer: CE

C E. multi-sites need separate APIC cluster in each site .
upvoted 1 times

Gab99 2 years, 2 months ago

Selected Answer: CE

yes true
upvoted 1 times

korthab 2 years, 4 months ago

Selected Answer: CE

The correct answer is C, E
upvoted 3 times

eddyedwards257 2 years, 5 months ago

MSO can be connected to anywhere as long as it as out of band reachability to the APIC's
upvoted 1 times

Xixmk 3 years, 2 months ago

CE, A is not correct. MSO can be connected in external network
upvoted 3 times

Leviatan_93 3 years, 5 months ago

Correct answer C E
upvoted 2 times

mrpeet 3 years, 6 months ago

Correct Answer: CE
The MSO connects to the APIC clusters via their OOB networks and requires separate APIC clusters per site.
upvoted 4 times

mrpeet 3 years, 6 months ago

Correct Answer: CE
The MSO connects to the APIC clusters via their OOB networks and requires separate APIC clusters per site.
upvoted 4 times

What are two requirements for the IPN network when implementing a Multi-Pod ACI fabric? (Choose two.)

- A. EIGRP routing
- B. PIM ASM multicast routing
- C. BGP routing
- D. VLAN ID 4
- E. OSPF routing

Correct Answer: AE

Community vote distribution

DE (100%)

 **hybersat** Highly Voted 4 years, 5 months ago

I believe DE is correct..

According to <https://haystacknetworks.com/cisco-aci-multi-pod-ipn-inter-pod-network-configuration-verification>


See these snips:

The links (interfaces) between the IPN devices and the spine switches must have the following OSPF configuration on the interfaces (as discussed, these are actually the sub-interfaces for VLAN-4).

.....

Cisco ACI requires\recommends (works best with) bi-dir multicast as we have many sources and many receivers

upvoted 21 times

 **hybersat** 4 years, 5 months ago

Also, it's PIM Bidir and not PIM ASM that are required for multi pod

upvoted 5 times

 **memelas** 4 years, 4 months ago

Correct - D and E

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html>

upvoted 4 times

 **prospio971** Most Recent 4 months, 2 weeks ago

Selected Answer: BE

When implementing a Cisco ACI Multi-Pod fabric, the Inter-Pod Network (IPN) must meet specific requirements to ensure proper communication and functionality across pods. These requirements include:

PIM ASM Multicast Routing:

Multicast is required in the IPN to handle BUM (Broadcast, Unknown Unicast, and Multicast) traffic across pods. Specifically, PIM ASM (Protocol Independent Multicast - Any Source Multicast) or PIM Bidirectional is used to replicate BUM traffic between the pods efficiently.

OSPF Routing:

The IPN must support OSPF (Open Shortest Path First) routing for establishing connectivity between spine nodes and IPN devices. OSPF is used to advertise the TEP (Tunnel Endpoint) pool prefixes between pods and ensure end-to-end IP reachability for VXLAN tunnels.

upvoted 1 times

 **Kastel** 7 months, 3 weeks ago

I believe BE is the correct answer.

upvoted 1 times

 **slogoheinzy** 1 year, 1 month ago

answer is DE - BUT as of 5.2, BGP is allowed as underlay protocol, so strictly speaking, this question has 2 possible answers, DE and CD

upvoted 1 times

🗲️ 👤 **Smoothey** 3 years, 2 months ago

Selected Answer: DE

Uses dot1.q with Vlan4 and OSPF for routing:

```
interface Ethernet1/1.4
```

```
description 40G Link to POD1-SPINE-101(1/1)
```

```
mtu 9150
```

```
encapsulation dot1q 4
```

```
vrf member fabric-mpod-vrf
```

```
ip address 192.168.1.253/30
```

```
ip ospf network point-to-point
```

upvoted 2 times

🗲️ 👤 **thomyohan** 4 years ago

Answer : DE

The spine interfaces are connected to the IPN devices through point-to-point routed interfaces. However, traffic originated from the spine interfaces is always tagged with an 802.1q VLAN 4 value

PIM Bidir used in the IPN network -

Since BUM traffic can be originated by all the leaf nodes deployed across Pods, the use of a different PIM protocol (like PIM ASM, for example) would results in the creation of multiple individual (S, G) entries on the IPN devices that may exceed the specific platform capabilities. With PIM Bidir, a single (*, G) entry must be created for a given BD, independently from the overall number of leaf nodes.

upvoted 4 times

🗲️ 👤 **thiyagas** 4 years, 5 months ago

Correct Answer is BE

upvoted 4 times

🗲️ 👤 **[Removed]** 4 years, 5 months ago

Correct Answer: BE. Multi-pod needs PIM bidirectional and it works with OSPF routing protocol.

upvoted 3 times

🗲️ 👤 **hybersat** 4 years, 5 months ago

They ask if PIM ASM are required not PIM Bidir :-)

upvoted 7 times

A Solutions Architect is asked to design two data centers based on Cisco ACI technology that can extend L2/L3, VXLAN, and network policy across locations. ACI

Multi-Pod has been selected. Which two requirements must be considered in this design? (Choose two.)

- A. ACI underlay protocols, i.e. COOP, IS-IS and MP-BGP, spans across pods. Create QoS policies to make sure those protocols have higher priority.
- B. A single APIC Cluster is required in a Multi-Pod design. It is important to place the APIC Controllers in different locations in order to maximize redundancy and reliability.
- C. ACI Multi-Pod requires an IP Network supporting PIM-Bidir.
- D. ACI Multi-Pod does not support Firewall Clusters across Pods. Firewall Clusters should always be local.
- E. Multi-Pod requires multiple APIC Controller Clusters, one per pod. Make sure those clusters can communicate to each other through a highly available connection.

Correct Answer: BC

Community vote distribution

BC (100%)

🗳️ 👤 **[Removed]** Highly Voted 👍 2 years, 11 months ago

Correct Answer: BC .
upvoted 26 times

🗳️ 👤 **thiyagas** Highly Voted 👍 2 years, 11 months ago

Correct Answer is BC
upvoted 13 times

🗳️ 👤 **jecq** Most Recent 🕒 1 year ago

Selected Answer: BC
B and C.
upvoted 2 times

🗳️ 👤 **korthab** 1 year, 4 months ago

Selected Answer: BC
The correct answer is B, C
upvoted 2 times

🗳️ 👤 **MrSaint** 2 years ago

A is not correct, altho COOP, IS-IS and MP-BGP spans across pods. this is a standard behavior which you dont have to worry about , so its not a matter of consider this as a requirement, is how it works. Have said that, you dont have to create QoS policies to achieve the requirements in the question per se.

C - You must enable PIM-Bidir in the IPN in order to L2 BUM traffic pass between PODs, otherwise the L2 extension will face some issues when talk to each other POD.

Reference <https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2019/pdf/BRKACI-2003.pdf> slide 28

Answer is B and C
upvoted 2 times

🗳️ 👤 **Rododendron2** 7 months, 4 weeks ago

altho COOP, IS-IS and MP-BGP spans : No - it does not span pods. They run separate instances of that.
upvoted 1 times

🗳️ 👤 **mrpeet** 2 years, 6 months ago

Correct answers are BC
Multipod is with only a single APIC cluster - multisite requires multiple clusters.

"The entire network hence runs as a single large fabric from an operational perspective; however, ACI Multi-Pod introduces specific enhancements to

isolate as much as possible the failure domains between Pods, contributing to increase the overall design resiliency. This is achieved by running separate instances of fabric control planes (IS-IS, COOP, MP-BGP) across Pods."

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html>

upvoted 4 times

An engineer configures a Multi-Pod system with the default gateway residing outside of the ACI fabric for a bridge domain. Which setting should be configured to support this requirement?

- A. disable Limit IP Learning to Subnet
- B. disable IP Data-plane Learning
- C. disable Unicast Routing
- D. disable Advertise Host Routes

Correct Answer: A

Community vote distribution

C (100%)

🗳️ 👤 **[Removed]** Highly Voted 4 years, 5 months ago

Correct Answer :C , Unicast routing should be disable to configure the default gateway outside the ACI fabric .
upvoted 14 times

🗳️ 👤 **hybersat** Highly Voted 4 years, 5 months ago

Yes C is correct, see this:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/migration_guides/migrating_existing_networks_to_aci.html#_Toc27467766

upvoted 8 times

🗳️ 👤 **designated** Most Recent 7 months, 2 weeks ago

Selected Answer: C

C is correct.

Unicast routing is enabled by default, and is required when you configure a default gateway for a bridge domain inside Cisco ACI fabric. If you configure the default gateway outside the fabric (for example, on a firewall), you should disable unicast routing and enable ARP flooding.

upvoted 1 times

🗳️ 👤 **GeorgeFortiGate** 2 years, 3 months ago

Selected Answer: C

If the DG is outside the ACI Fabric , you should disable the Unicast Routing for sure . If not the you can have unexpected outages. C for sure 100%

upvoted 1 times

🗳️ 👤 **Gab99** 2 years, 8 months ago

Selected Answer: C

they picked C

upvoted 1 times

🗳️ 👤 **eddyedwards257** 2 years, 11 months ago

Should be C

upvoted 1 times

🗳️ 👤 **thiyagas** 4 years, 5 months ago

shouldn't this be C .?

upvoted 4 times

What do Pods use to allow Pod-to-Pod communication in a Cisco ACI Multi-Pod environment?

- A. over Layer 3 directly connected back-to-back spines
- B. over Layer 3 Out connectivity via border leafs
- C. over Layer 3 IPN connectivity via border leafs
- D. over Layer 3 IPN connectivity via spines

Correct Answer: D

Reference:

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html>

Community vote distribution

D (100%)

🗨️ 👤 **zelya19** 9 months, 4 weeks ago

Selected Answer: D

The answer is D

From a physical perspective, the different Pods are interconnected by leveraging an "Inter-Pod Network" (IPN). Each Pod connects to the IPN through the spine nodes; the IPN can be as simple as a single Layer 3 device, or can be built with a larger Layer 3 network infrastructure.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html#Introduction>

upvoted 1 times

🗨️ 👤 **freed** 1 year ago

I think this question is outdated:

Beginning with Cisco APIC Release 5.2(3), the ACI Multi-Pod architecture is enhanced to support connecting the spines of two Pods directly with back-to-back ("B2B") links. With this solution, called Multi-Pod Spines Back-to-Back, the IPN requirement can be removed for small ACI Multi-Pod deployments. Multi-Pod Spines Back-to-Back also brings operational simplification and end-to-end fabric visibility, as there are no external devices to configure. For more information, please refer to the paper at the link below: <https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/kb/cisco-multipod-b2b.html>

upvoted 4 times

🗨️ 👤 **slogoheinzy** 7 months, 1 week ago

i agree - back -to-back is only for 2 pods but the questions doesnt specify that anyway

upvoted 1 times

An engineer must advertise a selection of external networks learned from a BGP neighbor into the ACI fabric. Which L3Out subnet configuration option creates an inbound route map for route filtering?

- A. External Subnets for the External EPG
- B. Shared Route Control Subnet
- C. Import Route Control Subnet
- D. Shared Security Import Subnet

Correct Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **Marinheiro** 7 months, 3 weeks ago

Selected Answer: C

Route Control for Routing Protocol

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/guide-c07-743150.html>

upvoted 1 times

An engineer must set up a Cisco ACI fabric to send Syslog messages related to hardware events, such as chassis line card failures. The messages should be sent to a dedicated Syslog server. Where in the Cisco APIC should the policy be configured to meet this requirement?

- A. uni/tn-common/monepg-default
- B. uni/infra/monifra-default
- C. uni/fabric/monfab-default
- D. uni/fabric/moncommon

Correct Answer: C

Community vote distribution

C (71%)

D (29%)

🗳️ 👤 **zelya19** 9 months, 4 weeks ago

Selected Answer: C

The Cisco Application Policy Infrastructure Controller (APIC) includes the following four classes of default monitoring policies:

monCommonPol (uni/fabric/moncommon): applies to all fabric, access, and tenant hierarchies

monFabricPol (uni/fabric/monfab-default): applies to fabric hierarchies

monInfraPol (uni/infra/monifra-default): applies to the access infrastructure hierarchy

monEPGPol (uni/tn-common/monepg-default): applies to tenant hierarchies

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/aci-fundamentals/Cisco-ACI-Fundamentals-401/Cisco-ACI-Fundamentals-401_chapter_01100.html#concept_DA7F29B8549E40E4BDE4495565F1FDCA

upvoted 2 times

🗳️ 👤 **sailorsoul** 11 months, 1 week ago

Selected Answer: C

I would pick C.

Since the question only asks hardware events.

upvoted 1 times

🗳️ 👤 **Ron_Berserker** 1 year, 3 months ago

Selected Answer: C

uni/fabric/moncommon also would work because includes both fabric and access, but the question only requeries hardware events and that is only Fabric, which is monfab.

Fabric: fabric ports, cards, chassis, fans, and so on

monFabricPol (uni/fabric/monfab-default): applies to fabric hierarchies https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/aci-fundamentals/Cisco-ACI-Fundamentals-401/Cisco-ACI-Fundamentals-401_chapter_01100.html

upvoted 2 times

🗳️ 👤 **S_1292_A** 1 year, 5 months ago

Selected Answer: D

tested with Visore, path is

dn=uni/fabric/moncommon

upvoted 2 times

🗳️ 👤 **Rododendron2** 1 year, 7 months ago

Not clear to me

monCommonPol (uni/fabric/moncommon): applies to both fabric and access infrastructure hierarchies

- monFabricPol (uni/fabric/monfab-default): applies to fabric hierarchies
- monInfraPol (uni/infra/monifra-default): applies to the access infrastructure hierarchy
- monEPGPol (uni/tn-common/monepg-default): applies to tenant hierarchies

upvoted 1 times

The existing network and ACI fabric have been connected to support workload migration. Servers will physically terminate at the Cisco ACI, but their gateway must stay in the existing network. The solution needs to adhere to Cisco's best practices. The engineer started configuring the relevant Bridge Domain and needs to complete the configuration. Which group of settings are required to meet these requirements?

A. L2 Unknown Unicast: Hardware Proxy
L3 Unknown Multicast Flooding: Flood
Multi Destination Flooding: Flood in BD

ARP Flooding: Enable -

B. L2 Unknown Unicast: Flood -
L3 Unknown Multicast Flooding: Flood
Multi Destination Flooding: Flood in BD

ARP Flooding: Enable -

C. L2 Unknown Unicast: Flood -
L3 Unknown Multicast Flooding: Optimize Flood
Multi Destination Flooding: Flood in BD

ARP Flooding: Disable -

D. L2 Unknown Unicast: Hardware Proxy
L3 Unknown Multicast Flooding: Optimize Flood
Multi Destination Flooding: Flood in BD
ARP Flooding: Disable

Correct Answer: B

Community vote distribution

B (100%)

 **Marinheiro** 7 months, 3 weeks ago

Selected Answer: B

To enable this behavior, you must configure specific properties on the bridge domain defined in the Cisco ACI fabric and associated to the legacy VLAN 10 (shown in the following diagram). The default settings for a bridge domain have ARP Flooding disabled and Unicast Routing enabled. For Layer 2 communication to work, you must adjust these settings from the defaults.


https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/migration_guides/migrating_existing_networks_to_aci.html
upvoted 1 times

An engineer must implement management policy and data plane separation in the Cisco ACI fabric. Which ACI object must be created in Cisco APIC to accomplish this goal?

- A. Application profile
- B. Tenant
- C. Contract
- D. Bridge domain

Correct Answer: B

Community vote distribution



Aamir_Inamdar 6 months ago

Selected Answer: D

As tenant already created for mgmt and question is must be created to separate management and data plane so I think BD us correct
upvoted 1 times

udo2020 6 months, 3 weeks ago

Selected Answer: D

A Bridge Domain (BD) is a Layer 2 forwarding construct that provides separation between management, policy, and data planes. By configuring different bridge domains, you can segment traffic within the fabric, effectively achieving the goal of separating management policy and the data plane.

upvoted 1 times

designated 7 months, 2 weeks ago

Selected Answer: B

B is correct = management policy and data plane separation

upvoted 2 times

Ceara 12 months ago

Selected Answer: B

Tenant

upvoted 1 times

Ceara 12 months ago

Selected Answer: B

Tenant should be the answer

upvoted 1 times

Marinheiro 1 year, 1 month ago

Selected Answer: C

Contracts

In addition to EPGs, contracts (vzBrCP) are key objects in the policy model. EPGs can only communicate with other EPGs according to contract rules. The following figure shows the location of contracts in the management information tree and their relation to other objects in the tenant.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/policy-model-guide/b-Cisco-ACI-Policy-Model-Guide.html#id_107453

upvoted 1 times

An engineer is implementing a Cisco ACI environment that consists of more than 20 servers. Two of the servers support only Cisco Discovery Protocol with no other link discovery protocol. The engineer wants the servers to be discovered automatically by the Cisco ACI fabric when connected. Which action must be taken to meet this requirement?

- A. Create an override policy that enables Cisco Discovery Protocol after LLDP is enabled in the default policy group.
- B. Configure a higher order interface policy that enables Cisco Discovery Protocol for the interface on the desired leaf switch.
- C. Configure a lower order policy group that enables Cisco Discovery Protocol for the interface on the desired leaf switch.
- D. Create an interface profile for the interface that disables LLDP on the desired switch that is referenced by the interface policy group.

Correct Answer: A

🗲️ 👤 **prospio971** 5 months ago

Selected Answer: B

To enable automatic discovery of the two servers that support only Cisco Discovery Protocol (CDP) by the Cisco ACI fabric, the engineer must: Enable CDP on the ACI leaf switch interfaces where these servers will be connected.

By enabling CDP on the relevant leaf switch interfaces, the ACI fabric will be able to discover and gather information about the connected servers using CDP. This action ensures that even though these servers do not support LLDP (Link Layer Discovery Protocol), which is the default discovery protocol used in ACI, they can still be automatically discovered and integrated into the fabric.

upvoted 1 times

🗲️ 👤 **2eb1ea8** 5 months, 1 week ago

Selected Answer: A

A leaf interface override policy allows interfaces that have interface policy group assignments to apply an alternate interface policy group

upvoted 1 times

🗲️ 👤 **udo2020** 6 months, 3 weeks ago

Selected Answer: B

In Cisco ACI, interface policies determine the behavior of physical interfaces, including protocols like Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP). Since LLDP is enabled by default in ACI, configuring a higher-order interface policy that explicitly enables CDP will override the default behavior, ensuring that CDP is used on the interfaces connected to the servers that only support CDP.

upvoted 2 times

🗲️ 👤 **Narbleddeath** 1 year, 2 months ago

A is correct.

