

使用ACI多站點結構配置站點間L3out

目錄

[簡介](#)

[必要條件](#)

[需求](#)

[採用元件](#)

[背景資訊](#)

[站點間L3out配置支援的架構](#)

[設定](#)

[網路圖](#)

[物理拓撲](#)

[邏輯拓撲](#)

[組態](#)

[配置Schema-config1](#)

[配置交換矩陣策略](#)

[設定RTEP/ETEP](#)

[配置延伸租戶](#)

[配置方案](#)

[建立架構](#)

[建立Site-A模板](#)

[配置模板](#)

[建立延伸模板](#)

[附加模板](#)

[配置靜態埠繫結](#)

[配置BD](#)

[配置主機A\(N9K\)](#)

[建立Site-B模板](#)

[配置站點B L3out](#)

[建立外部EPG](#)

[配置外部N9K\(Site-B\)](#)

[將站點B L3out連線到站點A EPG\(BD\)](#)

[設定合約](#)

[建立合約](#)

[驗證](#)

[終端學習](#)

[ETEP/RTEP驗證](#)

[ICMP可達性](#)

[路由驗證](#)

[疑難排解](#)

[Site2 Leaf1](#)

[站點2 骨幹](#)

[站點1 骨幹](#)

[瞭解路由識別符號條目](#)

[站點1 葉1](#)

[驗證ELAM\(Site1 Spine\)](#)

[站點1 骨幹驗證路由對映](#)

簡介

本檔案介紹使用思科以應用為中心的基礎設施(ACI)多站點交換矩陣的站點間L3out配置步驟。

必要條件

需求

思科建議您瞭解以下主題：

- 功能性ACI多站點交換矩陣設定
- 外部路由器/連線

採用元件

本檔案中的資訊是根據：

- 多站點協調器(MSO)版本2.2(1)或更高版本
- ACI版本4.2(1)或更高版本
- MSO節點
- ACI交換矩陣
- Nexus 9000系列交換機(N9K) (終端主機和L3out外部裝置模擬)
- Nexus 9000系列交換器(N9K)(站點間網路(ISN))

本文中的資訊是根據特定實驗室環境內的裝置所建立。文中使用到的所有裝置皆從已清除 (預設) 的組態來啟動。如果您的網路運作中，請確保您瞭解任何指令可能造成的影響。

背景資訊

站點間L3out配置支援的架構

Schema-config1

- 租戶在站點 (A和B) 之間延伸。
- 虛擬路由和轉送(VRF)在站點 (A和B) 之間延伸。
- 一個站點(A)的本地終端組(EPG)/網橋域(BD)。
- L3out本地到另一個站點(B)。
- L3out的外部EPG本地到站點(B)。
- 從MSO完成合約建立和配置。

Schema-config2

- 租戶在站點 (A和B) 之間延伸。

- VRF在站點 (A和B) 之間延伸。
- EPG/BD在站點 (A和B) 之間延伸。
- L3out本地到一個站點(B)。
- L3out的外部EPG本地到站點(B)。
- 合約配置可以從MSO完成，或者每個站點都從應用策略基礎設施控制器(APIC)建立本地合約，並在延伸的EPG和L3out外部EPG之間本地連線。在這種情況下，由於本地合約關係和策略實施需要影子External_EPG，因此影子External_EPG出現在站點A。

Schema-config3

- 租戶在站點 (A和B) 之間延伸。
- VRF在站點 (A和B) 之間延伸。
- EPG/BD在站點 (A和B) 之間延伸。
- L3out本地到一個站點(B)。
- L3out的外部EPG在站點 (A和B) 之間延伸。
- 合約配置可以從MSO完成，或者每個站點都從APIC建立本地合約，並在延伸型EPG和延伸型外部EPG之間本地連線。

Schema-config4

- 租戶在站點 (A和B) 之間延伸。
- VRF在站點 (A和B) 之間延伸。
- 一個站點(A)的本地EPG/BD或每個站點的EPG/BD (站點A中的EPG-A和站點B中的EPG-B)。
- L3out本地到一個站點(B)，或者為了獲得外部連線的冗餘，您可以將L3out本地到每個站點 (本地到站點A，本地到站點B)。
- L3out的外部EPG在站點 (A和B) 之間延伸。
- 合約配置可以從MSO完成，或者每個站點都從APIC建立本地合約，並在延伸的EPG和延伸的外部EPG之間本地連線。

Schema-config5 (傳輸路由)

- 租戶在站點 (A和B) 之間延伸。
- VRF在站點 (A和B) 之間延伸。
- L3out local to each site(local to site A and local to site B)。
- 每個站點 (A和B) 的本地外部EPG。
- 合約配置可以從MSO完成，或者每個站點都從APIC建立本地合約，並在外部EPG本地和卷影外部EPG本地之間本地連線。

Schema-config5 (InterVRF傳輸路由)

- 租戶在站點 (A和B) 之間延伸。
- 每個站點 (A和B) 本地的VRF。
- L3out local to each site(local to site A and local to site B)。
- 每個站點 (A和B) 的本地外部EPG。
- 合約配置可以從MSO完成，或者每個站點都從APIC建立本地合約，並在外部EPG本地和卷影外部EPG本地之間本地連線。