<https://www.tejasjain.com/2019/10/cisco-aci-cdp-configuration.html>


upvoted 3 times

An engineer wants to monitor all configuration changes, threshold crossing, and link-state transitions in a Cisco ACI fabric. Which action must be taken to receive the required messages?

- A. Add Faults and Events to the monitor policy.
- B. Add Session Logs and Audit Logs to the monitor policy.
- C. Include Audit Logs and Events in the Syslog source policy.
- D. Include Events and Session Logs in the Syslog source policy.

Correct Answer: C

Community vote distribution



🗳️ 👤 **zelya19** 9 months, 4 weeks ago

Selected Answer: C

Events – Holds records of system related events (i.e., link state transitions, Logged Contract hits)

Audit Logs – Records user-initiated events (i.e., logins, configuration changes)

<https://unofficialaciguide.com/2018/08/11/configuring-syslog-for-aci/>

upvoted 2 times

🗳️ 👤 **Marinheiro** 7 months, 3 weeks ago

That's it, thanks!

upvoted 1 times

🗳️ 👤 **Huberttheman** 11 months ago

Selected Answer: C

Faults are default on the syslog policy, so you would add Audit and Events

upvoted 1 times

🗳️ 👤 **sailorsoul** 11 months, 1 week ago

Selected Answer: A

I would pick A. Audit logs are only user events. While the question states it should also include link state changes and etc.

upvoted 1 times

🗳️ 👤 **Macc10** 1 year, 1 month ago

Unsure on this

"Monitoring configuration changes" - you can see within the events logs which suggest they have occurred due to configuration change, the audit logs of course give more detail which makes this question confusing as you would think Audit logs are required to view all the configuration changes.

When reading Cisco whitepaper there is a statement which says "Any configuration or state change in any MO is considered an event".

a threshold crossing. This is something which I think is specific to the faults which made me think it was A.

However these selections I believe are made within the Syslog Source Policy. You tick each one.

upvoted 1 times

🗳️ 👤 **hebdeb** 1 year, 5 months ago

Selected Answer: A

It is A

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/faults/guide/b_APIC_Faults_Errors/b_IFC_Faults_Errors_chapter_011.html

upvoted 1 times

🗳️ 👤 **S_1292_A** 1 year, 5 months ago

Selected Answer: C

C is correct

upvoted 1 times

🗨️ 👤 **John771** 1 year, 8 months ago

Selected Answer: C

Correct answer: C

<https://unofficialaciguide.com/2018/08/11/configuring-syslog-for-aci/>
upvoted 2 times

🗨️ 👤 **John771** 1 year, 8 months ago

Correct answer: C

<https://unofficialaciguide.com/2018/08/11/configuring-syslog-for-aci/>
upvoted 2 times

🗨️ 👤 **cooljit** 1 year, 8 months ago

Correct Answer : A .

<https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2020/pdf/BRKACI-2271.pdf>

Syslog Overview

35BRKACI-2271

- Syslog is a method to collect messages (Faults, Events, Audit, Session) from devices to a server running a syslog daemon.

- During operation, a Fault or Event in ACI system can trigger the sending of a system log (syslog) message to the console and/or to an external logging server.

- Fault-generated system messages are triggered by:

- A fault rule.
- A threshold crossing
- A failure of a task or Finite State Machine (FSM) sequence.

- Event-generated system messages are triggered by:

- An event rule.
- An event in NX-OS of Leaf/Spine.

📌 Syslog like SNMP Traps are also enabled

upvoted 2 times

An organization has encountered many STP-related issues in the past due to failed hardware components. They are in the process of long-term migration to a newly deployed ACI fabric. Senior engineers are worried that spanning-tree loops in the existing network may be extended to the ACI fabric. Which feature must be enabled on the ACI leaf ports to protect the fabric from spanning-tree loops?

- A. BPDU Guard
- B. per-VLAN MCP
- C. Storm Control
- D. BPDU Filter

Correct Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **udo2020** 6 months, 3 weeks ago

Selected Answer: A

BPDU Guard is a feature that protects the ACI fabric from potential spanning-tree loops by disabling ports that receive BPDUs. When enabled on the ACI leaf ports, BPDU Guard shuts down the port immediately if it detects BPDUs, preventing spanning-tree issues from propagating into the ACI fabric.

This is especially critical during migration scenarios where legacy networks using Spanning Tree Protocol (STP) are connected to the ACI fabric, as it ensures the integrity of the ACI environment by isolating potential STP-related problems.

upvoted 1 times

🗨️ 👤 **Marinheiro** 1 year, 1 month ago

Selected Answer: B

The ACI Fabric does not participate in spanning tree protocol (STP) but instead acts as hub with respect to STP.

<https://unofficialaciguide.com/2018/03/27/using-mcp-miscabling-protocol-for-aci/#:~:text=MCP%20supports%2056%20VLANs%20per,from%20APIC%20version%203.2%20onward.>)

upvoted 1 times

🗨️ 👤 **imamus** 2 years, 1 month ago

Selected Answer: B

i think the key is failed hardware components - MCP controls that

upvoted 2 times

A network engineer must design a method to allow the Cisco ACI to redirect traffic to the firewalls. Only traffic that matches specific L4-L7 policy rules should be redirected. The load must be distributed across multiple firewalls to scale the performance horizontally. Which action must be taken to meet these requirements?

- A. Configure ACI Service Graph with Unidirectional PBR.
- B. Implement ACI Service Graph with GIPo.
- C. Implement ACI Service Graph Two Nodes with GIPo.
- D. Configure ACI Service Graph with Symmetric PBR.

Correct Answer: D

Community vote distribution

D (100%)

🗨️ 👤 **marceaubueno** 12 months ago

D. Because L4-L7 devices perform connection tracking (stateful firewall), they must see both directions of a flow. Therefore, you need to make sure that incoming and return traffic are redirected to the same PBR node.

upvoted 2 times

🗨️ 👤 **S_1292_A** 1 year, 5 months ago

Selected Answer: D

D is correct

upvoted 1 times

🗨️ 👤 **Mr_Certifiable** 1 year, 7 months ago

One of the main features of the service graph is Policy-Based Redirect (PBR).

With PBR, the Cisco ACI fabric can redirect traffic between security zones to L4-L7 devices, such as a firewall, Intrusion-Prevention System (IPS), or load balancer, without the need for the L4-L7 device to be the default gateway for the servers or the need to perform traditional networking configuration such as Virtual Routing and Forwarding (VRF) sandwiching or VLAN stitching

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739971.html>

upvoted 2 times

An engineer created two interface protocol policies called Pol_CDP40275332 and Pol_LLDP46783451. The policies must be used together in a single policy. Which ACI object must be used?

- A. interface policy group
- B. switch policy group
- C. switch profile
- D. interface profile

Correct Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Tamerelsayed** 8 months, 1 week ago

D

interface Profile

upvoted 1 times

🗳️ 👤 **zelya19** 9 months, 4 weeks ago

Selected Answer: A

Fabric > Access Policies > Interfaces > Leaf Interfaces > Policy Groups >...

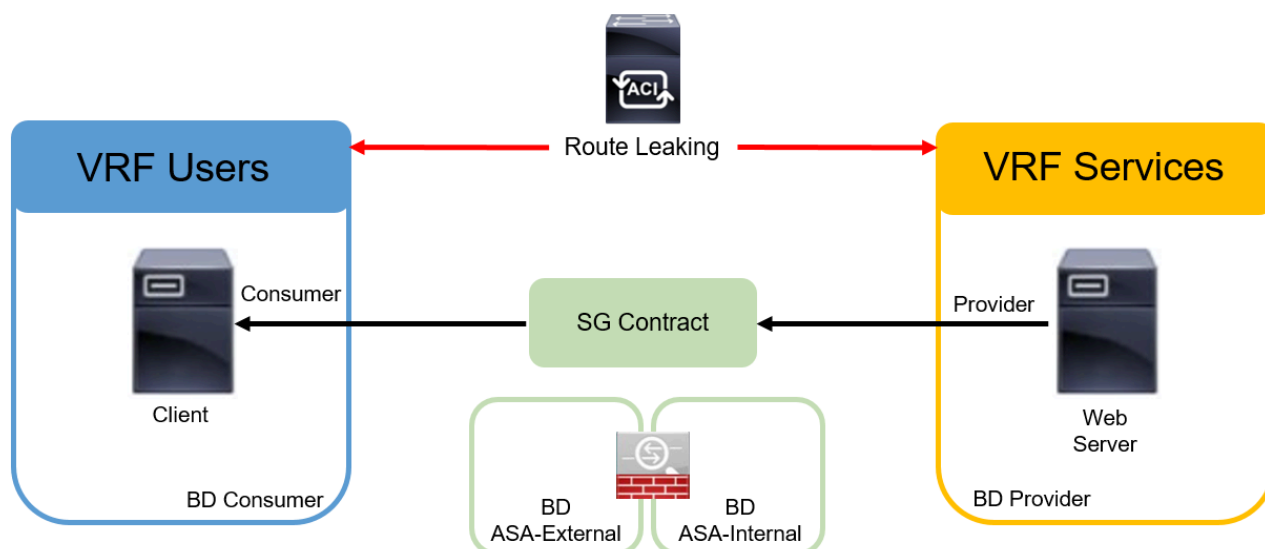
upvoted 3 times

What is the minimum number of APICs does Cisco recommend to deploy in a production cluster?

- A. 1
- B. 3
- C. 4
- D. 5

Correct Answer: *B*

Currently there are no comments in this discussion, be the first to comment!



Refer to the exhibit. An engineer must implement the inter-tenant service graph. Which set of actions must be taken to accomplish this goal?

- A. • Define the contract in the provider tenant and export it to the consumer tenant.
• Define the Layer 4 to Layer 7 device, service graph template, and ASA bridge domains in the provider tenant.
- B. • Define the contract in the provider tenant and export it to the consumer tenant.
• Define the Layer 4 to Layer 7 device and service graph template in the provider tenant and the ASA bridge domains in the consumer tenant.
- C. • Define the contract in the consumer tenant and export it to the provider tenant.
• Define the Layer 4 to Layer 7 device and service graph template in the provider tenant and the ASA bridge domains in the consumer tenant.
- D. • Define the contract in the consumer tenant and export it to the provider tenant.
• Define the Layer 4 to Layer 7 device, service graph template, and ASA bridge domains in the consumer tenant.

Correct Answer: A

Community vote distribution

A (100%)

sailorsoul 11 months, 1 week ago

Selected Answer: A

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739971.html>

Figure 120 shows a configuration example in which the provider EPG is in VRF2 in a user tenant and the consumer EPG is in VRF1 in another user tenant:

- The contract is defined in the provider tenant and is exported to the consumer tenant so that it is visible from both the consumer and provider EPGs.
- The device selection policy is defined in the provider tenant because the provider EPG is in the provider tenant.
- The L4-L7 device and service graph template are defined in the provider tenant so that the contract can refer to the service graph template.
- PBR bridge domains in VRF2 are defined in the provider tenant so that the device selection policy can refer to the cluster interfaces in the L4-L7 device and the PBR bridge domains.

upvoted 3 times

RUT 1 year ago

Selected Answer: A

The answer should be A

upvoted 1 times

  **examz_collector** 1 year, 7 months ago

Selected Answer: A

If the provider and the consumer are in different user tenants, BDs for the service device must be in the provider user tenant.

upvoted 2 times

All workloads in VLAN 1001 have been migrated into EPG-1001. The requirement is to move the gateway address for VLAN 1001 from the core outside the Cisco ACI fabric into the Cisco ACI fabric. The endpoints in EPG-1001 must route traffic to endpoints in other EPGs and minimize flooded traffic in the fabric. Which configuration set is needed on the bridge domain to meet these requirements?

- A. Enable Flood -
Enable Unicast Routing
- B. Disable Local IP Learning Limit -
Disable Unicast Routing
- C. Disable ARP Flood -
Disable Limit Endpoint Learning
- D. Enable Hardware Proxy -
Enable Unicast Routing

Correct Answer: D

Community vote distribution

D (100%)

 **Dash** Highly Voted 2 years, 2 months ago

Selected Answer: D

D is correct for me!
upvoted 5 times

 **designated** Most Recent 7 months, 2 weeks ago

Selected Answer: D

D is correct
upvoted 1 times

 **Macc10** 1 year, 8 months ago

D is correct. Spine proxy utilized by enabling hardware proxy and enabling Unicast Routing turns on L3 capability on the BD.
upvoted 3 times

An engineer must advertise a bridge domain subnet out of the ACI fabric to an OSPF neighbor. Which two configuration steps are required? (Choose two.)

- A. Add External Subnet for External EPG flag under External EPG
- B. Configure Subnet scope to Advertised Externally
- C. Configure the Subnet under the EPG level
- D. Create Route Control Profile with the export direction under External EPG
- E. Add L3Out profile to the bridge domain using Associated L3Outs section

Correct Answer: BE

Community vote distribution

BE (100%)

🗳️ 👤 **zelya19** 9 months, 4 weeks ago

Selected Answer: BE

Tenant > Networking > Bridge Domains > BD > L3 Configurations tab:

- Mark a BD subnet with an "Advertised Externally" scope.
 - Associate the BD with the L3Out(s) that need(s) to advertise the BD subnet to the outside.
- upvoted 2 times

🗳️ 👤 **thingtanklearningDOTcom** 1 year, 7 months ago

Selected Answer: BE

Documented clearly here: <https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/guide-c07-743150.html#3AdvertiseinternalroutesBDsubnetsstoexternaldevices>

The key points here are as follows:

- Mark a BD subnet with an "Advertised Externally" scope.
 - Associate the BD with the L3Out(s) that need(s) to advertise the BD subnet to the outside.
- upvoted 3 times

🗳️ 👤 **Dash** 1 year, 8 months ago

Selected Answer: BE

Correct

upvoted 4 times

🗳️ 👤 **GuKe** 1 year, 9 months ago

Selected Answer: BE

It's BE

upvoted 2 times

🗳️ 👤 **frztt** 1 year, 9 months ago

Selected Answer: BE

Correct, BE

upvoted 2 times

🗳️ 👤 **VY01** 1 year, 9 months ago

Selected Answer: BE

I think should be B & E the same

upvoted 2 times

🗳️ 👤 **KrishEngineer** 1 year, 9 months ago

Selected Answer: BE

It should be B & E

upvoted 2 times

An engineer must connect a new host to port 1/1 on Leaf 101. A Cisco ACI fabric has an MCP policy configured but experiences excessive Layer 2 loops. The engineer wants the Cisco ACI fabric to detect and prevent Layer 2 loops in the fabric. Which set of actions accomplishes these goals?

- A. Enable MCP locally.
Associate the MCP policy with an interface policy group.
- B. Enable MCP locally.
Associate the MCP policy with an interface profile.
- C. Enable MCP globally.
Associate the MCP policy with an interface selector.
- D. Enable MCP globally.
Associate the MCP policy with an interface policy group.

Correct Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **designated** 7 months, 2 weeks ago

Selected Answer: D

D is correct

upvoted 1 times

🗳️ 👤 **zelya19** 1 year, 3 months ago

Selected Answer: D

Fabric > Access Policies > Policies > Global > MCP Instance Policy default

Fabric > Access Policies > Policies > Interface > MCP Interface

upvoted 2 times

🗳️ 👤 **marceaubueno** 1 year, 6 months ago

D is correct. For MCP to work, it must be enabled globally and per-interface (attach the MCP policy under the interface policy group you have configured). If the ACI fabric receives an MCP packet on any interface, it will take action (could be just alerting with a fault, or err-disabling the interface on which the MCP packet was received).

upvoted 2 times



An ACI engineer is implementing a Layer 3 Out inside the Cisco ACI fabric that must meet these requirements:

- The data center core switch must be connected to one of the leaf switches with a single 1G link.
- The routes must be exchanged using a link-state routing protocol that supports hierarchical network design.
- The data center core switch interface must be using 802.1Q tagging, and each VLAN will be configured with a dedicated IP address.

Which set of steps accomplishes this goals?

- A. Set up the EIGRP Protocol policy with the selected Autonomous System number.
Create the Routed Outside object and Node Profile, selecting EIGRP.
Configure the Interface profile, selecting Routed Interface and the appropriate interfaces.
Create the External Network object with a network of 0.0.0 0/0.
- B. Set up the EIGRP Protocol policy with the selected Autonomous System number.
Set up the Routed External Network object and Node Profile, selecting EIGRP.
Create the Switch profile, selecting Port-channel and the appropriate interfaces.
Create the default network and associate it with the Routed Outside object.
- C. Set up the BGP Protocol policy with the Autonomous System number of 0.
Configure an Interface policy and an External Bridged Domain.
Create an External Bridged Network using the configured VLAN pool.
Build the Leaf profile, selecting Routed sub-interface and the appropriate VLAN.
- D. Configure the OSPF Protocol policy with an area of 0.
Create Routed Outside object and Node Profile, selecting OSPF as the routing protocol.
Build the Interface profile, selecting Routed Sub-interface and the appropriate VLAN.
Configure the External Network object with a network of 0 0.0 0/0.

Correct Answer: D

  **udo2020** 6 months, 3 weeks ago

Selected Answer: D

Ony OSPF is a link state routing protocoll


upvoted 1 times

An engineer is in the process of discovering a new Cisco ACI fabric consisting of two spines and four leaf switches. The discovery of leaf 1 has just been completed. Which two nodes are expected to be discovered next? (Choose two.)

- A. leaf 2
- B. leaf 4
- C. spine 1
- D. leaf 3
- E. spine 2

Correct Answer: CE

Community vote distribution



gentlesol 3 months ago

Selected Answer: CE

After the APIC discovers leaf1, next it discovers both spines and then the remaining leafs
upvoted 1 times

Macc10 7 months, 2 weeks ago

Selected Answer: CE

CE is the right answer - as per the Cisco Links below.

"Each APIC instance in the cluster first discovers only the leaf switch to which it is directly connected. After the leaf switch is registered with the APIC, the APIC discovers all spine switches that are directly connected to the leaf switch."

Leaf 1 will have a connection into both spines and as per the Cisco white paper text above I am sticking with CE.

upvoted 3 times

imamus 10 months, 3 weeks ago

Selected Answer: AC

Correct answers are A and C.