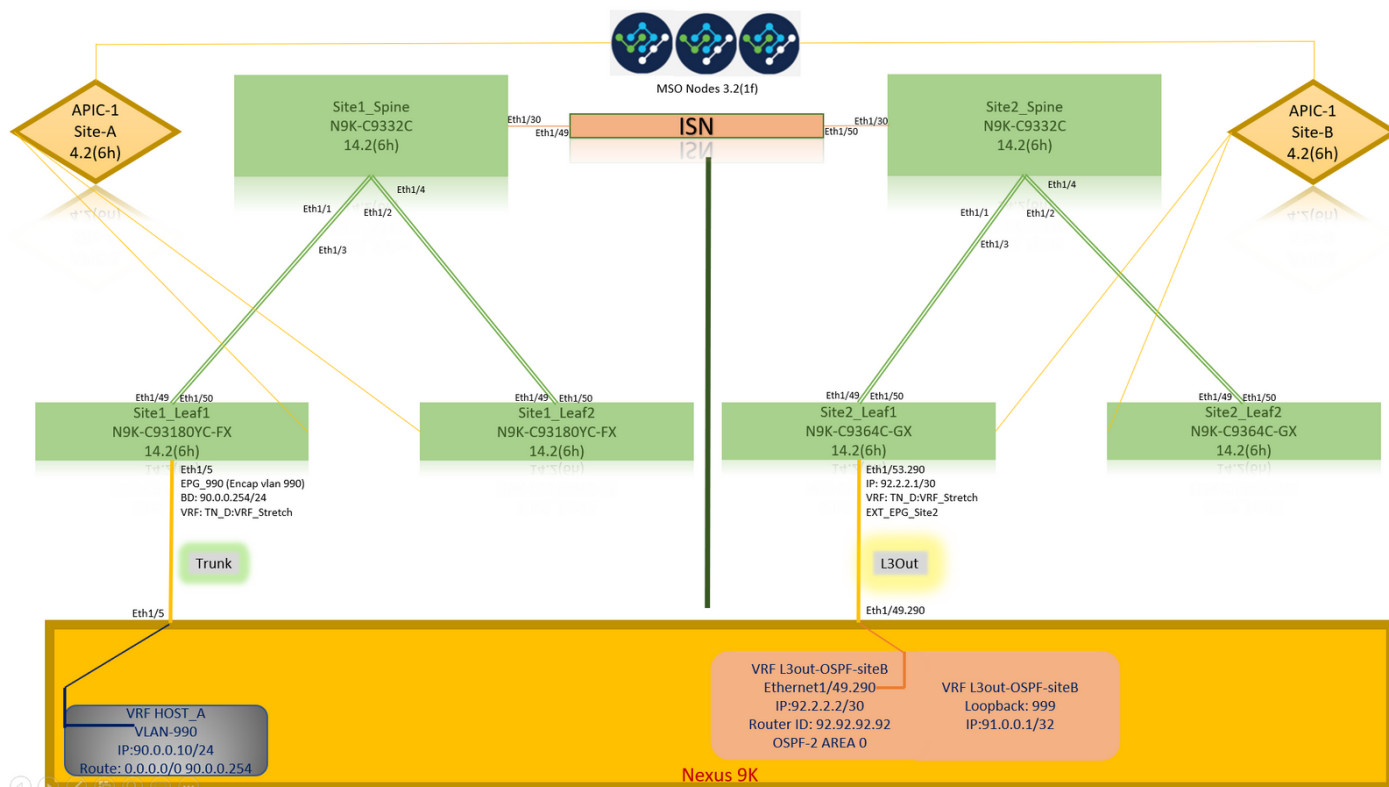
注意：本文檔提供基本的站點間L3out配置步驟和驗證。在本示例中，使用了Schema-

config1.