In a scenario where you have two spines and four leafs, logic assumes that leaf1 is connected to spine1, where discovery process starts. Once spine1 is discovered next visible devices will be closest connected downlinks (e.g. leaf2), and only then spine2 and so on.

upvoted 2 times

Mr_Certifiable 1 year, 1 month ago

Follow the bread crumbs in the Apic discovery process - what is the next hop after spine 1 -> likely leaf2 [no direct links between spines] assuming the seed leaf is leaf1 where the Apic would most likely be situated when starting provisioning.

[https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/3-](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/3-x/getting_started/b_APIC_Getting_Started_Guide_Rel_3_x/b_APIC_Getting_Started_Guide_Rel_3_x_chapter_0100.html)

[x/getting_started/b_APIC_Getting_Started_Guide_Rel_3_x/b_APIC_Getting_Started_Guide_Rel_3_x_chapter_0100.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/3-x/getting_started/b_APIC_Getting_Started_Guide_Rel_3_x/b_APIC_Getting_Started_Guide_Rel_3_x_chapter_0100.html)

therefore AC

upvoted 3 times

thinqtanklearningDOTcom 1 year, 1 month ago

Selected Answer: CE

Switch Discovery

About Switch Discovery with the APIC

The APIC is a central point of automated provisioning and management for all the switches that are part of the ACI fabric. A single data center might include multiple ACI fabrics; each data center might have its own APIC cluster and Cisco Nexus 9000 Series switches that are part of the fabric. To ensure that a switch is managed only by a single APIC cluster, each switch must be registered with that specific APIC cluster that manages the fabric.

The APIC discovers new switches that are directly connected to any switch it currently manages. Each APIC instance in the cluster first discovers

only the leaf switch to which it is directly connected. After the leaf switch is registered with the APIC, the APIC discovers all spine switches that are directly connected to the leaf switch. As each spine switch is registered, that APIC discovers all the leaf switches that are connected to that spine switch. This cascaded discovery allows the APIC to discover the entire fabric topology in a few simple steps.

upvoted 1 times

🗲️ 👤 **Rododendron2** 1 year, 1 month ago

Selected Answer: CE

C & E is correct ... leaf connected to spines only...

upvoted 1 times

🗲️ 👤 **7korn7** 1 year, 2 months ago

Selected Answer: CE

[https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/3-](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/3-x/getting_started/b_APIC_Getting_Started_Guide_Rel_3_x/b_APIC_Getting_Started_Guide_Rel_3_x_chapter_0100.html)

[x/getting_started/b_APIC_Getting_Started_Guide_Rel_3_x/b_APIC_Getting_Started_Guide_Rel_3_x_chapter_0100.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/3-x/getting_started/b_APIC_Getting_Started_Guide_Rel_3_x/b_APIC_Getting_Started_Guide_Rel_3_x_chapter_0100.html)

The APIC discovers new switches that are directly connected to any switch it currently manages. Each APIC instance in the cluster first discovers only the leaf switch to which it is directly connected. After the leaf switch is registered with the APIC, the APIC discovers all spine switches that are directly connected to the leaf switch.

upvoted 3 times

🗲️ 👤 **VY01** 1 year, 3 months ago

Selected Answer: CE

C & E is correct

upvoted 3 times

🗲️ 👤 **Maya25** 1 year, 3 months ago

Selected Answer: CE

CE is correct answer.

upvoted 2 times

A bridge domain for an EPG called "Web Servers" must be created in the Cisco APIC. The configuration must meet these requirements:

- Only traffic to known MAC addresses must be allowed to reduce noise.
- The multicast traffic must be limited to the ports that are participating in multicast routing.
- The endpoints within the bridge domain must be kept in the endpoint table for 20 minutes without any updates.

Which set of steps configures the bridge domain that satisfies the requirements?

A. Switch L2 Unknown Unicast to Flood.

Select the default Endpoint Retention Policy and set the Local Endpoint Aging to 20 minutes.

Set Multicast Destination Flooding to Flood in Encapsulation.

B. Set L2 Unknown Unicast to Hardware Proxy.

Configure L3 Unknown Multicast Flooding to Optimized Flood.

Create an Endpoint Retention Policy with a Local Endpoint Aging Interval of 1200 seconds.

C. Multicast Destination Flooding should be set to Flood in BD.

Set L3 Unknown Multicast Flooding to Flood.

Select the default Endpoint Retention Policy with a Local Endpoint Aging Interval of 1200 seconds.

D. Select the ARP Flooding checkbox.

Create an Endpoint Retention Policy with a Remote Endpoint Aging Interval of 20 minutes.

Set L3 Unknown Multicast Flooding to Optimized Flooding.

Correct Answer: B

  **Mr_Certifiable** 11 months, 1 week ago

B

endpoint retention policy defined in seconds

Endpoint Retention Policy

The Endpoint Retention Policy configuration is located at Tenant > Policies > Protocol > End Point Retention (Figure 35) and is referred from a Bridge Domain (BD) or a VRF (Figure 36). By default, a BD or a VRF refers to the default policy defined in the common tenant is used

Local End Point Aging Interval: The amount of time in seconds that a leaf node can keep each local endpoint in its endpoint table without further updates. The default interval is 900 seconds. If 75 percent of the interval is reached, the leaf node sends three ARP requests to verify the presence of the endpoint. If no response is received, the endpoint is deleted.

- Remote End Point Aging Interval: The amount of time in seconds that a leaf node can keep each remote endpoint in its endpoint table without further updates. The default interval is 300 seconds.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

upvoted 1 times

An engineer is troubleshooting fabric discovery in a newly deployed Cisco ACI fabric and analyzes this output:

LEAF101# show ip int brief vrf overlay-1

(...output truncated for brevity...)

IP Interface Status for VRF "overlay-1"(4)

Interface	Address	Interface Status
lo1023	10.233.44.32/32	protocol-up/link-up/admin-up

LEAF101# show vlan extended

VLAN	Name	Encap	Ports
8	infra:default	vxlan-41174821,	Eth1/1, Eth1/2, Eth1/47 vlan-3600

Which ACI fabric address is assigned to interface lo1023?

- A. VXLAN tunnel endpoint
- B. Physical tunnel endpoint
- C. Fabric tunnel endpoint
- D. Dynamic tunnel endpoint

Correct Answer: C

Community vote distribution

C (100%)

 **sailorsoul** 11 months, 1 week ago

Selected Answer: C

Lo1023 is FTEP address exists on every leaf.

upvoted 1 times

 **examz_collector** 1 year, 7 months ago

Selected Answer: C

Fabric TEP

upvoted 1 times

 **NoUserName1234** 1 year, 8 months ago

Selected Answer: C

<https://community.cisco.com/t5/application-centric-infrastructure/default-ip-interfaces-on-fabric-nodes/td-p/2651678>

upvoted 1 times

 **7korn7** 1 year, 8 months ago

Selected Answer: C

https://www.mvankleij.nl/post/aci_topology_hardware/

FTEP answer is C

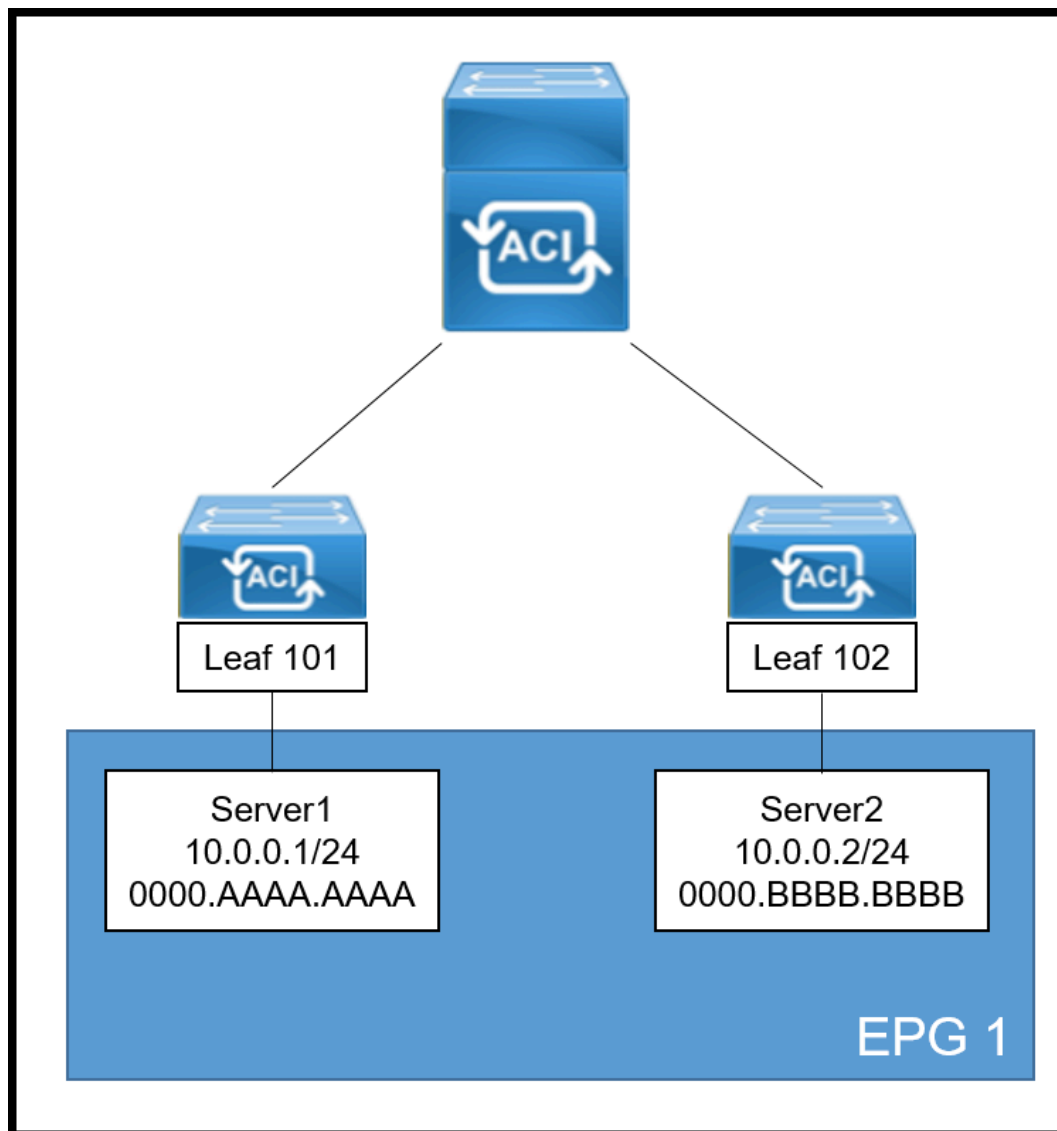
upvoted 3 times

The company's Cisco ACI fabric hosts multiple customer tenants. To meet a service level agreement, the company is constantly monitoring the Cisco ACI environment. Syslog is one of the methods used for monitoring. Only events related to leaf and spine environmental information without specific customer data should be logged. To which ACI object must the configuration be applied to meet these requirements?

- A. infra tenant
- B. access policy
- C. switch profile
- D. fabric policy

Correct Answer: *D*

Currently there are no comments in this discussion, be the first to comment!



Refer to the exhibit. A systems engineer is implementing the Cisco ACI fabric. However, the Server2 information is missing from the Leaf 101 endpoint table and the COOP database of the spine. The requirement is for the bridge domain configuration to enforce the ACI fabric to forward the unicast packets generated by Server1 destined to Server2. Which action must be taken to meet these requirements?

- A. Enable ARP Flooding
- B. Set L2 Unknown Unicast to Flood
- C. Set IP Data-Plane Learning to No
- D. Enable Unicast Routing

Correct Answer: B

prospio971 5 months ago

Selected Answer: B

By setting L2 Unknown Unicast to "Flood", the ACI fabric will flood unknown unicast traffic within the bridge domain. This ensures that packets from Server1 destined for Server2 will be forwarded to all ports in the bridge domain, allowing Server2 to receive the traffic even if its location is not known to the fabric

upvoted 1 times

An engineer must allow multiple external networks to communicate with internal ACI subnets. Which action should the engineer take to assign the prefix to the class ID of the external Endpoint Group?

- A. Enable the Export Route Control Subnet for the External Endpoint Group flag.
- B. Enable an L3Out with Shared Route Control Subnet.
- C. Configure subnets with the External Subnets for External EPG flag enabled.
- D. Configure subnets with the Import Route Control Subnet flag enabled.

Correct Answer: C

Community vote distribution

C (83%)

B (17%)

🗳️ 👤 **zelya19** 9 months, 4 weeks ago

Selected Answer: C

External Subnets for the External EPG: This defines which subnets belong to this external EPG for the purpose of defining a contract between EPGs. This is the same semantics as for an ACL in terms of prefix and mask.

<https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-application-centric-infrastructure-design-guide.html#ExternalNetworkExternalEPGConfigurationOptions>

upvoted 1 times

🗳️ 👤 **Macc10** 1 year, 1 month ago

Selected Answer: C

Answer is C, To allow communication between internal and external EPG's the External EPG Flag must be enabled.

upvoted 2 times

🗳️ 👤 **Rododendron2** 1 year, 6 months ago

assign the prefix to the class ID ???

upvoted 1 times

🗳️ 👤 **7korn7** 1 year, 8 months ago

Selected Answer: C

Sorry, my fault. answer is C

upvoted 2 times

🗳️ 👤 **7korn7** 1 year, 8 months ago

Selected Answer: B

multiple external networks is the key point. answer is B

upvoted 1 times

🗳️ 👤 **Anantharajesh** 1 year, 4 months ago

C is correct : "shared" is always meant for intra tenant routes... not for external

upvoted 1 times

An engineer must ensure that Cisco ACI flushes the appropriate endpoints when a topology change notification message is received in an MST domain. Which three steps are required to accomplish this goal? (Choose three.)

- A. Enable the BPDU interface controls under the spanning tree interface policy.
- B. Configure a new STP interface policy.
- C. Bind the spanning tree policy to the switch policy group.
- D. Associate the STP interface policy to the appropriate interface policy group.
- E. Create a new region policy under the spanning tree policy.
- F. Map VLAN range to MST instance number.

Correct Answer: CEF

Community vote distribution

CEF (100%)

🗳️ 👤 **Rododendron2** 6 months, 3 weeks ago

Selected Answer: CEF

<https://learnduty.com/cisco-aci/cisco-aci-and-mst-interaction/>
upvoted 3 times

🗳️ 👤 **frzzt** 9 months, 2 weeks ago

Selected Answer: CEF

I think it is CEF.
upvoted 3 times

🗳️ 👤 **VY01** 9 months, 1 week ago

Agree with Ref. <https://unofficialaciguide.com/2019/03/28/spanning-tree-stp-and-aci/>
upvoted 1 times

🗳️ 👤 **imamus** 8 months ago

spanning tree policy is part of the interface policy group, not the switch policy group, thus C is for sure not correct
upvoted 1 times

🗳️ 👤 **imamus** 7 months, 2 weeks ago

sorry, C is right, you need a switch policy group
upvoted 1 times

A Cisco ACI bridge domain and VRF are configured with a default data-plane learning configuration. Which two endpoint attributes are programmed in the leaf switch when receiving traffic? (Choose two.)

- A. Remote MAC, IP
- B. Remote Subnet
- C. Local IP, not MAC
- D. Local MAC, IP
- E. Local Subnet
- F. Remote IP



Correct Answer: DF

Community vote distribution

DF (100%)

 **frzzt**  2 years, 3 months ago

It can also be Remote MAC,IP because it depends if it is Routerd or switched traffic.. The question sucks
upvoted 7 times

 **designated**  7 months, 1 week ago

Selected Answer: AD

A and D are correct:

When dataplane IP address learning is enabled (which is the default setting) for a VRF instance, local and remote MAC addresses are learned using an endpoint-to-endpoint ARP request.

upvoted 1 times

 **designated** 7 months, 1 week ago

Notes: When dataplane IP address learning is disabled, all the remote IP address entries in the tenant VRF instance are removed. The local IP entries are aged out and, subsequently, will not be re learned through the dataplane, but can still be learned from the control plane.

• When dataplane IP address learning is disabled, already learned local IP endpoints are retained and require control plane refreshes (ARP) to be kept alive (assuming IP aging is also enabled). Dataplane Layer 3 traffic will not keep IP endpoints alive.

<https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/6x/l3-configuration/cisco-apic-layer-3-networking-configuration-guide-60x/dataplane-ip-learning-per-vrf-layer3-config-60x.pdf>


upvoted 1 times

 **sailorsoul** 1 year, 5 months ago

Selected Answer: DF

remote point can only be either ip address or mac address.

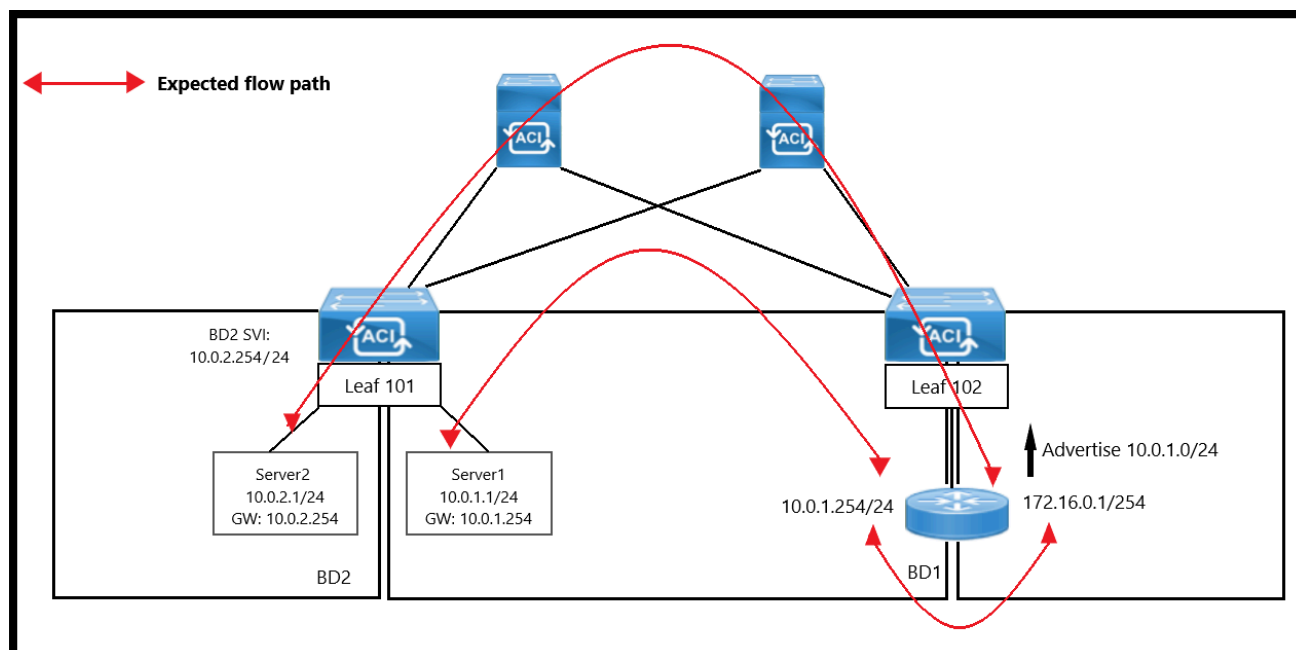
upvoted 2 times

 **7korn7** 2 years, 2 months ago

Selected Answer: DF

The key point is default data-plane learning configuration

upvoted 3 times



Refer to the exhibit. An engineer wants to initiate an ICMP ping from Server1 to Server2. The requirement is for the BD1 to enforce ICMP replies that follow the expected path. The packets must be prevented from taking the direct path from Leaf1 to Server1. Which action must be taken on BD1 to meet these requirements?

- A. Set L2 Unknown Unicast to Flood.
- B. Set L2 Unknown Unicast to Hardware Proxy.
- C. Disable Unicast Routing.
- D. Enable ARP Flooding.

Correct Answer: C

crooks_1988 Highly Voted 1 year, 9 months ago

B - when using BD as L3 anycast GW, like in this case, L2 unknown should be set to Hardware Proxy
upvoted 5 times

prospio971 Most Recent 6 months, 2 weeks ago

Selected Answer: C

To meet the requirements of enforcing ICMP replies to follow the expected path and prevent packets from taking the direct path from Leaf1 to Server1, the engineer must take the following action on BD1:

Disable unicast routing on the bridge domain (BD1).

upvoted 1 times

Rollizo 8 months, 1 week ago

Selected Answer: C

C. Disable Unicast Routing.

Explanation:

When you disable unicast routing on a bridge domain in Cisco ACI, it prevents the ACI fabric from routing unicast traffic directly. Instead, all traffic must go through the appropriate path as defined by the fabric's policies. In this scenario, by disabling unicast routing:

Traffic Redirection: ICMP replies from Server2 to Server1 will not take the direct path between Leaf1 and Server1. Instead, they will be redirected to follow the intended path as defined by the ACI policies and the configurations of the Bridge Domain.

Path Enforcement: This ensures that the packets must go through the designated path (possibly going through a service graph, firewall, or specific switches) as required by the architecture, instead of taking a shortcut that bypasses the intended flow.

upvoted 2 times

  **marceaubueno** 1 year, 6 months ago

That is right. However in this topology the gateway for server1 is outside of the ACI fabric, so C is correct.

upvoted 2 times

  **Marinheiro** 1 year, 1 month ago

I'm going with C

upvoted 1 times

A network engineer must configure a Cisco ACI system to detect network loops for untagged and tagged traffic. The loop must be detected and stopped by disabling an interface within 4 seconds. Which configuration must be used?

A.

Admin State: ☐ Disabled ☒ Enabled

Controls: ☐ Enable MCP PDU per VLAN

Key:

Confirm Key:

Loop Detect Multiplication Factor:

Loop Protection Action: ☒ Port Disable

Initial Delay (sec):

Transmission Frequency (sec): (msec):

B.

Admin State: ☐ Disabled ☒ Enabled

Controls: ☒ Enable MCP PDU per VLAN

Key:

Confirm Key:

Loop Detect Multiplication Factor:

Loop Protection Action: ☒ Port Disable

Initial Delay (sec):

Transmission Frequency (sec): (msec):

C.

Admin State: ☐ Disabled ☒ Enabled

Controls: ☒ Enable MCP PDU per VLAN

Key:

Confirm Key:

Loop Detect Multiplication Factor:

Loop Protection Action: ☒ Port Disable

Initial Delay (sec):

Transmission Frequency (sec): (msec):

D.

Admin State: ☐ Disabled ☒ Enabled

Controls: ☒ Enable MCP PDU per VLAN

Key:

Confirm Key:

Loop Detect Multiplication Factor:

Loop Protection Action: ☐ Port Disable

Initial Delay (sec):

Transmission Frequency (sec): (msec):

Correct Answer: B

Community vote distribution

B (100%)

  **designated** 7 months, 2 weeks ago

Selected Answer: B

Port, MCP and loop protection must be enabled and be disabled when reach 4 seconds.

upvoted 1 times

  **mdriraa** 11 months ago

Selected Answer: B

4sec x 1 so B is correct

upvoted 1 times

An engineer must configure a group of servers with a contract that uses TCP port 80. The EPG that contains the web servers requires an external Layer 3 cloud to initiate communication. Which action must be taken to meet these requirements?

- A. Configure the EPG as a provider and L3 out as consumer of the contract.
- B. Configure OSPF to exchange routes between the L3 out and EPG.
- C. Configure a taboo contract and apply it to the EPG.
- D. Configure the EPG as a consumer and L3 out as a provider of the contract.

Correct Answer: A

  **gentlesol** 3 months ago

Selected Answer: D

The right answer is D
upvoted 1 times

  **designated** 7 months, 1 week ago

Selected Answer: D

I'm wrong.. the correct answer is D.

- Option A suggests that the EPG is the provider, meaning it would assume the web servers are offering a service to the L3Out. However, that is not the case here—the web servers need access to the external network for their operations.

- Also, option A assumes the external cloud is consuming a service provided by the EPG, which contradicts the question's requirement.

upvoted 2 times

  **designated** 7 months, 2 weeks ago

Selected Answer: A

The EPG that contains the web servers (EPG provider) requires an external Layer 3 cloud (EPG consumer) to initiate communication.

upvoted 2 times

The unicast routing feature is enabled on the bridge domain. Which two conditions enable the Cisco ACI leaf to learn a source IP as a local endpoint? (Choose two.)

- A. Through Ethernet traffic received in a bridge domain.
- B. IP traffic routed through an SVI.
- C. Through VXLAN traffic received on the uplink.
- D. IP traffic routed through a Layer 3 Out.
- E. Through ARP received on an SVI.

Correct Answer: BE

Community vote distribution

BE (100%)

🗲️ 👤 **zelya19** 9 months, 4 weeks ago

Selected Answer: BE

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html#Localendpointlearning>

upvoted 2 times

🗲️ 👤 **sailorsoul** 11 months, 1 week ago

Selected Answer: BE

A Cisco ACI leaf switch follows these steps to learn a local endpoint MAC address and IP address:

1. The Cisco ACI leaf receives a packet with a source MAC Address (MAC A) and source IP Address (IP A).
2. The Cisco ACI leaf learns MAC A as a local endpoint.
- 3a. If the packet is an ARP request, the Cisco ACI leaf learns IP A tied to MAC A based on the ARP header.
- 3b. If the packet is an IP packet and routing is performed by the Cisco ACI leaf, the Cisco ACI leaf learns IP A tied to MAC A based on the IP header.

Thus, if the packet is switched and not an ARP packet, the Cisco ACI leaf never learns the IP address but only the MAC address. This behavior is the same as traditional MAC address learning behavior on a traditional switch.

upvoted 4 times

When does the Cisco ACI leaf learn a source IP or MAC as a remote endpoint?

- A. When VXLAN traffic arrives on a leaf fabric port from the spine and outer source IP is in the Layer 3 Out EPG subnet range.
- B. When VXLAN traffic arrives on a leaf fabric port from the spine and outer source IP is in the bridge domain subnets range.
- C. When VXLAN traffic arrives on a leaf fabric port from the spine and inner source IP is in the Layer 3 Out EPG subnet range.
- D. When VXLAN traffic arrives on a leaf fabric port from the spine and inner source IP is in the bridge domain subnets range.

Correct Answer: D

Community vote distribution

D (89%)

11%

 **designated** 7 months, 2 weeks ago

Selected Answer: B

Hardware-Proxy behavior


Leaf local > Spine (COOP) > Leaf Remote

upvoted 1 times

 **designated** 7 months, 2 weeks ago

D is correct!!!

upvoted 1 times


 **benguela** 1 year, 1 month ago

B

Remote endpoint learning

Cisco ACI learns a MAC or IP address as a remote endpoint when a packet comes into a Cisco ACI leaf switch from another leaf switch through a spine switch. When a packet is sent from one leaf to another leaf, Cisco ACI encapsulates the original packet with an outer header representing the source and destination leaf Tunnel Endpoint (TEP) and the Virtual Extensible LAN (VXLAN) header, which contains the bridge domain or VRF information of the original packet.

upvoted 2 times

 **zelya19** 1 year, 3 months ago

Selected Answer: D

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html#Remoteendpointlearning>

upvoted 1 times

 **sailorsoul** 1 year, 5 months ago

Selected Answer: B

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

When a packet is sent from one leaf to another leaf, Cisco ACI encapsulates the original packet with an outer header representing the source and destination leaf Tunnel Endpoint (TEP) and the Virtual Extensible LAN (VXLAN) header, which contains the bridge domain or VRF information of the original packet.

It is still B. The bridge domain information is contained on the outer IP header.

upvoted 1 times

 **sailorsoul** 1 year, 5 months ago

A Cisco ACI leaf switch follows these steps to learn a remote endpoint MAC or IP address:

1. The Cisco ACI leaf receives a packet with source MAC A and source IP A from a spine switch.
2. The Cisco ACI leaf learns MAC A as a remote endpoint if VXLAN contains bridge domain information.
3. The Cisco ACI leaf learns IP A as a remote endpoint if VXLAN contains VRF information.

upvoted 2 times

  **asd248402** 2 years ago

Selected Answer: D

Inner IP



upvoted 3 times

  **Rododendron2** 2 years, 1 month ago

Selected Answer: D

Need inner source IP - D

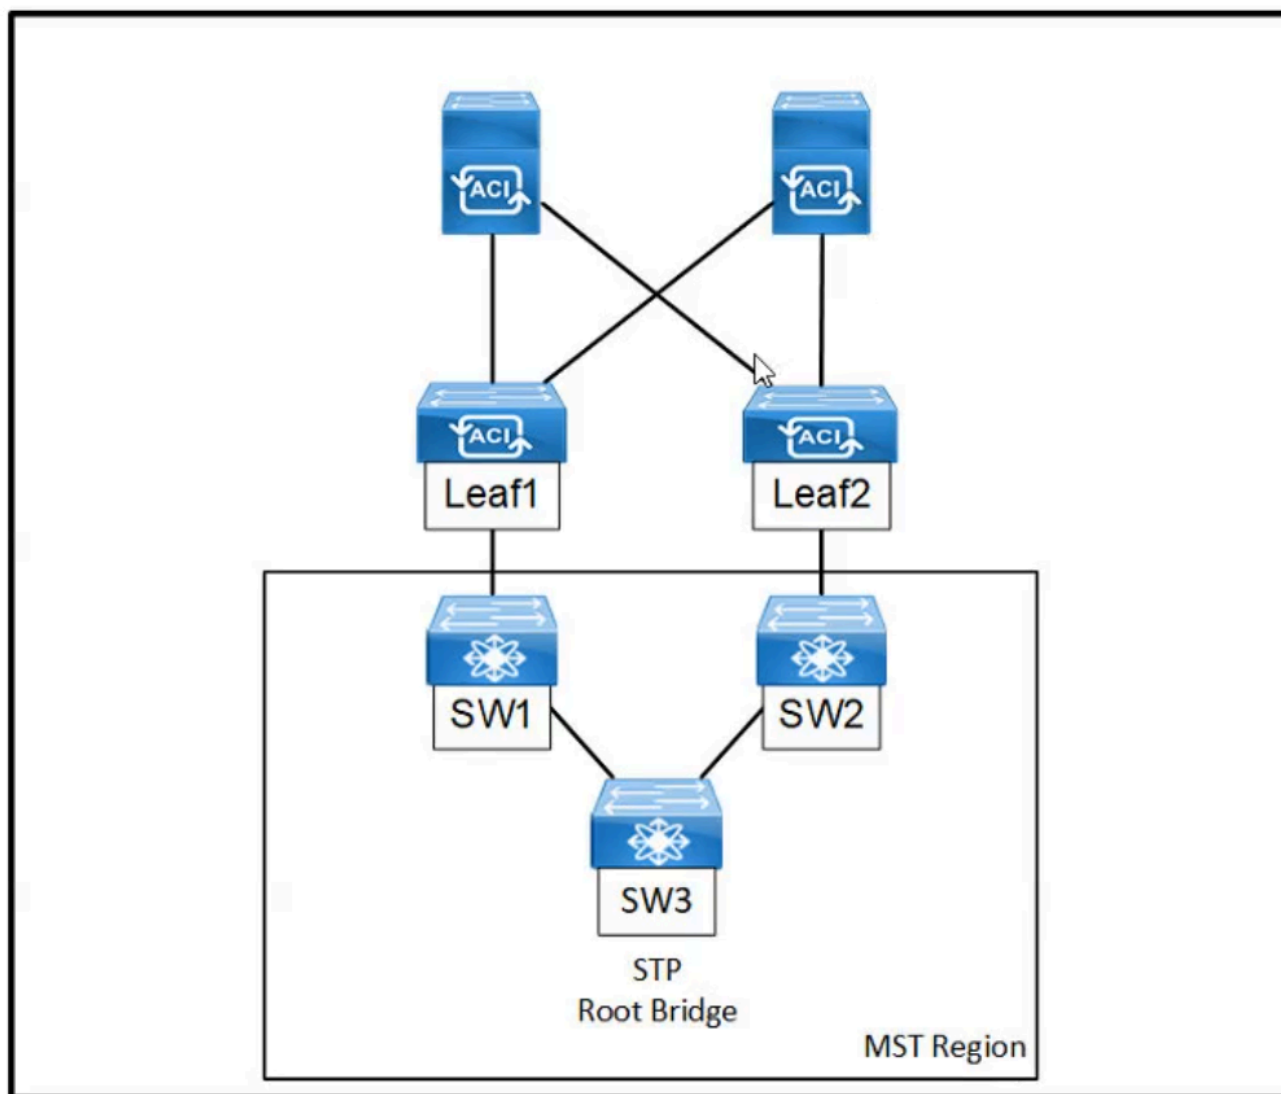
upvoted 3 times

  **frztt** 2 years, 3 months ago

Selected Answer: D

It is D. Inner source is BD subnet. Outer address is VXLAN TEP.

upvoted 4 times



Refer to the exhibit. An engineer is deploying a Cisco ACI environment but experiences a STP loop between SW1 and SW2. Which configuration step is needed to break the STP loop?

- A. Configure a Layer 2 external bridged network on the interfaces facing the MST switches.
- B. Enable the native VLAN on the interfaces facing the MST switches using static ports in a dedicated EPG.
- C. Enable BPDU filter under the STP interface policy on the interfaces facing the MST switches.
- D. Configure the STP instance to VLAN mapping under the switch STP policy.

Correct Answer: B

Community vote distribution

B (100%)

designated 7 months, 2 weeks ago

Selected Answer: B

MST does not use VLAN TAG and it should be native VLAN for this case.

upvoted 1 times

Ross99 1 year, 11 months ago

Selected Answer: B

Should be B

upvoted 1 times

thingtanklearningDOTcom 2 years, 1 month ago

Selected Answer: B

BPDUs in MST are sent on the Native Vlan (these are sent untagged), and not on a per-vlan basis as is the case with RPVST and PVST. This means we will have to make configuration specifically to accommodate for these differences in the ACI Fabric.

To deal with the untagged MST BPDUs, you'll need to configure a couple of things differently than what you do when dealing with RPVST or PVST.

Create a Switch Policy Group for your MST region – Under Fabric Access Policies, you'll need to create a Switch Policy Group (note – I'm not talking about Interface Policy Groups). For this Policy Group, you'll create a Spanning-tree policy and add in information about your MST region.



Native Vlan EPG – For ACI, in your Tenant where your Layer-2 connection resides, you will need to create a specific EPG to carry the MST BPDUs. The static path bindings for your Legacy switches will be configured as dot1p (native). Failure to do this could very likely result in a loop.

upvoted 4 times

  **Narbleddeath** 2 years, 2 months ago


Picking B.

upvoted 1 times

  **frzst** 2 years, 3 months ago

This should be B&D. BPDU Filter is not helping when it is MST. With MST BPDU are blocked when entering the fabric because they are sent on native vlan. You need to map EPG to native vlan on all ports facing MST switches. And also map vlans to MST instance..

upvoted 2 times

  **samfab** 4 months, 2 weeks ago

yes, Should be B&D

upvoted 1 times

What is MP-BGP used for in Cisco ACI fabric?

- A. MP-BGP VPNv4 AF is used as protocol on L3Out between a border leaf and an external router
- B. MP-BGP Layer 2 VPN EVPN AF is used to propagate L3Out routes that are received from a border leaf
- C. MP-BGP VPNv4 AF is used to propagate L3Out routes that are received from a border leaf to the fabric
- D. MP-BGP VPNv4 AF is used between spines in an ACI Multi-Pod fabric to propagate the endpoint

Correct Answer: C

Community vote distribution

C (80%)

D (20%)

🗳️ 👤 **Narbledeath** 8 months, 2 weeks ago

Selected Answer: C

C.

<https://unofficialaciguide.com/2017/09/25/configuring-aci-fabric-bgp-route-reflector/>
upvoted 2 times

🗳️ 👤 **Dash** 8 months, 3 weeks ago

Selected Answer: C

ACI uses Multi-Protocol BGP (MP-BGP) with VPNv4 in the ACI infra VRF (overlay-1 VRF) to distribute external routes from a border leaf to other leaf switches.

upvoted 2 times

🗳️ 👤 **7korn7** 8 months, 3 weeks ago

Selected Answer: D

D is the correct

upvoted 1 times

🗳️ 👤 **imanus** 7 months, 2 weeks ago

no, this happens in multi-site and sentence says multi-pod

upvoted 2 times

What are two descriptions of ACI Multi-Site? (Choose two.)