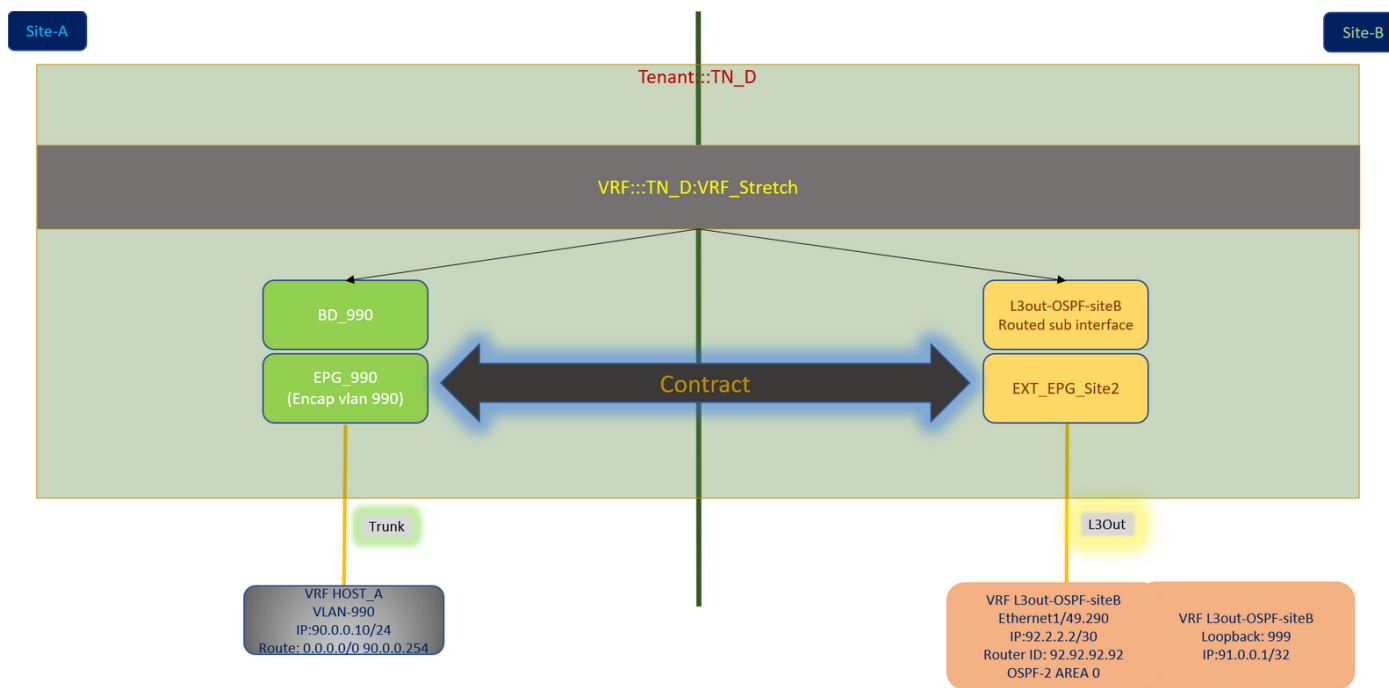
設定

網路圖

物理拓撲



邏輯拓撲



組態

在本例中，我們使用Schema-config1。但是，對於其他支援的架構配置，此配置可以類似方式完成（根據合約關係進行細微更改），不同之處在於，延伸對象需要在延伸模板中而不是特定站點模板中。

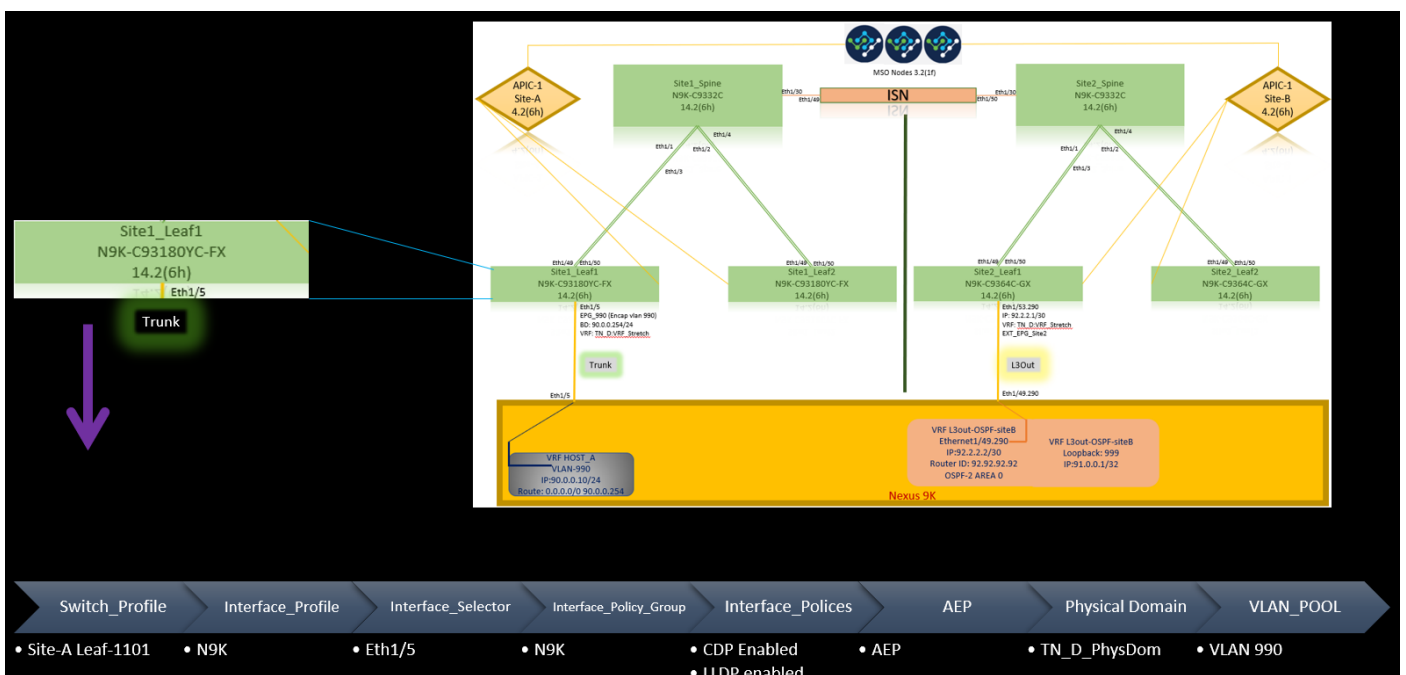
配置Schema-config1

- 租戶在站點（A和B）之間延伸。
- VRF在站點（A和B）之間延伸。
- 一個站點本地的EPG/BD(A)。
- L3out本地到另一個站點(B)。
- L3out的外部EPG本地到站點(B)。
- 從MSO完成合約建立和配置。
檢視站點間L3Out准則和限制。
- 不支援站點間L3out的配置：站點中的組播接收器，通過另一個站點L3out從外部源接收組播。站點中從外部源接收的組播不會傳送到其他站點。當站點中的接收器收到來自外部源的組播時，必須在本地L3out上接收該組播。內部多點傳送來源使用PIM-SM任意來源多點傳送(ASM)將多點傳送傳送到外部接收器。內部多點傳送來源必須能夠從本地L3out到達外部集結點(RP)。Giant OverLay Fabric（高爾夫）。外部EPG的首選組。