- A. Routers in the Inter-Site network must run OSPF, DHCP relay, and MP-BGP
- B. ACI Multi-Site is a solution that allows one APIC cluster to manage multiple ACI sites
- C. The Multi-Site orchestrator must be directly attached to one ACI leaf
- D. ACI Multi-Site is a solution that supports a dedicated APIC cluster per site
- E. The Inter-Site network routers should run OSPF to establish peering with the spines

Correct Answer: DE

Community vote distribution

DE (75%)

BE (25%)

🗳️ 👤 **designated** 7 months, 2 weeks ago

Selected Answer: DE

D, E are correct
upvoted 1 times

🗳️ 👤 **zelya19** 1 year, 3 months ago

Selected Answer: DE

Initially, OSPF was the only supported control plane, but from NDO Release 3.5(1) BGP was also added as a viable option.

Note: The use of BGP also requires the deployment of Cisco ACI Release 5.2(1).

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html#CiscoACIMultiSiteunderlaycontrolplane>

upvoted 1 times

🗳️ 👤 **Huberttheman** 1 year, 4 months ago

Selected Answer: DE

Multi site has a apic cluster per site but both can be managed from MSO
upvoted 1 times

🗳️ 👤 **sailorsoul** 1 year, 5 months ago

Selected Answer: DE

DE
multi site, not multi pod
upvoted 1 times

🗳️ 👤 **shark1989** 1 year, 5 months ago

Selected Answer: BE

This should be the right answer since b would be the most appropriate answer instead of D as APIC cluster need not be site specific but instead spread between multiple sites.

upvoted 1 times

DRAG DROP

-

An engineer must configure VMM domain integration on a Cisco UCS B-Series server that is connected to a Cisco ACI fabric. Drag and drop the products used to create VMM domain from the bottom into the sequence in which they should be implemented at the top. Products are used more than once.

On the interface, create a dynamic VLAN pool.
On the interface, create a VMware vCenter domain.
On the interface, create a vCenter/vShield controller.
On the user interface, verify that the VMware vDS is created.

APIC

vCenter

UCS Manager

Correct Answer:

On the interface, create a dynamic VLAN pool.
On the interface, create a VMware vCenter domain.
On the interface, create a vCenter/vShield controller.
On the user interface, verify that the VMware vDS is created.

APIC

vCenter

UCS Manager

 **designated** 7 months, 2 weeks ago

It is correct

APIC

APIC

APIC

vCenter

upvoted 1 times

The company ESXi infrastructure is hosted on the Cisco UCS-B Blade Servers. The company decided to take advantage of ACI VMM integration to enable consistent enforcement of policies across virtual and physical workloads. The requirement is to prevent the packet loss between the distributed virtual switch and the ACI fabric. Which setting must be implemented on a vSwitch policy to accomplish this goal?

- A. Static Channel
- B. MAC Pinning
- C. LACP
- D. LLDP

Correct Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **designated** 7 months, 2 weeks ago

Selected Answer: B

distributed virtual switch and the ACI fabric = MAC Pinning

upvoted 1 times

🗳️ 👤 **Marinheiro** 1 year, 1 month ago

Selected Answer: B

If you configure a mac-pinning policy, it programs the port groups to use this mechanism. This is very important in order to prevent packet loss.

<https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/118965-config-vmm-aci-ucs-00.html>

upvoted 2 times

🗳️ 👤 **SysAd** 1 year, 4 months ago

not C?

upvoted 1 times

An engineer is configuring ACI VMM domain integration with Cisco UCS-B Series. Which type of port channel policy must be configured in the vSwitch policy?

- A. LACP Active
- B. MAC Pinning
- C. LACP Passive
- D. MAC Pinning-Physical-NIC-load

Correct Answer: B

Community vote distribution

B (100%)

🗳️ **designated** 7 months, 1 week ago

Selected Answer: B

UCS-B = Mac-Pinning

"[...] the only supported load balancing mechanism when UCS B series is used is Route Based on Originating Virtual Port. If you configure a mac-pinning policy, it programs the port groups to use this mechanism. This is very important in order to prevent packet loss."

upvoted 1 times

🗳️ **Mr_Certifiable** 2 years, 1 month ago

By default on the DVS, the Discovery Protocol used is LLDP. This is fine for any servers that support LLDP, but the UCS B series blades only support LLDP on UCSM version 2.2(4b) and later. Because of this, ESXi cannot report LLDP information to the APIC, unless you are on the correct code.

As an alternative to LLDP, use CDP in order to discover the hosts. In order to get the DVS to use CDP, configure a vSwitch policy on the VMM Domain that has CDP enabled and LLDP disabled.

Along with this, the only supported load balancing mechanism when UCS B series is used is Route Based on Originating Virtual Port. If you configure a mac-pinning policy, it programs the port groups to use this mechanism. This is very important in order to prevent packet loss.

<https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/118965-config-vmm-aci-ucs-00.html>

upvoted 2 times

🗳️ **imamus** 2 years, 1 month ago

Selected Answer: B

<https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/118965-config-vmm-aci-ucs-00.html>

upvoted 2 times

A Cisco ACI is integrated with a VMware vSphere environment. The port groups must be created automatically in vSphere and propagated to hypervisors when created in the ACI environment. Which action accomplishes this goal?

- A. Create the port groups on the vCenter that reflect the EPG names in the APIC
- B. Assign the uplinks of the ESXi hosts to the vDS that the APIC created
- C. Configure contracts for the EPGs that are required on the ESXi hosts
- D. Associate the VMM domain with the EPGs that must be available in vCenter

Correct Answer: D

Community vote distribution

D (100%)

🗉 👤 **zelya19** 9 months, 4 weeks ago

Selected Answer: D

Cisco APIC Term - Endpoint group (EPG), VMware Term - Port group

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/virtualization/Cisco-ACI-Virtualization-Guide-401/Cisco-ACI-Virtualization-Guide-401_chapter_011.html#reference_AADBC61F737940D0A5C0342053E5942C

upvoted 1 times

A network engineer is integrating a new Hyperflex storage duster into an existing Cisco ACI fabric. The Hyperflex cluster must be managed by vCenter, so a new vSphere Distributed switch must be created. In addition, the hardware discovery must be performed by a vendor-neutral discovery protocol. Which set of steps meets these requirements?

- A. Configure an Interface Policy group, select CDP, and apply it to the desired interfaces.
Enter the vCenter IP and credentials in the Create vCenter Controller dialog box.
In the Create VMware VMM domain dialog box, select Read-Only Mode.
- B. Configure an Interface Policy group, select LLDP, and apply it to the selected interfaces.
Create a VLAN pool, add it to the VMware VMM domain, and include the appropriate interfaces.
Enter the vCenter IP and credentials in the Create vCenter Controller dialog box.
- C. Configure a Switch Policy group, select LLDP, and apply it to the indicated interfaces.
Set up a VMware VMM domain and apply it to the appropriate interfaces.
Enter the APIC management IP and credentials in the Create vCenter Controller dialog box.
- D. Configure an Interface Policy group, select CDP, and apply it to the designated interfaces.
Create a VMware VMM domain, add it to the VLAN pool, and associate it to the designated interfaces.
Select Read Only Mode in the Create VMware VMM domain dialog box.

Correct Answer: B

  **prospio971** 5 months ago

Selected Answer: B

LLDP is the vendor-neutral discovery protocol that should be used for hardware discovery in this scenario, as it allows for the discovery of non-Cisco devices and is supported by both ACI and Hyperflex.


upvoted 2 times

In-band is currently configured and used to manage the Cisco ACI fabric. The requirement is for leaf and spine switches to use out-of-band management for NTP protocol. Which action accomplishes this goal?

- A. Select Out-of-Band as Management EPG in the default DateTimePolicy.
- B. Create an Override Policy with NTP Out-of-Band for leaf and spine switches.
- C. Change the interface used for APIC external connectivity to ooband.
- D. Add a new filter to the utilized Out-of-Band-Contract to allow NTP protocol.

Correct Answer: A

Community vote distribution



🗳️ 👤 **zelya19** 9 months, 4 weeks ago

Selected Answer: A

Fabric > Fabric Policies > Policies > Pod > Date and Time > select OOB Management EPG when adding NTP servers
upvoted 1 times

🗳️ 👤 **supunit21** 1 year, 5 months ago

<https://www.linkedin.com/pulse/connecting-cisco-aci-mstp-lan-islam-elbaz/>
upvoted 1 times

🗳️ 👤 **asd248402** 1 year, 7 months ago

Selected Answer: A

Check on my apic the A is correct
upvoted 1 times

🗳️ 👤 **frzzt** 1 year, 9 months ago

Selected Answer: A

Ill go with A. As per this reference doc:

In the Management EPG drop-down list, if the NTP server is reachable by all nodes on the fabric through out-of-band management, choose Out-of-Band. If you have deployed in-band management, see the details about In-Band Management NTP. Click OK.

In-band IP addressing used within the ACI fabric is not reachable from anywhere outside the fabric. To leverage an NTP server external to the fabric with in-band management, construct a policy to enable this communication..

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/3-x/basic_config/b_APIC_Basic_Config_Guide_3_x/b_APIC_Basic_Config_Guide_3_x_chapter_0111.html

Since questions says use OOB not INB, the answer should be A.
upvoted 3 times

🗳️ 👤 **frzzt** 1 year, 9 months ago

Tricky, it could be both. Looking for more input
upvoted 1 times

🗳️ 👤 **VY01** 1 year, 9 months ago

Selected Answer: C

I think answer is C Ref. https://www.cisco.com/c/en/us/td/docs/security/workload_security/tetration-analytics/sw/config/cisco-aci-in-band-management-configuration-for-cisco-tetration.html
upvoted 1 times

DRAG DROP

An engineer must configure RADIUS authentication with Cisco ACI for remote authentication with out-of-band management access. Drag and drop the RADIUS configuration steps from the left into the required implementation order on the right. Not all steps are used.

Specify and set the Cisco APIC connectivity preferences to ooband

Create the RADIUS provider group

Set the Cisco APIC connectivity preferences to ooband

Create the login domain for RADIUS

Set the Cisco APIC connectivity preferences to inband

Create the RADIUS provider

step 1

step 2

step 3

step 4

Correct Answer:

Specify and set the Cisco APIC connectivity preferences to ooband

Create the RADIUS provider group

Set the Cisco APIC connectivity preferences to ooband

Create the login domain for RADIUS

Set the Cisco APIC connectivity preferences to inband

Create the RADIUS provider

step 1

step 2


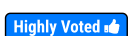
step 3

step 4

 **frzzt**  2 years, 3 months ago

1. OOB interface
2. Radius provider
3. Radius provider group
4. Radius login domain

This is correct
upvoted 6 times

 **Macc10**  1 year, 8 months ago

- Step 1. Create the radius provider
- Step 2. Create the radius provider group
- Step 3. Create the login domain for radius
- Step 4. Set the Cisco APIC connectivity preferences to oob.

<https://howtoaci.com/2018/05/21/tacacs-configuration-in-aci/>

Step 4 - strange answers to specify and set or just set - you simply change the connectivity preference - System -> System Settings -> APIC Connectivity Preferences so I'll go with just set
upvoted 5 times

🗨️ 👤 **designated** Most Recent 7 months, 1 week ago

Step "Set the Cisco APIC connectivity preferences to oob" can be done at first or last.. doesn't matter this order.
upvoted 1 times

🗨️ 👤 **imamus** 2 years, 1 month ago

https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/basic-configuration/cisco-apic-basic-configuration-guide-51x/m_aaa.html
upvoted 1 times

🗨️ 👤 **imamus** 2 years, 1 month ago

Hope this question is better framed in the exam. I'd chose:

- Step 1. Create the radius provider
- Step 2. Create the radius provider group
- Step 3. Create the login domain for radius
- Step 4. Specify and set the Cisco APIC connectivity preferences to oob. OR Set the Cisco APIC connectivity preferences to oob.

Honestly the only important thing is that the provider and provider group need to be created prior to the login domain, as they are necessary fields to be completed in it. Cant understand the difference between those two oob preferences... the only difference is one more verb in the sentence
upvoted 3 times

🗨️ 👤 **Rododendron2** 2 years, 1 month ago

What Q is this ? implementation order absolutely does not matter
upvoted 3 times

🗨️ 👤 **VY01** 2 years, 3 months ago

Agree with Ref.https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/basic-configuration/cisco-apic-basic-configuration-guide-51x/m_aaa.html
upvoted 1 times

An administrator must migrate the vSphere Management VMkernel of all ESXi hosts in the production cluster from the standard default virtual switch to a VDS that is integrated with APIC in a VMM domain. Which action must be completed in this scenario?

- A. The Management VMkernel EPG resolution must be set to Pre-Provision.
- B. The administrator must create an in-band VMM Management EPG before performing the migration.
- C. The administrator must set the Management VMkernel BD resolution immediacy to On-Demand.
- D. The VMkernel Management BD must be located under the Management Tenant.

Correct Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **Marinheiro** 7 months, 2 weeks ago

Selected Answer: A

When using pre-provision immediacy, policy is downloaded to Cisco ACI leaf switch regardless of CDP/LLDP neighborship. Even without a hypervisor host that is connected to the VMM switch.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/aci-fundamentals/Cisco-ACI-Fundamentals-42x/Cisco-ACI-Fundamentals-41X_chapter_01001.html

upvoted 1 times

A customer implements RBAC on a Cisco APIC using a Windows RADIUS server that is configured with network control policies. The APIC configuration is as follows:

- Tenant = TenantX
- Security Domain = TenantX-SD
- User = X

The customer requires User X to have access to TenantX only, without any extra privilege in the Cisco ACI fabric domain. Which Cisco AV pair must be implemented on the RADIUS server to meet these requirement?


- A. shell:domains = TenantX-SD/fabric-admin/,common//read-all
- B. shell:domains = TenantX-SD/tenant-admin
- C. shell:domains = TenantX-SD/tenant-ext-admin/,common//read-all
- D. shell:domains = TenantX-SD/tenant-admin/,common//read-all

Correct Answer: C

Community vote distribution

C (86%)

14%

 **kamel86** Highly Voted 1 year, 8 months ago

Selected Answer: C

As mentioned "without any extra privilege in the Cisco ACI fabric domain" so it should be C to manage the tenant only

<https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/security-configuration/cisco-apic-security-configuration-guide-release-52x/access-authentication-and-accounting-52x.html>

upvoted 6 times

 **udo2020** Most Recent 6 months, 2 weeks ago

Selected Answer: D

It is "shell:domains = TenantX-SD/tenant-admin/,common//read-all" because the role tenant-ext-admin has additional rights and this is not allowed.

upvoted 1 times

 **designated** 7 months, 1 week ago

Selected Answer: B

It is B!!!

upvoted 2 times

 **Rollizo** 8 months, 1 week ago

Selected Answer: B

It is B: tenant-admin "When assigned to a limited security domain, allows configuration of most attributes inside a tenant but does not allow changes to fabricwide settings that can potentially impact other tenants"

upvoted 2 times

 **Rollizo** 7 months, 3 weeks ago

Has to be C, tenant-admin allows access to fabric configuration inside the tenant itself

upvoted 1 times

 **zelya19** 1 year, 3 months ago

Selected Answer: C

Also work:

shell:domains = TenantX-SD/admin/,common//read-all

upvoted 1 times

 **sailorsoul** 1 year, 5 months ago

Selected Answer: C

tenant-admin have fabric privileges as well.

upvoted 1 times

 **Rododendron2** 2 years, 1 month ago

so practically, none of options, correct is:

shell:domains = TenantX-SD/admin/,common//read-all

upvoted 3 times

  **Rododendron2** 2 years, 1 month ago

Selected Answer: D

D. shell:domains = TenantX-SD/tenant-admin/,common//read-all

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/5-x/security/cisco-apic-security-configuration-guide-50x/m_access_authentication_and_accounting.html

upvoted 1 times

A network engineer must backup the PRODUCTION tenant. The configuration backup should be stored on the APIC using a markup language and contain all secure information. Which export policy must be used to meet these requirement?

A.

Name:

Description:

Format: ☒ json ☐ xml

Start Now: ☒ Yes ☐ No

Target DN:

Snapshot: ☐

Scheduler: ▼

Export Destination: ▼

Modify Global AES Encryption Settings: **Enabled**

B.

Name:

Description:

Format: ☒ json ☐ xml

Start Now: ☒ Yes ☐ No

Target DN:

Snapshot: ☐

Scheduler: ▼

Export Destination: ▼

Modify Global AES Encryption Settings: **Enabled**

C.

Name:

Description:

Format: ☒ json ☐ xml

Start Now: ☒ Yes ☐ No

Target DN:

Snapshot: ☒

Scheduler: ▼

Modify Global AES Encryption Settings: **Enabled**

D.

Name:

Description:


Format: ☐ json ☐ xml

Start Now: ☐ Yes ☐ No

Target DN:

Snapshot: ☒


Scheduler:

Modify Global AES Encryption Settings: **Enabled** 

Correct Answer: D


Community vote distribution



 **ashwind123** 3 months, 1 week ago

Selected Answer: D

D is right and B is wrong since question mentioned store in APIC
upvoted 1 times

 **udo2020** 6 months, 2 weeks ago

Selected Answer: D

Backup is a Snapshot when saved on APIC.
upvoted 1 times

 **sailorsoul** 11 months, 1 week ago

Selected Answer: D

Correction, it is D. Snapshot needs to be taken means it is saved locally.

https://www.labminutes.com/dc0022_aci_backup_restore_1

upvoted 1 times


 **sailorsoul** 11 months, 1 week ago

Selected Answer: B

It is B. Snapshot is not configuration.
upvoted 1 times

 **Rododendron2** 1 year, 7 months ago

D) DN format for tenant is: uni/tn-<name>
but shall be there encryption, otherwise no passwd imo
upvoted 1 times

 **7korn7** 1 year, 8 months ago

Selected Answer: D

Snapshots created for a single tenant do not include configurations from anywhere else in the fabric
upvoted 2 times

An engineer must create a backup of the Cisco ACI fabric for disaster recovery purposes. The backup must be transferred over a secure and encrypted transport. The backup file must contain all user and password related information. The engineer also wants to process and confirm the backup file validity by using a Python script. This requires the data structure to have a format similar to a Python dictionary. Which configuration set must be used to meet these requirements?