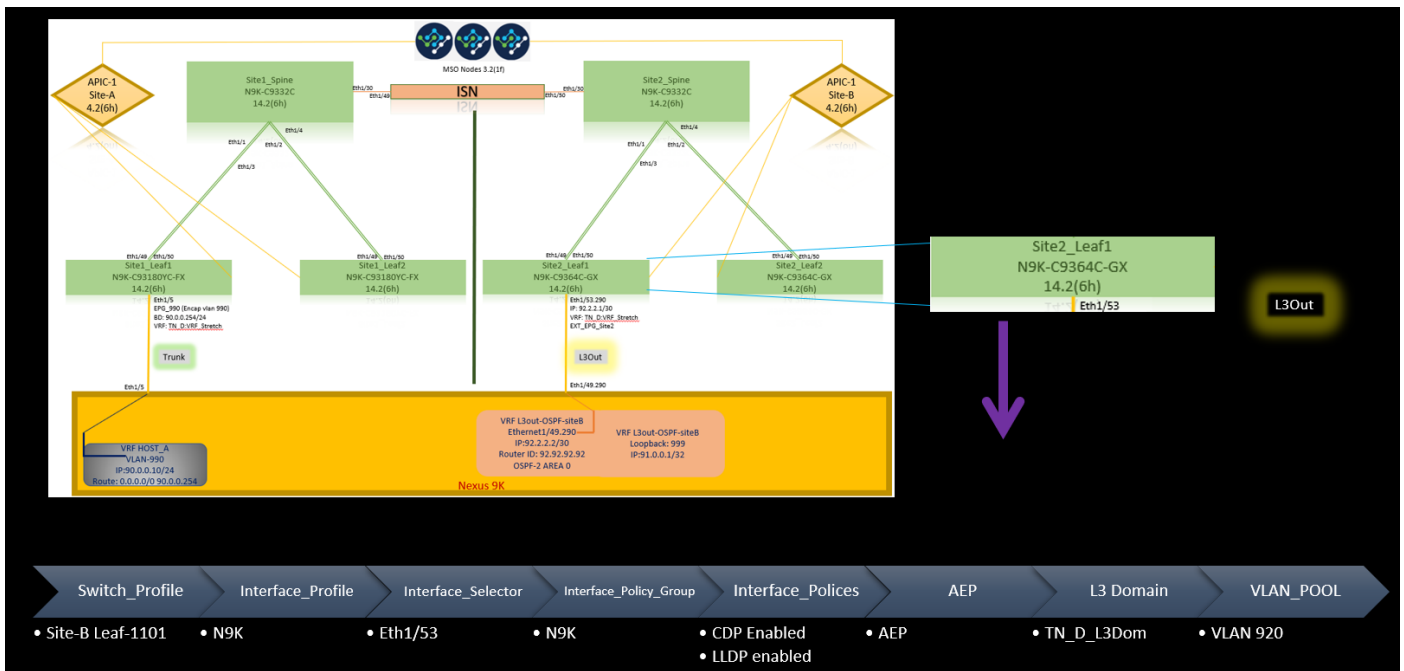
配置交換矩陣策略

每個站點上的交換矩陣策略是基本配置，因為這些策略配置連結到特定租戶/EPG/靜態埠繫結或L3out物理連線。交換矩陣策略的任何配置錯誤都可能導致來自APIC或MSO的邏輯配置失敗，因此提供的交換矩陣策略配置在實驗室設定中使用。它有助於瞭解哪個對象連結到MSO或APIC中的哪個對象。

站點A上的主機A連線結構策略



站點B的L3out連線結構策略



可選步驟

一旦為各個連線設定了交換矩陣策略，就可以確保從各自的APIC集群發現並訪問所有枝葉/主幹。接下來，您可以驗證兩個站點（APIC集群）均可從MSO訪問，並且多站點設定可以正常運行（並且IPN連線）。

設定RTEP/ETEP

可路由通道端點池(RTEP)或外部通道端點池(ETEP)是站點間L3out所需的配置。較舊版本的MSO顯示「可路由TEP池」，而較新版本的MSO顯示「外部TEP池」，但兩者都是同義詞。這些TEP池通過VRF「Overlay-1」用於邊界網關協定(BGP)乙太網VPN(EVPN)。

從L3out的外部路由通過BGP EVPN通告到另一個站點。此RTEP/ETEP還用於遠端枝葉配置，因此，如果您的ETEP/RTEP配置已經存在於APIC中，則必須在MSO中匯入它。

以下是從MSO GUI配置ETEP的步驟。由於版本是3.X MSO，因此它會顯示ETEP。ETEP池在每個站點必須是唯一的，並且不能與每個站點的任何內部EPG/BD子網重疊。

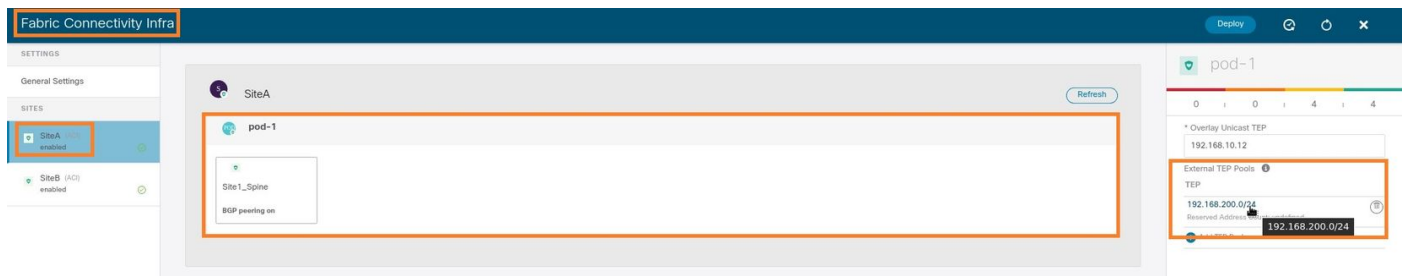
站點A

步驟1. 在MSO GUI頁面（在網頁中開啟多站點控制器）中，選擇Infrastructure > Infra Configuration。按一下「Configure Infra」。




步驟2. 在Configure Infra內，選擇Site-A，在Site-A內，選擇pod-1。然後在Pod-1內，使用站點A的外部TEP IP地址配置外部TEP池。（在本例中為192.168.200.0/24）。如果您在站點A中有多個

POD，請對其他POD重複此步驟。



步驟3.若要驗證APIC GUI中ETEP池的配置，請選擇Fabric > Inventory > Pod Fabric Setup Policy > Pod-ID (按兩下以開啟[Fabric Setup Policy a POD-Pod-x]) > External TEP。



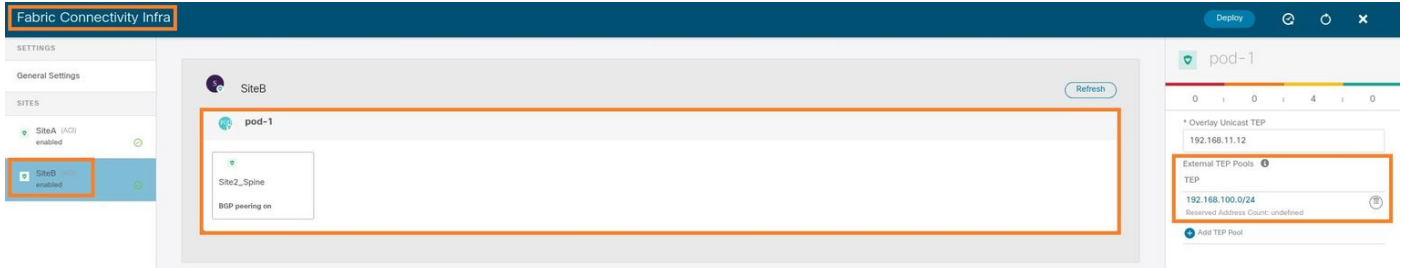
您也可以使用以下命令驗證設定：

```
moquery -c fabricExtRoutablePodSubnet
moquery -c fabricExtRoutablePodSubnet -f 'fabric.ExtRoutablePodSubnet.pool=="192.168.200.0/24"'
```

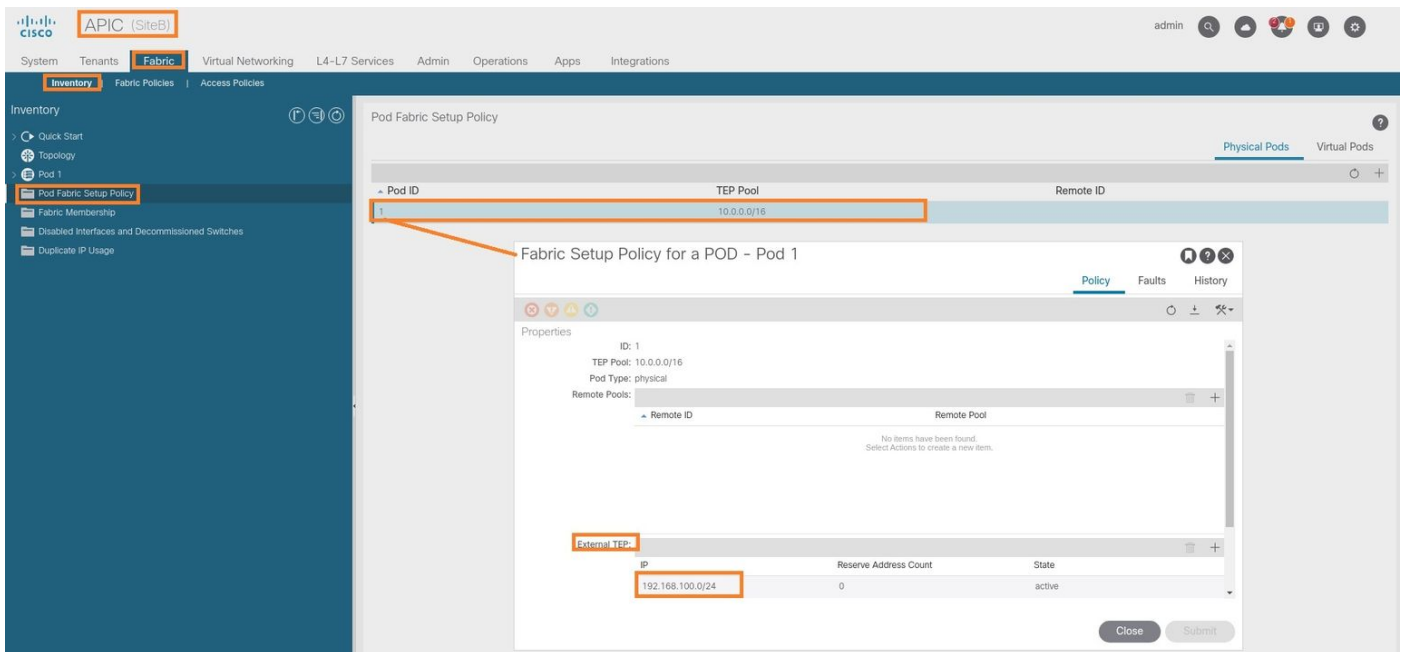
```
APIC1# moquery -c fabricExtRoutablePodSubnet
Total Objects shown: 1
# fabric.ExtRoutablePodSubnet
pool                : 192.168.200.0/24
annotation         : orchestrator:misc
childAction        :
descr              :
dn                 : uni/controller/setuppod/setupp-1/extrtpodsubnet-[192.168.200.0/24]
extMngdBy          :
lcOwn               : local
modTs              : 2021-07-19T14:45:22.387+00:00
name               :
nameAlias          :
reserveAddressCount : 0
rn                 : extrtpodsubnet-[192.168.200.0/24]
state              : active
status             :
uid                : 0
```