A. Under the Create Remote location settings, select Protocol: FTP

Under the Export policy, select -

- Format: XML

- Modify Global AES Encryption Settings: Enabled

B. Under the Create Remote location settings, select Protocol: FTP

Under the Export policy, select -

- Format: XML

- Modify Global AES Encryption Settings: Disabled

C. Under the Create Remote location settings, select Protocol: SCP

Under the Export policy, select -

- Format: JSON

- Modify Global AES Encryption Settings: Disabled

D. Under the Create Remote location settings, select Protocol: SCP

Under the Export policy, select -



- Format: JSON

- Modify Global AES Encryption Settings: Enabled

Correct Answer: D

Community vote distribution

D (100%)

  **hebdeb** 11 months, 2 weeks ago

Selected Answer: D

It is D


Secure and encrypted transport means: SCP

Under the Export policy, select -

Python format mean: JSON

Include user and password mean: AES Encryption Settings: Enabled

upvoted 2 times

  **Narbleddeath** 1 year, 2 months ago

D is correct

upvoted 2 times

The Application team reports that a previously existing port group has disappeared from vCenter. An engineer confirms that the VMM domain association for the EPG is no longer present. Which action determines which user is responsible for the change?

- A. Check the EPG audit logs for the "deletion" action and compare the affected object and user.
- B. Evaluate the potential faults that are raised for that EPG.
- C. Examine the health score and drill down to an object that affects the EPG combined score.
- D. Inspect the server logs to see who was logging in to the APIC during the last few hours.

Correct Answer: A

  **prospio971** 4 months, 2 weeks ago

Selected Answer: A

*Audit Logs: Audit logs in Cisco ACI record all user-initiated configuration changes, including actions such as creating, updating, or deleting objects like EPGs. These logs include details such as the name of the user who performed the action, a timestamp, and the specific object affected by the change.

*Relevance to the Scenario: Since the VMM domain association for the EPG has disappeared, checking the audit logs for a "deletion" action will reveal which user made this change and when it occurred. This is the most direct way to identify responsibility.

upvoted 1 times

An application team tells the Cisco ACI network administrator that it wants to monitor the statistics of the unicast and BUM traffic that are seen in a certain EPG. Which statement describes the collection statistics?

- A. All EPGs in the Cisco ACI tenant object must be enabled for statistics to be collected.
- B. Cisco ACI does not capture statistics at the EPG level. Only statistics that pass through ACI contracts can be monitored.
- C. EPG statistics can be collected only for VMM domains. If a physical domain exists, statistics are not collected.
- D. The collection of statistics is enabled on the EPG level by enabling the statistics for unicast and BUM traffic.

Correct Answer: D

Community vote distribution

D (75%)

B (25%)

🗳️ **zelya19** 9 months, 4 weeks ago

Selected Answer: D

Tenant > Application Profiles > Application EPGs > <EPG>: Stats
upvoted 2 times

🗳️ **sailorsoul** 11 months, 1 week ago

Selected Answer: B

I think it is B.

I have not found any document from Cisco supports monitoring on a EPG level.

upvoted 1 times

🗳️ **asd248402** 1 year, 6 months ago

Selected Answer: D

EPG-Stats tab can be configured so D
upvoted 1 times

🗳️ **Mr_Certifiable** 1 year, 7 months ago

Statistics data are gathered from a variety of sources, including interfaces, VLANs, EPGs, application profiles, ACL rules, tenants, or internal APIC processes. Statistics accumulate data in 5-minute, 15-minute, 1-hour, 1-day, 1-week, 1-month, 1-quarter, or 1-year sampling intervals. Shorter duration intervals feed longer intervals.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_01110.html

upvoted 1 times

🗳️ **Rododendron2** 1 year, 7 months ago

Selected Answer: D

It is D

upvoted 1 times

🗳️ **Narbledeath** 1 year, 8 months ago

Is this not D?

upvoted 1 times

An engineer must securely export Cisco APIC configuration snapshots to a secure, offsite location. The exported configuration must be transferred using an encrypted tunnel and encoded with a platform-agnostic data format that provides namespace support. Which configuration set must be used?

A. Policy: Export Policy -

Protocol: TLS -

Format: JSON

B. Policy: Import Policy -

Protocol: TLS -

Format: XML

C. Policy: Import Policy -

Protocol: SCP -

Format: JSON

D. Policy: Export Policy -

Protocol: SCP -

Format: XML

Correct Answer: D

  **gentlesol** 3 months ago

Selected Answer: D

Namespace support is XML

upvoted 1 times

A network engineer must allow secure access to the Cisco ACI out-of-band (OOB) management only from external subnets 10.0.0.0/24 and 192.168.20.0/25. Which configuration set accomplishes this goal?

- A. Create a L3Out in the MGMT tenant in OOB VRF.
Set External Management Network Instance Profile as a consumer of the OOB contract.
Create an External EPG with two subnet entries with the external subnets.
- B. Create a PBR service graph in the MGMT tenant.
Create a management Profile with the required OOB EPG.
Redirect all traffic going into ACI management to the external firewall.
Create two subnet entries under the OOB Bridge domain with the required subnets.
- C. Create an EPG and BD in the MGMT tenant in OOB VRF.
Set OOB VRF to provide the contract.
Set a new EPG to consume the OOB contract.
- D. Create an OOB contract that allows the required ports.
Provide the contract from the OOB EPG.
Consume the contract by the OOB External Management Network Instance Profile.
Create two subnet entries in the External Management Network Profile with the required subnets.

Correct Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Mr_Certifiable** 7 months ago

D -

Step 1 – Providing the contract

Tenant > Tenant mgmt > Node Management EPGs > Out-of-Band EPG default

- Under the “Provided Out-of-Band Contracts” in the policy window, provide the appropriate contract. (This could be a the default/common contract or a specific contract you have created and modified). Click Submit.

Tenant > Tenant mgmt > Node Management EPGs > Out-of-Band EPG default

Step 2 – Consuming the contract

Tenant > Tenant mgmt > External Management Network Instance Profiles > YourInstanceProfile

- Consume the same contract that you provided in the previous step.
- Enter the subnets that are allowed to have access to the APIC. (0.0.0.0/0 will permit all).

<https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/aci-guide-configuring-out-of-band-access-for-your-fabric.pdf>

upvoted 1 times

🗳️ 👤 **Narbleddeath** 8 months, 2 weeks ago

Selected Answer: D

<https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/aci-guide-configuring-out-of-band-access-for-your-fabric.pdf>

upvoted 3 times

Properties

Name: Rostan-oob-mgmt-01

Node Blocks:

From	To
1	1
101	101
102	102
201	201

Nodes Within The Policy:

ID	Name	Out-of-Band Management IP	Out-of-Band Management Gateway	In-Band Management IP	In-Band Management Gateway
1	apic 1	172.16.31.85 fe80::200:ff:fe0:0	172.16.31.254/24	192.168.11.1 fc00::1	
101	leaf-1	20.0.254.101 ::	20.0.254.1	::	
102	leaf-2	20.0.254.102 ::	20.0.254.1	::	
201	spine-1	172.16.31.84 ::	172.16.31.254/24	::	

Refer to the exhibit. A Cisco ACI fabric is using out-of-band management connectivity. The APIC must access a routable host with an IP address of 192.168.11.2. Which action accomplishes this goal?

- A. Change the switch APIC Connectivity Preference to in-band management
- B. Modify the Pod Profile to use the default Management Access Policy
- C. Add a Fabric Access Policy to allow management connections
- D. Remove the in-band management address from the APIC

Correct Answer: D

 **SvelterPrice97** 2 months, 1 week ago

Selected Answer: D

I think some info is missing. But I go for D.

If we choose A, changing to INBAND will not involve ROUTING as requested (You're in same network - And no gateway is configured infact).

Before choosing D you're using OOB to exit from the APIC, but if there is traffic returing I think it will use the In-Band, so removing the In-band IP could be the way.

The key word in the question is that need to route towards 192.168.x.x but a lot of info missing.

upvoted 1 times

 **prospio971** 5 months ago

Selected Answer: A

It is A:

You can make out-of-band management access the default management connectivity mode for the Cisco APIC server by clicking System > System Settings > APIC Connectivity Preferences. Then on the Connectivity Preferences page, click inband.

upvoted 1 times

 **designated** 7 months, 1 week ago

Selected Answer: D

no.. it is D.. replacing .1 to .2


upvoted 2 times

 **designated** 7 months, 1 week ago

Selected Answer: A

A. Change the switch APIC Connectivity Preference to in-band management following zelya19 steps.

upvoted 1 times

 **zelya19** 1 year, 3 months ago

System -> System Settings -> APIC Connectivity Preferences - set to "inband" by default

upvoted 1 times


A Cisco APIC is configured with RADIUS authentication as the default. The network administrator must ensure that users can access the APIC GUI with a local account if the RADIUS server is unreachable. Which action must be taken to accomplish this goal?

- A. Associate console authentication with the "RADIUS" realm
- B. Reference the "local" realm in the fallback domain
- C. Create an additional login domain that references local accounts
- D. Enable the fallback check with the default authentication domain

Correct Answer: C

Community vote distribution

C (100%)

 **sailorsoul** Highly Voted 11 months, 1 week ago

Selected Answer: C

Horrible question from Cisco again, both B & C will work.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_ACI-TACACS-config.html

But in the document, it says best practice is to create another local domain.

upvoted 5 times

 **prospio971** Most Recent 4 months, 2 weeks ago

Selected Answer: D

In Cisco ACI, you can configure fallback authentication to allow local user access in case the primary authentication method (such as RADIUS) is unreachable. To enable this feature, you need to configure the APIC to fallback to the local authentication domain if the RADIUS server is not available.

By enabling fallback authentication, when the APIC cannot reach the RADIUS server, it will automatically check local user credentials as a fallback.

upvoted 1 times

 **2eb1ea8** 5 months ago

Selected Answer: B

Allow access with a local account (if the RADIUS server is unreachable).

If we create an additional login domain that references local accounts it will give any user the possibility to choose to login with local credentials even if the RADIUS server is reachable.

upvoted 2 times

Create Configuration Export Policy

Name: !

Description:

Format: ☒ json ☐ xml

Start Now: ☒ Yes ☐ No

Target DN:

Snapshot: ☐

Scheduler:

Export Destination: !

Modify Global AES Encryption Settings: **Enabled** ☒

Refer to the exhibit. A customer must back up the current Cisco ACI configuration securely to the remote location using encryption and authentication. The backup job must run once per day. The customer's security policy mandates that any sensitive information including passwords must not be exported from the device. Which set of steps meets these requirements?

- A. Export destination using FTP protocol.
Use XML format.
- B. Export destination using FTP protocol.
Disable Global AES Encryption.
- C. Export destination using SCP protocol.
Disable Global AES Encryption.
- D. Export destination using SCP protocol.
Use XML format.

Correct Answer: C

Community vote distribution

C (100%)

Narbledeath Highly Voted 8 months, 2 weeks ago

Picking C.

AES is disabled because you don't want passwords or sensitive info to be exported.

upvoted 7 times

thinqtanklearningDOTcom Most Recent 7 months, 1 week ago

Selected Answer: C

Absolutely C based on the requirements. SCP provides authentication and encryption. And per Cisco documentation "To export hashed secure properties (passwords and certificates), AES encryption must be configured and enabled. While encryption is not enabled, any secure fields will not be exported. In this case re-importing the configuration will require all secure properties to be re-configured."

upvoted 4 times

Dash 8 months, 3 weeks ago

Shouldn't this be D? If you disable AES then data will not be encrypted!

upvoted 1 times

  **thingtanklearningDOTcom** 7 months, 1 week ago



See my comments above.

upvoted 1 times

An engineer must perform a Cisco ACI fabric upgrade that minimizes the impact on user traffic and allows only permitted users to perform an upgrade. Which two configuration steps should be taken to meet these requirements? (Choose two.)

- A. Grant tenant-ext-admin access to a user who performs an upgrade
- B. Divide Cisco APIC controllers into two or more maintenance groups
- C. Combine all switches into an upgrade group
- D. Grant the fabric administrator role to a user who performs an upgrade
- E. Divide switches into two or more maintenance groups

Correct Answer: *DE*

  **crooks_1988** 9 months, 3 weeks ago

D and E: E for obvious reasons, D because fabric-admin role can deal with firmware management
upvoted 4 times

Which routing protocol is supported between Cisco ACI spines and IPNs in a Cisco ACI Multi-Pod environment?

- A. OSPF
- B. IS-IS
- C. BGP
- D. EIGRP

Correct Answer: A

Community vote distribution

A (90%)

10%

 **KrishEngineer** Highly Voted 1 year, 3 months ago

OSPF is right choice
upvoted 6 times

 **hebdeb** Most Recent 11 months, 2 weeks ago

Selected Answer: A

It is A
<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html#InterPodConnectivityDeploymentConsiderations>
upvoted 1 times

 **S_1292_A** 12 months ago

Selected Answer: A

OSPF for sure as everybody else mentioning
upvoted 2 times

 **asd248402** 1 year ago

Selected Answer: A

OSPF for sure
upvoted 2 times

 **Rododendron2** 1 year ago


Selected Answer: A

between Cisco ACI spines and IPNs - sorry, yes OSPF
upvoted 2 times

 **Rododendron2** 1 year, 1 month ago

Selected Answer: C

It is MP-BGP
OSPF is for multi-site ACI , multi-pod runs MP-BGP between spines across IPN
<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html#InterPodsMPBGPCControlPlane10>
upvoted 1 times

 **7korn7** 1 year, 3 months ago

Selected Answer: A

OSPF is correct answer
upvoted 2 times

 **Dash** 1 year, 3 months ago

Definitely OSPF
upvoted 2 times

An engineer must deploy Cisco ACI across 10 geographically separated data centers. Which ACI site deployment feature enables the engineer to control which bridge domains contain Layer 2 flooding?

- A. GOLF
- B. Multi-Site
- C. Multi-Pod
- D. Stretched Fabric

Correct Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **designated** 7 months, 1 week ago

Selected Answer: B

Multi-Site is a Cisco ACI feature designed for managing multiple geographically dispersed ACI fabrics. It allows control over Layer 2 flooding at the bridge domain level, which is essential for maintaining network efficiency and isolation in large-scale deployments spanning multiple sites.

upvoted 1 times

🗨️ 👤 **zelya19** 1 year, 3 months ago

Selected Answer: B

The ACI stretched fabric design has been validated, and is hence supported, on up to three interconnected sites.



https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_kb-aci-stretched-fabric.html

upvoted 1 times

A customer creates Layer 3 connectivity to the outside network. However, only border leaf switches start receiving destination updates to other networks from the newly created L3Out. The updates must also be propagated to other Cisco ACI leaf switches. The L3Out is linked with the EPGs via a contract. Which action must be taken in the pod policy group to accomplish this goal?

- A. Apply a BGP route reflector policy.
- B. Enable a COOP policy.
- C. Configure an IS-IS policy.
- D. Implement an access management policy.

Correct Answer: A

  **ashwind123** 3 months, 2 weeks ago

Selected Answer: A

A is correct!


upvoted 1 times

A network administrator configures AAA inside the Cisco ACI fabric. The authentication goes through the local users if the TACACS+ server is not reachable. If the Cisco APIC is out of the cluster, the access must be granted through the fallback domain. Which configuration set meets these requirements?

- A. Ping Check: True -
Default Authentication Realm: Local
Fallback Check: True
- B. Ping Check: True -
Default Authentication Realm: TACACS+
Fallback Check: False
- C. Ping Check: False -
Default Authentication Realm: Local
Fallback Check: False
- D. Ping Check: False -
Default Authentication Realm: TACACS+
Fallback Check: True

Correct Answer: B

Community vote distribution



sailorsoul 11 months, 1 week ago

Selected Answer: B

Fallback check must be false.

upvoted 2 times

sailorsoul 11 months ago

https://www.labminutes.com/dc0019_aci_aaa_radius_tacacs_3

12:20

upvoted 1 times

thingtanklearningDOTcom 1 year, 7 months ago

Selected Answer: B

<https://howtoaci.com/2018/05/21/tacacs-configuration-in-aci/>

It is all documented here. And Cisco documentation also specifically states: Note: Make sure to leave/set the Fallback Check property to false. Setting the Fallback Check property to true may cause local logins to fail.

upvoted 2 times

Rododendron2 1 year, 7 months ago

Selected Answer: B

My tip is B

Unclear what is supposed be the ping check

Anyway Default Authentication Realm hall be TACACS+ and you shall not put fallback check on if required working fallback local authentication

Note: Make sure to leave/set the Fallback Check property to false. Setting the Fallback Check property to true may cause local logins to fail.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_ACI-TACACS-config.html

upvoted 3 times

Rododendron2 1 year, 7 months ago

How is done AAA server check if enabled check, but icmp check disabled ? What is the "secret check" that will allow the APIC out of cluster switch to fallback authentication ?

upvoted 1 times

thingtanklearningDOTcom 1 year, 7 months ago

Note: Make sure to leave/set the Fallback Check property to false. Setting the Fallback Check property to true may cause local logins to fail.

upvoted 1 times

  **Narbleddeath** 1 year, 8 months ago

Selected Answer: D

Correct

<https://community.cisco.com/t5/application-centric-infrastructure/should-you-always-be-able-to-login-using-fallback-domain/m-p/4502626>
upvoted 1 times

A Cisco ACI environment consists of multiple silent hosts that are often relocated between leaf switches. When the host is relocated, the bridge domain takes more than a few seconds to relearn the host's new location. The requirement is to minimize the relocation impact and make the ACI fabric relearn the new location of the host faster. Which action must be taken to meet these requirements?

- A. Set Unicast Routing to Enabled.
- B. Configure ARP Flooding to Enabled.
- C. Set L2 Unknown Unicast to Hardware Proxy.
- D. Configure IP Data-Plane Learning to No.