站點B

步驟1.為站點B配置外部TEP池（與站點A相同的步驟。）在MSO GUI頁面（在網頁中開啟多站點控制器）中，選擇**Infrastructure > Infra Configuration**。按一下**Configure Infra**。在Configure Infra內，選擇**Site-B**。在站點B內部，選擇**pod-1**。然後在pod-1內部，使用站點B的外部TEP IP地址配置**外部TEP池**。（在本例中為192.168.100.0/24）。如果您在站點B中有多個POD，請對其他POD重複此步驟。



步驟2.若要驗證APIC GUI中ETEP池的配置，請選擇**Fabric > Inventory > Pod Fabric Setup Policy > Pod-ID**(按兩下以開啟[Fabric Setup Policy a POD-Pod-x]) > **External TEP**。




對於Site-B APIC，輸入以下命令以驗證ETEP地址池。

```
apic1# moquery -c fabricExtRoutablePodSubnet -f
'fabric.ExtRoutablePodSubnet.pool=="192.168.100.0/24"'
Total Objects shown: 1
# fabric.ExtRoutablePodSubnet
pool                : 192.168.100.0/24
annotation          : orchestrator:mso <<< This means, configuration pushed from MSO.
childAction         :
descr               :
dn                  : uni/controller/setuppod/setupp-1/extrtpodsubnet-[192.168.100.0/24]
extMngdBy           :
lcOwn               : local
modTs               : 2021-07-19T14:34:18.838+00:00
name                :
nameAlias           :
reserveAddressCount : 0
rn                  : extrtpodsubnet-[192.168.100.0/24]
```

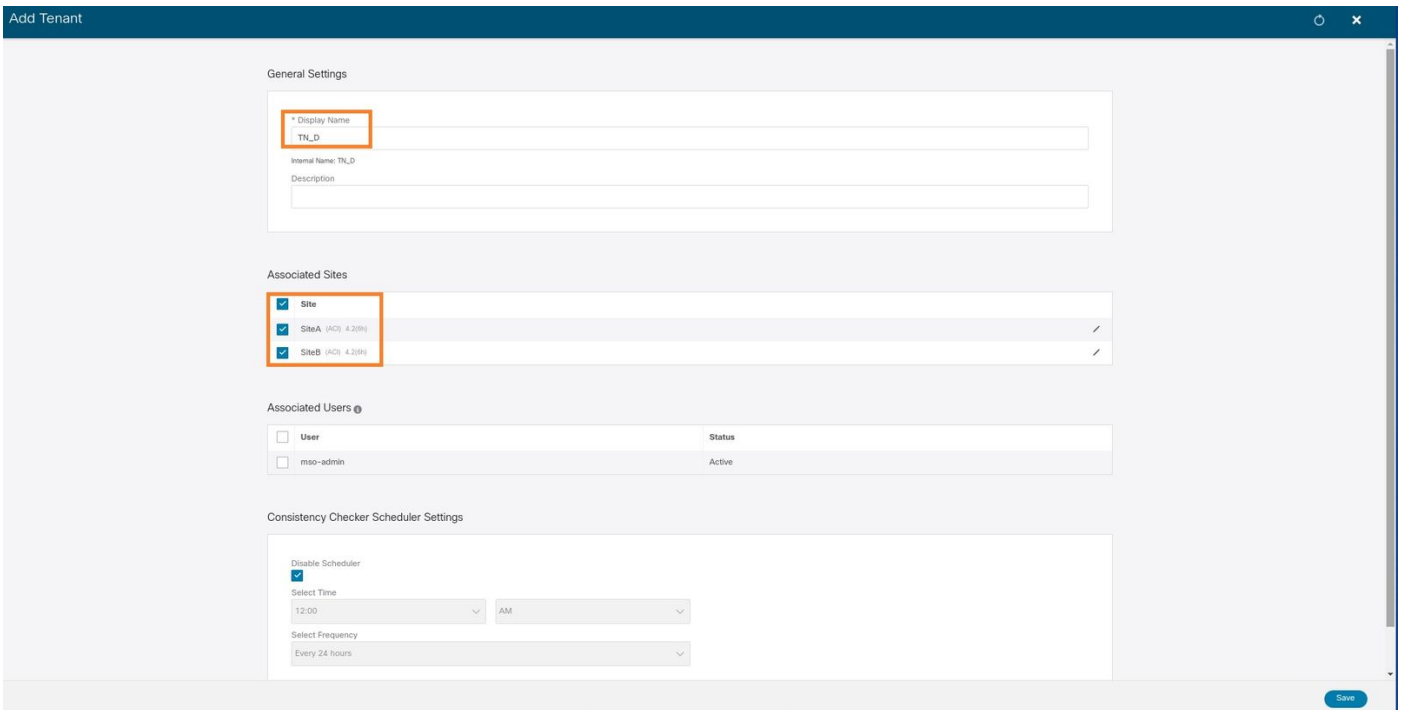
```
state           : active
status          :
uid             : 0
```