Correct Answer: B

Community vote distribution

B (89%)

11%


 **sailorsoul** 11 months, 1 week ago

Selected Answer: B

B

ARP glean only works for Unknown L3 packets.


upvoted 3 times

 **imamus** 1 year, 4 months ago

Selected Answer: B

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html#Silenthostsconsiderations>

upvoted 1 times

 **hebddeb** 1 year, 5 months ago

Selected Answer: B

It is B

A tangible difference between enabling and disabling ARP flooding occurs with silent host movements. Suppose that hardware proxy has been enabled on a bridge domain, ARP flooding has been disabled, and ACI has already learned a silent host in the BD through ARP gleaning. If the silent host moves from one location to another without notifying the new ACI leaf via GARP or some other mechanism, ACI switches continue to forward traffic intended for the silent IP address to the previous location until retention timers clear the endpoint from COOP. Until that point, if an endpoint sends ARP requests toward this silent host, ARP gleaning is not triggered because COOP considers the destination endpoint to be known. On the other hand, with ARP flooding enabled on the BD, ARP requests are flooded, and the silent host responds at its new location, enabling the new local leaf to learn the silent host and update COOP.

upvoted 3 times

 **thinqtanklearningDOTcom** 1 year, 7 months ago

Selected Answer: C

Hardware Proxy

By default, Layer 2 unknown unicast traffic is sent to the spine proxy. This behaviour is controlled by the hardware proxy option associated with a bridge domain: if the destination is not known, send the packet to the spine proxy; if the spine proxy also does not know the address, discard the packet (default mode).

The advantage of the hardware proxy mode is that no flooding occurs in the fabric. The potential disadvantage is that the fabric has to learn all the endpoint addresses.

With Cisco ACI, however, this is not a concern for virtual and physical servers that are part of the fabric: the database is built for scalability to millions of endpoints. However, if the fabric had to learn all the IP addresses coming from the Internet, it would clearly not scale.

upvoted 1 times

 **thinqtanklearningDOTcom** 1 year, 7 months ago

Flooding Mode

Alternatively, you can enable flooding mode: if the destination MAC address is not known, flood in the bridge domain. By default, ARP traffic is not flooded but sent to the destination endpoint. By enabling ARP flooding, ARP traffic is also flooded. A good use case for enabling ARP flooding



would be when the Default Gateway resides outside of the ACI Fabric. This non-optimal configuration will require ARP Flooding enabled on the BD.

This mode of operation is equivalent to that of a regular Layer 2 switch, except that in Cisco ACI this traffic is transported in the fabric as a Layer 3 frame with all the benefits of Layer 2 multi-pathing, fast convergence, and so on.

Hardware proxy and unknown unicast and ARP flooding are two opposite modes of operation. With hardware proxy disabled and without unicast and ARP flooding, Layer 2 switching would not work.

This option does not have any impact on what the mapping database actually learns; the mapping database is always populated for Layer 2 entries regardless of this configuration.



upvoted 1 times

  **thingtanklearningDOTcom** 1 year, 7 months ago
ARP Gleaning...



Forwarding is optimized:

- L2 Unknown Unicasts are sent to the Hardware Proxy
- L3 Unknown Multicasts are flooded
- Multi Destination frames are flooded within the BD
- ARP flooding is disabled

upvoted 1 times

  **thingtanklearningDOTcom** 1 year, 7 months ago
For this reason, I would choose option C (L2 Unknown Unicasts are sent to the Hardware Proxy)

upvoted 1 times

  **7korn7** 1 year, 8 months ago

Selected Answer: B

If you have silent hosts in your network, then ARP Flooding ON is a must.

upvoted 2 times

  **thingtanklearningDOTcom** 1 year, 7 months ago

Not necessarily. ARP flooding is only required if the following two conditions are met:

There is a silent host in a Bridge Domain

There is no IP address configured for the bridge domain in the same subnet as the silent host.

The reason for this is because ACI does ARP Gleaning.

ARP Gleaning ONLY works if the Bridge Domain (or EPG associated with the Bridge Domain) has been assigned an IP address on the same subnet with which it can source a Gleaning ARP.

The IP address assigned to the Bridge Domain does not have to be the default gateway IP – if you have a router or firewall attached that serves as a default gateway for an EPG and you DON'T want to turn on ARP flooding, assigning any IP address on that subnet to the Bridge Domain will ensure your hosts will find their default gateway.

upvoted 2 times

```

ID: 4295457803
Description: Failed to form relation to MO qosdppol-default of class qosDppPol in context
Severity: cleared
Affected Object: uni/tn-Raccoon_City/out-L3Out_Demo/lnodep-L3Out_Demo_nodeProfile/lifp-L3Out_Demo_interfaceProfile/rsingressQosDppPol
Delegated From:
Created: 2021-02-05T19:29:08.730+00:00
Code: F2044
Type: Config
Cause: resolution-failed
Change Set: state (Old: missing-target, New: formed), tDn (Old: , New: uni/tn-common/qosdppol-default)
Action: modification
Domain: Infra
Life Cycle: Retaining
Count Fault Occurred: 1
Acknowledgement Status: false

```

Refer to the exhibit. An engineer configures an L3Out but receives the error presented. Which action clears the fault?


- A. Acknowledge the QoS-related error.
- B. Associate a custom QoS class.
- C. Create a custom QoS policy.
- D. Set the QoS policy to Level 3.

Correct Answer: A

Community vote distribution

A (75%)

B (25%)

  **prospio971** 4 months, 3 weeks ago

Selected Answer: C

The Code F2044 error typically occurs when there is a mismatch or conflict in QoS policies applied to the L3Out configuration. To address this issue, a custom QoS policy should be created and applied directly to the L3Out interface. This ensures that the QoS settings align with the specific requirements of the traffic being handled.

Why Not the Other Options?

A. Acknowledge the QoS-related error:

Acknowledging the error in the APIC GUI won't resolve the underlying configuration issue. The fault will persist until the QoS policy is corrected.

B. Associate a custom QoS class:

While associating a QoS class is part of the solution, you first need to create the custom QoS policy itself, which defines how the classes behave.

D. Set the QoS policy to Level 3:

Setting a QoS policy to a specific level (e.g., Level 3) may not align with your specific traffic requirements and won't address the root cause if a custom policy is needed.

upvoted 1 times

  **designated** 7 months, 1 week ago



Selected Answer: B

1. Fault Lifecycle: The fault is marked as "Retaining". This indicates that the system has recognized the issue but has not yet fully resolved it.

2. Cause: The cause is "resolution-failed", specifically due to a missing QoS policy relation (qosDppPol-default).

3. Fault Type: It's a configuration fault.

upvoted 1 times

  **zelya19** 1 year, 3 months ago

Selected Answer: A

Life Cycle: Retaining, Severity: cleared - the fault will be deleted upon acknowledgement

upvoted 1 times

  **Defilet** 2 years ago

It is A because the severity level is cleared.

You need to acknowledge it in order to clear it before the retention timer expires.

upvoted 1 times



  **Narbleddeath** 2 years, 2 months ago

Selected Answer: A

Id say A.



Because performing those actions will prevent the fault from appearing again in the future, but will not clear the fault that is already present.

upvoted 2 times

  **frzzt** 2 years, 3 months ago

or C? Create custom policy?

upvoted 2 times

  **VY01** 2 years, 3 months ago

Selected Answer: B

B is correct? <https://pubhub.devnetcloud.com/media/apic-mim-ref-411/docs/FAULT-F2044.html> the fault about missing-target so need to associate a new policy QoS?

upvoted 1 times

A customer must upgrade the Cisco ACI fabric to use a feature from the new code release. However, there is no direct path from the current release to the desired one. Based on the Cisco APIC Upgrade/Downgrade Support Matrix, the administrator must go through one intermediate release.

Which set of steps must be taken to upgrade the fabric to the new release?

- A. 1. Upgrade the APICs to an interim release.
2. Upgrade the switches to an interim release.
3. Upgrade the APICs to the targeted release.
4. Upgrade the leaf and spine switches to the targeted release.
- B. 1. Upgrade the APICs to an interim release and then switches to an interim release.
2. When all switches are operational, upgrade leaf switches to the targeted release.
3. Upgrade the spine switches to the targeted release.
4. Upgrade the APICs to the targeted release.
- C. 1. Upgrade the APICs to an interim release.
2. Upgrade the leaf switches directly to the targeted release.
3. Upgrade the spine switches directly to the targeted release.
4. Upgrade the APICs to the targeted release.
- D. 1. Upgrade the APICs directly to the targeted release.
2. Upgrade the switches to an interim release.
3. When all switches are operational, upgrade the leaf switches to the targeted release.
4. Upgrade the spine switches to the targeted release.

Correct Answer: A

Community vote distribution

A (100%)

 **thingtanklearningDOTcom** 7 months, 1 week ago

Selected Answer: A

The basic process is:

Get files from Cisco onto a HTTP/SCP server and then uploaded to APIC

Get APICs upgraded

Wait for things to stabilise.

Get Leaf/Spines upgraded

Wait for things to stabilise.

When upgrading or downgrading to a release that does not have a direct path from your current release, you must upgrade or downgrade all the APICs and switches to an intermediate supported release to which there is a direct path, then upgrade or downgrade from that release to your desired release. Sometimes, you must move through multiple intermediate releases before being able to get to your desired release, upgrading or downgrading both the APICs and switches to the same release each time.

upvoted 1 times

```
leaf-102# show interface brief
!snip
```

Port-channel Interface	VLAN	Type	Mode	Status	Reason	Speed	Protocol
Po3	46	eth	trunk	down	mac-pinning	inherit(D	lACP
Po11	--	eth	fabric	up	none	10G(D)	none
Po12	0	eth	trunk	down	mcp-loop-err-disable	inherit(D	none

Refer to the exhibit. Which two configuration steps are completed before this output is generated? (Choose two.)

- A. MCP policy for the interface policy group for Port-channel 12 is enabled.
- B. MCP Instance Policy default in the global access policies is enabled.
- C. Error Disabled Recovery Policy for Loop Indication by MCP is set to True.
- D. BPDU Guard is enabled for the interface policy group for Port-channel 12.
- E. Spanning Tree Policy Region STP_4CAF232E48FF20 is added to the spanning-tree policy of the switch.

Correct Answer: AB

Community vote distribution

AB (100%)

as248402 7 months, 2 weeks ago

Selected Answer: AB

a and b

upvoted 2 times

Rododendron2 7 months, 3 weeks ago

Selected Answer: AB

agree with A and B

upvoted 1 times

7korn7 8 months, 3 weeks ago

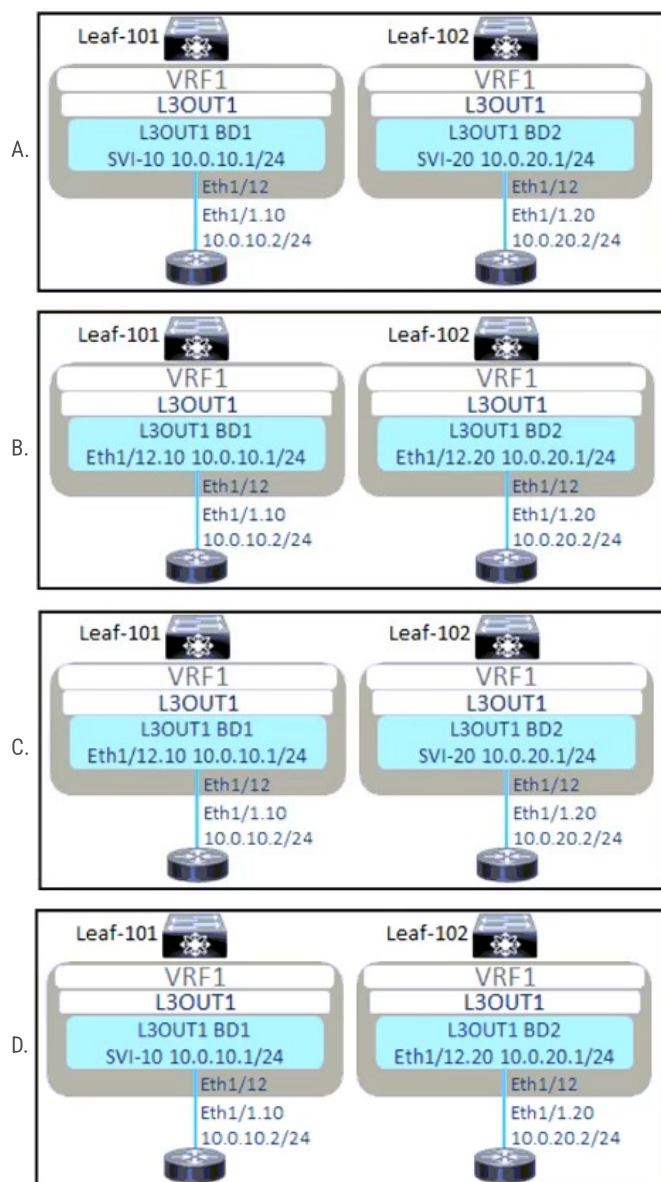
Selected Answer: AB

A and B

<https://unofficialaciguide.com/2018/03/27/using-mcp-miscabling-protocol-for-aci/>

upvoted 3 times

The customer is looking for redundant interconnection of the existing network to the new ACI fabric. Unicast and multicast traffic must be routed between the two networks. Which L3Out implementation meets these requirements?



Correct Answer: A

Community vote distribution

B (100%)

frzzt 2 years, 3 months ago

Isn't this B? In older versions (which exam references mostly) SVI's don't allow multicast routing?
upvoted 8 times

imanus 2 years, 1 month ago

I agree with you:

PIMv4/PIM6 is supported on Layer 3 Out routed interfaces and routed subinterfaces including Layer 3 port-channel interfaces. In the 5.2(2) release and earlier, PIMv4/PIM6 is not supported on Layer 3 Out SVI interfaces. In the 5.2(3) release and later, PIMv4/PIM6 is supported on Layer 3 Out SVI interfaces

https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/l3-configuration/cisco-apic-layer-3-networking-configuration-guide-52x/ip-multicast-layer3-config-52x.html#id_21570

upvoted 1 times

Larsom 4 months, 2 weeks ago

Selected Answer: A

Cisco ACI does not support multicast on Routed Subinterfaces or Layer 3 Port (Routed Physical Interfaces). Multicast traffic requires PIM, which is only supported on SVI-based L3Outs. Routed subinterfaces and direct routed interfaces do not support PIM in ACI.



upvoted 1 times

  **designated** 7 months, 1 week ago

Selected Answer: B

Only routed interface and subinterface allow multicast!!!!

upvoted 2 times

  **hebdeb** 1 year, 11 months ago

Selected Answer: B

Only routed interface and subinterface allow multicast

upvoted 2 times

A network engineer configures the Cisco ACI fabric to connect to vCenter with these requirements:

- Port groups must be automatically created on the distributed virtual switch.
- Port groups must use the VLAN allocation in the range between 20-30.
- The deployment must optimize the CAM space on the leaf switches.


Which set of actions meets these criteria?

- A. Create a dynamic VLAN pool with the VLAN range of 20-30.
Create a VMM domain and associate it with the VLAN pool.
Create the EPG and associate the domain.
Set the deployment immediacy to On Demand.
- B. Create a dynamic VLAN pool with the VLAN range of 20-30.
Create a physical domain and associate it with the VLAN pool.
Create the EPG and associate the domain.
Set the deployment immediacy to On Demand.
- C. Create a static VLAN pool with the VLAN range of 20-30.
Create a physical domain and associate it with the VLAN pool.
Create the EPG and associate the domain.
Set the deployment immediacy to Immediate.
- D. Create a static VLAN pool with the VLAN range of 20-30.
Create a VMM domain and associate it with the VLAN pool.
Create the EPG and associate the domain.
Set the deployment immediacy to Immediate.

Correct Answer: A

Community vote distribution

A (100%)

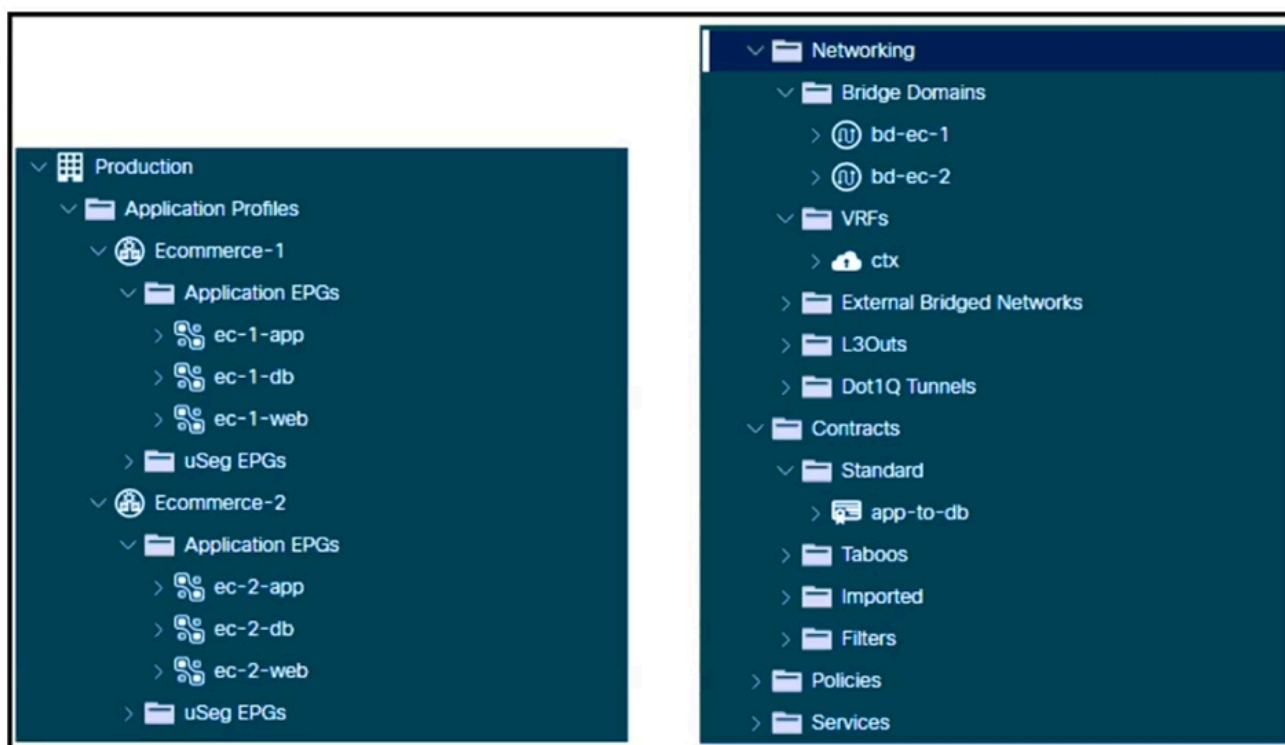
 **Narbleddeath** 8 months, 2 weeks ago

Selected Answer: A

A.