配置延伸租戶

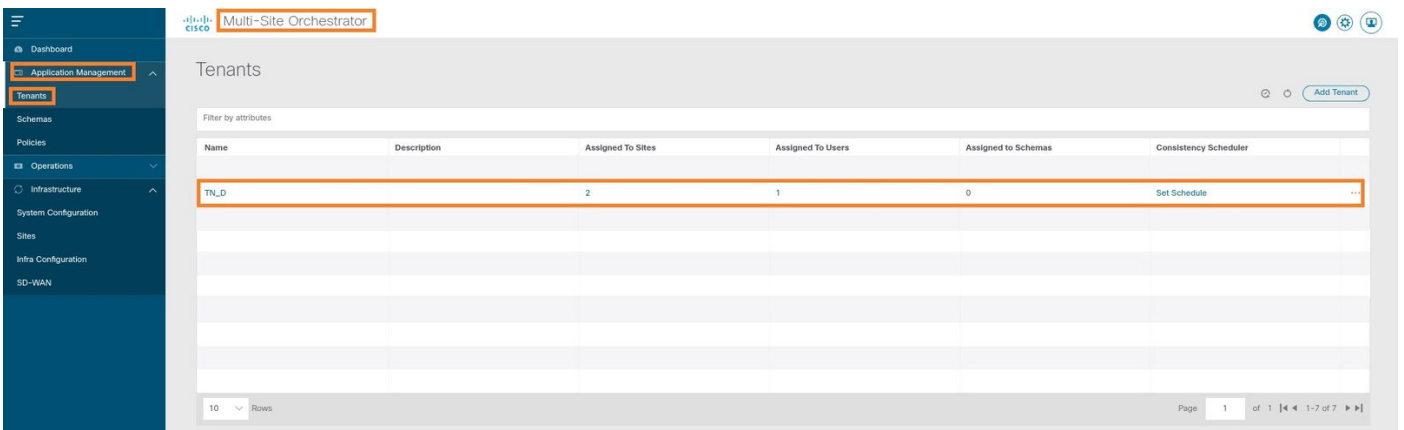
步驟1. 在MSO GUI中，選擇**Application Management > Tenants**。 按一下**Add Tenant**。 在本示例中，租戶名稱為「TN_D」。



步驟2. 在**顯示名稱**欄位中輸入租戶的名稱。在**Associated Sites**部分中，選中**Site A**和**Site B**覈取方塊。



步驟3. 驗證是否已建立新租戶「Tn_D」。



邏輯檢視

當我們從MSO建立租戶時，它基本上在站點A和站點B建立租戶。這是一個延伸租戶。此租戶的邏輯檢視如本示例所示。此邏輯檢視有助於理解租戶TN_D是站點A和站點B之間的延伸租戶。

The screenshot displays the Cisco Multi-Site Orchestrator interface. The left sidebar contains navigation options: Dashboard, Application Management, Tenants, Schemas, Policies, Operations, Infrastructure, and System Configuration. The main content area is titled 'Tenants' and includes a filter by attributes field and an 'Add Tenant' button. A table lists tenants with the following columns: Name, Description, Assigned To Sites, Assigned To Users, Assigned to Schemas, and Consistency Scheduler. The table contains one entry, 'TN_D', which is highlighted with an orange border. This entry shows 2 sites, 1 user, and 0 schemas assigned to it. Below the screenshot is a large light green rectangular area with the text 'Tenant:::TN_D' in the top left corner.

Name	Description	Assigned To Sites	Assigned To Users	Assigned to Schemas	Consistency Scheduler
TN_D		2	1	0	Set Schedule

您可以在每個站點的APIC中驗證邏輯檢視。您可以看到站點A和站點B都顯示已建立「TN_D」租戶。



This has been created from

TN_D

> Quick Start

∨ TN_D

> Application Profiles

> Networking

> Contracts

> Policies

> Services

在Site-B中還建立了相同的延伸租戶「TN_D」。



System

Tenants

Fabric

ALL TENANTS

Add Tenant

T



This has been created from

TN_D

> Quick Start

▼ TN_D

> Application Profiles

> Networking

> Contracts

> Policies

> Services

此命令顯示從MSO推送的租戶，您可以將其用於驗證目的。您可以在兩個站點的APIC中運行此命令。

```
APIC1# moquery -c fvTenant -f 'fv.Tenant.name=="TN_D"'
```

```
Total Objects shown: 1
```

```
# fv.Tenant
```

```
name          : TN_D
annotation    : orchestrator:misc
childAction   :
descr         :
dn            : uni/tn-TN_D
extMngdBy     : msc
lcOwn         : local
modTs         : 2021-09-17T21:42:52.218+00:00
monPolDn      : uni/tn-common/monepg-default
nameAlias     :
ownerKey      :
ownerTag      :
rn            : tn-TN_D
status        :
uid           : 0
```

```
apic1# moquery -c fvTenant -f 'fv.Tenant.name=="TN_D"'
```

```
Total Objects shown: 1
```

```
# fv.Tenant
```

```
name          : TN_D
```



```
annotation      : orchestrator:msc
childAction     :
descr           :
dn              : uni/tn-TN_D
extMngdBy      : msc
lcOwn          : local
modTs          : 2021-09-17T21:43:04.195+00:00
monPolDn       : uni/tn-common/monepg-default
nameAlias      :
ownerKey       :
ownerTag       :
rn             : tn-TN_D
status         :
uid            : 0
```