<https://community.cisco.com/t5/data-center-and-cloud-knowledge-base/apic-policy-deployment-and-resolution-immediacy-for-avs-vmm/ta-p/3134812>

upvoted 3 times



Refer to the exhibit. A Cisco ACI environment hosts two e-commerce applications. The default contract from a common tenant between different application tiers is used, and the applications work as expected. The customer wants to move to more specific contracts to prevent unwanted traffic between EPGs. A network administrator creates the app-to-db contract to meet this objective for the application and database tiers. The application EPGs must communicate only with their respective database EPGs. How should this contract be configured to meet this requirement?

- A. Set the app-to-db scope to Global.
- B. Set the app-to-db scope to Application Profile.
- C. Implement the app-to-db scope as VRF.
- D. Implement the app-to-db as a Taboo contract.

Correct Answer: B

Community vote distribution

B (100%)

zelya19 9 months, 4 weeks ago

Selected Answer: B

Scope: VRF would allow cross-AP communication
upvoted 1 times

Mr_Certifiable 1 year, 7 months ago

B
Taboo contracts can be used to deny specific traffic that is otherwise allowed by contracts. The traffic to be dropped matches a pattern (such as, any EPG, a specific EPG, or traffic matching a filter). Taboo rules are unidirectional, denying any matching traffic coming toward an EPG that provides the contract.

With Cisco APIC Release 3.2(x) and switches with names that end in EX or FX, you can alternatively use a subject Deny action or Contract or Subject Exception in a standard contract to block traffic with specified patterns.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401/b_Cisco_APIC_Security_Guide_chapter_01010.html



upvoted 1 times

Icarus7322 1 year, 7 months ago

Selected Answer: B

Application: A contract will only program rules between EPGs that are defined within the same application profile. Use of the same contract across other application profile EPGs will not allow for crosstalk between them.

upvoted 1 times



  **Lorygru** 1 year, 8 months ago

Selected Answer: B

"The application EPGs must communicate only with their respective database EPGs"

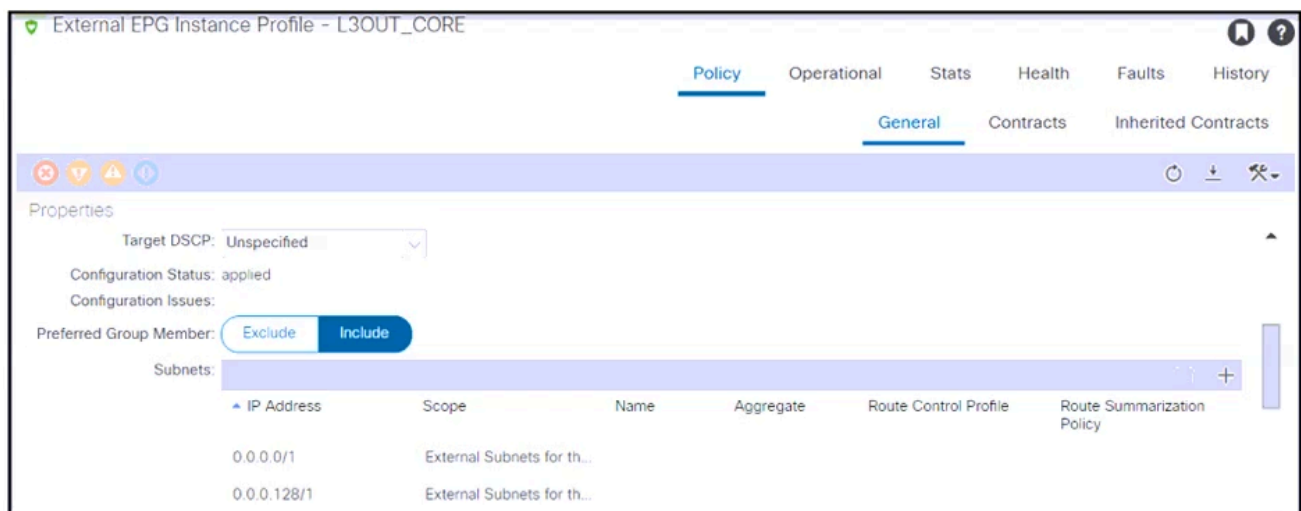
B is correct, setting the scope to application profile allow the communication between app and db on the same AP only

upvoted 4 times

  **frztt** 1 year, 9 months ago

Shouldn't this be C? apply the contract to appropriate EPG's inside the VRF. Why use taboo in this case if it only denies traffic?

upvoted 2 times



Refer to the exhibit. An engineer configured subnets on the external EPG called L3OUT_CORE. The external endpoints in the 10.1.0.0/24 subnet can reach internal endpoints, but the external endpoints in the 172.16.1.0/24 subnet are unreachable. Which set of actions enables the connectivity?

- A. Delete both external EPG subnets.
Create the 0.0.0.0/1 subnet.
- B. Delete the external EPG subnet 0.0.0.128/1.
Create the 128.0.0.0/1 subnet.
- C. Delete both external EPG subnets.
Create the 0.0.0.0/0 subnet.
- D. Delete the external EPG subnet 0.0.0.0/0.
Create the 0.0.0.0/128 subnet.

Correct Answer: B

Community vote distribution

B (60%)

C (40%)

prospio971 4 months, 3 weeks ago

Selected Answer: C

Create the 0.0.0.0/0 subnet.

Because Create 128.0.0.0/1: This only covers the second half of the IPv4 address space (128.0.0.0 to 255.255.255.255) and does not include 10.1.0.0/24, which falls in the first half.

upvoted 1 times

prospio971 4 months, 3 weeks ago

The 0.0.0.0/0 matches every IP address, whereas 0.0.0.0/1 only matches half of them (0.0.0.0-127.255.255.255) and requires 128.0.0.0/1 as its pair to match the rest (128.0.0.0-255.255.255.255).

upvoted 1 times

prospio971 4 months, 3 weeks ago

Sorry, it's B. Because, we don't delete 0.0.0.0/1.

upvoted 1 times

designated 7 months ago

Selected Answer: B

Hard question

Adding internal EPGs and External EPGs (associated to L3Outs) to the Preferred Group allows to enable free east-west and northsouth connectivity

- When adding the Ext-EPG to the Preferred Group:

• Can't use 0.0.0.0/0 for classification, needs more specific prefixes

As workaround it is possible to use 0.0.0.0/1 and 128.0.0.0/1 to achieve the same result.

<https://www.ciscolive.com/c/dam/r/ciscolive/global-event/docs/2022/pdf/BRKDCN-2480b.pdf>

upvoted 2 times

🗳️ 👤 **Huberttheman** 1 year, 4 months ago

Selected Answer: B

I vote b

upvoted 1 times

🗳️ 👤 **Macc10** 1 year, 8 months ago

Selected Answer: B

Answer is B - https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/L3-configuration/Cisco-APIC-Layer-3-Networking-Configuration-Guide-401/Cisco-APIC-Layer-3-Networking-Configuration-Guide-401_chapter_011000.html

upvoted 1 times

🗳️ 👤 **Macc10** 1 year, 9 months ago

Selected Answer: B

Answer is B -

<https://bst.cloudapps.cisco.com/bugsearch/bug/CSCve28564>

<https://unofficialaciguide.com/2017/10/24/transit-routing-use-case-eigrp-with-routed-interfaces/>

When the L3 Out Network EPG has 0.0.0.0/0 as a prefix entry, "Preferred Group Member" Include is not configurable - One workaround is to break 0.0.0.0/0 to 0.0.0.0/1 and 128.0.0.0/1

The correct Answer is B

upvoted 2 times

🗳️ 👤 **Ron_Berserker** 1 year, 9 months ago

Selected Answer: C

According this ACI allows whole range 0/0: "As part of the L3Out configuration, these subnets should be defined as external networks. Alternatively, an external network could be defined as 0.0.0.0/0 to cover all possible destinations."

https://www.networklife.net/images/sheets/Networklife_CheatSheet_ACI_04_L3out_v1.1.pdf

upvoted 1 times

🗳️ 👤 **S_1292_A** 1 year, 12 months ago

Selected Answer: B

We have preferred groups running, so the external network should be splitted

upvoted 1 times

🗳️ 👤 **Rododendron2** 2 years ago

Selected Answer: B

B is right

ACI does not allow whole range 0.0.0.0/0 , so required split ranges.

upvoted 1 times

🗳️ 👤 **Mr_Certifiable** 2 years, 1 month ago

B - is more granular

The 0.0.0.0/0 matches every IP address, whereas 0.0.0.0/1 only matches half of them (0.0.0.0-127.255.255.255) and requires 128.0.0.0/1 as its pair to match the rest (128.0.0.0-255.255.255.255).

<https://serverfault.com/questions/1100250/what-is-the-difference-between-0-0-0-0-0-and-0-0-0-0-1>

upvoted 1 times

🗳️ 👤 **imanus** 2 years, 1 month ago

Selected Answer: C

its C, this allows any subnet to communicate with internal EPGs (given the right contracts of course). Those two subnets in the screenshot dont even make any sense at all.

upvoted 1 times

🗳️ 👤 **Narbledeath** 2 years, 2 months ago

Selected Answer: C

is this not C??

<https://unofficialaciguide.com/2019/11/08/understanding-scope-of-prefixes-in-l3-out-external-epg-in-aci/>
upvoted 2 times

🗨️ 👤 **Lorygru** 2 years, 2 months ago
why not B? also B can work isn't it?
upvoted 2 times

An engineer deploys a two-pod Cisco ACI Multi-Pod environment. Why should no more than two Cisco APIC controllers be deployed in the same pod?

- A. to enable equal capacity to scale in each pod
- B. to avoid losing all replicas of a shard if a pod fails
- C. to avoid hair-pinning traffic that is destined for the primary APIC controller between pods
- D. to ensure that all nodes in all pods have local access to a controller

Correct Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **thingtanklearningDOTcom** 7 months, 1 week ago

Selected Answer: B

<https://www.ciscopress.com/articles/article.asp?p=3150964&seqNum=4>

"To ensure that a total loss of a given pod does not result in the loss of all shards for a given attribute, Cisco recommends that no more than two APICs be placed in a single pod."

upvoted 3 times

🗨️ 👤 **VY01** 9 months, 3 weeks ago

Selected Answer: B

It should B.

upvoted 3 times

Create Subnet

?✕

Gateway IP:
address/mask

Treat as virtual IP address: ☐

Make this IP address primary: ☐

Scope: ☒ Private to VRF
☐ Advertised Externally
☐ Shared between VRFs

Description:

Subnet Control: ☐ No Default SVI Gateway
☒ Querier IP

L3 Out for Route Profile: ▼

Route Profile: ▼

ND RA Prefix policy: ▼

CancelSubmit

Refer to the exhibit. An engineer configures communication between the EPGs in different tenants. Which action should be taken to create the subnet?

- A. Change Scope to Shared between VRFs.
- B. Leave Scope set to Private to VRF.
- C. Add the L3Out for Route Profile value.
- D. Change Scope to Advertised Externally.

Correct Answer: A

Community vote distribution

A (100%)

 **Mr_Certifiable** 7 months ago

A ->

Shared Service (Route Leaking) enables traffic between endpoints in different VRFs.

- A shared service EPG provider is an EPG that provides a contract consumed by an EPG in a different VRF

Restrictions

- Provider Subnet must be defined under the provider EPG
- Both provider and consumer subnets must have scope set to shared

- contract needs correct scope
- VzAny not supported as provider

Shared Service (Route Leaking) enables traffic between endpoints in different VRFs.

- A shared service EPG provider is an EPG that provides a contract consumed by an EPG in a different VRF

Restrictions

- Provider Subnet must be defined under the provider EPG
- Both provider and consumer subnets must have scope set to shared
- contract needs correct scope
- VzAny not supported as provider

Scope:

☒ Private to VRF

☒ Advertise Externally

-> Share Between VRFs

page 136

<https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2019/pdf/BRKACI-3101.pdf>

upvoted 2 times

  **thingtanklearningDOTcom** 7 months, 1 week ago

Selected Answer: A

Private to VRF = This subnet will only be routed internal to the ACI fabric. This is to say it won't be advertised via any L3outs.



Advertise Externally = The opposite of Private. This means you want this subnet to participate in any associated L3outs.

Shared Between VRFs = This is used when you want to do route leaking between different VRFs

Please note that Shared Services term which is widely used in most of the documents is nothing but about Inter VRF/Tenant Communication.

Define the subnets at BD level for both the EPGs in both VRF and mark them as Shared between VRFs.

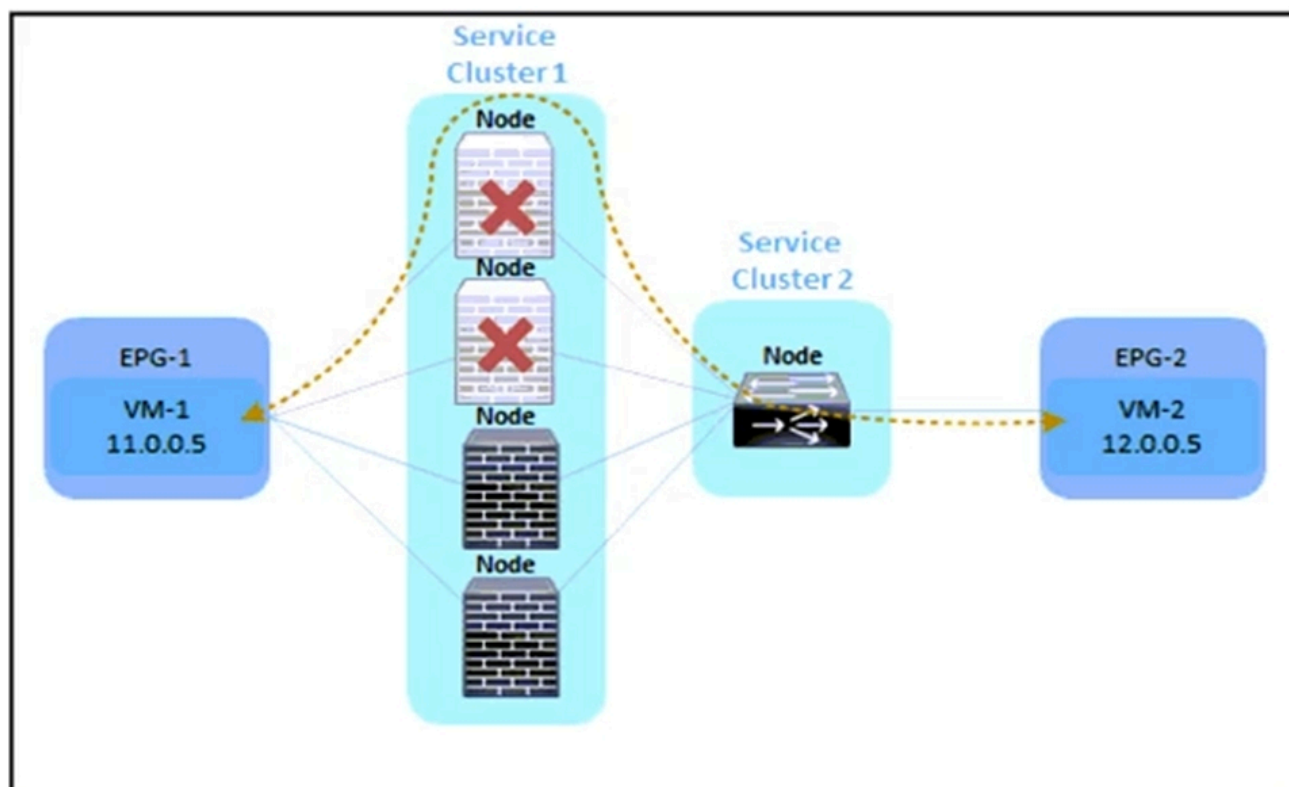
upvoted 1 times

  **Narbleddeath** 8 months, 2 weeks ago

Selected Answer: A

<https://community.cisco.com/t5/data-center-and-cloud-blogs/cisco-aci-inter-vrf-tenant-route-leaking-design-simplified/ba-p/3820919>

upvoted 1 times



Refer to the exhibit. An engineer must divert the traffic between VM-1 and VM-2 by using a Multi-Node service graph. The solution should prevent an insufficient number of available Layer 4 to Layer 7 devices in the first cluster. Which configuration set accomplishes this goal?

- A. • PBR node tracking
 - tracking threshold with action bypass
 - symmetric PBR
 - resilient hashing
- B. • PBR node tracking
 - tracking threshold with action permit
 - unidirectional PBR
 - resilient hashing
- C. • PBR node tracking
 - tracking threshold with action permit
 - symmetric PBR
 - resilient hashing
- D. • PBR node tracking
 - tracking threshold with action deny
 - symmetric PBR
 - unidirectional PBR

Correct Answer: A

Community vote distribution

A (80%)

C (20%)

👤 **zelya19** 9 months, 4 weeks ago

Selected Answer: A

Table 17: Down action - Bypass, Behavior - Traffic is redirected to next PBR node in the service graph, Use case - Skip over optional service node in Multi nodes service graph.

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739971.html#Optionalfeatures>

upvoted 1 times

🗨️ 👤 **kamel86** 1 year, 2 months ago

Selected Answer: A

Bypass: Traffic is redirected to next PBR node in the service graph.

permit action: Traffic is sent directly to the destination, and the rest of the service chain is skipped.

upvoted 2 times

🗨️ 👤 **imamus** 1 year, 4 months ago

Selected Answer: A

Answer is A, as bypass action is necessary to skip the nodes with lower threshold or the ones that are not up / unavailable.

upvoted 1 times

🗨️ 👤 **cooljit** 1 year, 8 months ago

Option C is correct:

- PBR node tracking: This allows PBR to track the status of the Layer 4 to Layer 7 devices in the service graph.
- Tracking threshold with action permit: This sets the threshold for device availability and permits traffic redirection to the available devices.
- Symmetric PBR: This ensures that both inbound and outbound traffic is redirected through the same set of Layer 4 to Layer 7 devices in the service graph.
- Resilient hashing: This load-balancing algorithm ensures that traffic is distributed evenly across all available devices, preventing any one device from becoming overloaded.

upvoted 2 times

🗨️ 👤 **frzzt** 1 year, 9 months ago

I think A, there is action bypass

https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/layer-4-to-layer-7-services-configuration/cisco-apic-layer-4-to-layer-7-services-deployment-guide-52x/configuring-policy-based-redirect-52x.html#id_109876

upvoted 4 times

🗨️ 👤 **VY01** 1 year, 9 months ago

Selected Answer: C

I thin C Ref.<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739971.html>

upvoted 1 times