配置方案

接下來，建立總共包含三個模板的架構：

1. 站點A的模板：站點A的模板僅與站點A關聯，因此該模板中的任何邏輯對象配置都只能推送到站點A的APIC。
2. 站點B的模板：站點B的模板僅與站點B關聯，因此該模板中的任何邏輯對象配置都只能推送到站點B的APIC。
3. 延伸模板：拉伸模板與兩個站點關聯，拉伸模板中的任何邏輯配置都可以推送到APIC的兩個站點。

建立架構

模式在MSO中是本地有效的，它不會在APIC中建立任何對象。架構配置是每個配置的邏輯分離。您可以為同一租戶使用多個架構，並且每個架構內還可以有多個模板。

例如，可以為租戶X的資料庫伺服器提供一個架構，而應用程式伺服器為同一個tenant-X使用不同的架構。這有助於區分每個特定的應用程式相關配置，並且在您需要調試問題時非常簡單。查詢資訊也很容易。

使用租戶名稱建立架構（例如TN_D_Schema）。但是，不需要將架構名稱以租戶名稱開頭，您可以建立具有任意名稱的架構。

步驟1.選擇Application Management > Schemas。按一下Add Schema。



步驟2.在名稱欄位中輸入方案的名稱。在本示例中，它是「TN_D_Schema」，但您可以保留任何適合您環境的名稱。按一下「Add」。

General
✕

* Name

TN_D_Schema

Description

Schema for Tenant TN_D

Add

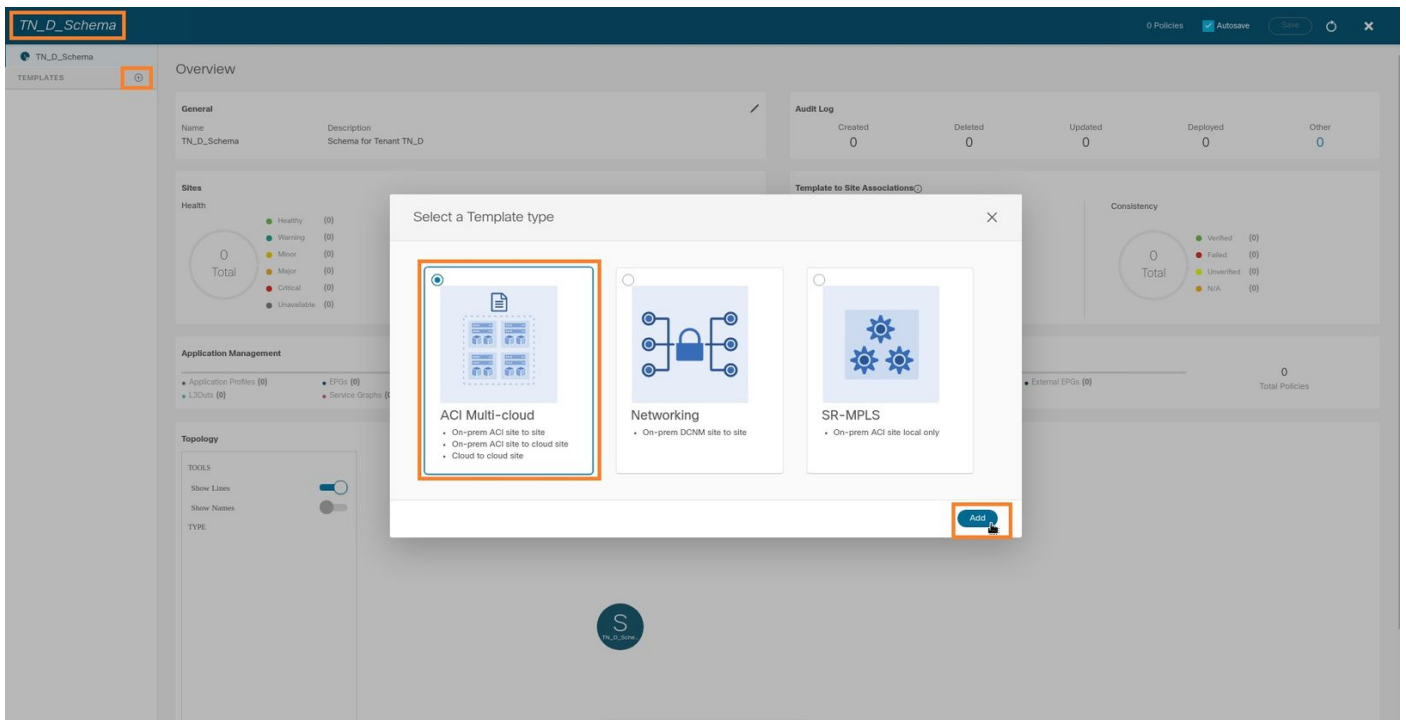
步驟3.驗證架構「TN_D_Schema」是否已建立。

The screenshot displays the 'Overview' page for the 'TN_D_Schema' template. The 'General' section is highlighted with orange boxes around the 'Name' (TN_D_Schema) and 'Description' (Schema for Tenant TN_D) fields. The 'Audit Log' section shows a table with columns for Created, Deleted, Updated, Deployed, and Other, all with a count of 0. The 'Sites' section features two donut charts: 'Health' (Total 0) and 'Type' (Total 0). The 'Application Management' section shows a 'Total Policies' count of 0. The 'Topology' section includes 'TOOLS' (Show Lines, Show Names) and 'TYPE' sections.

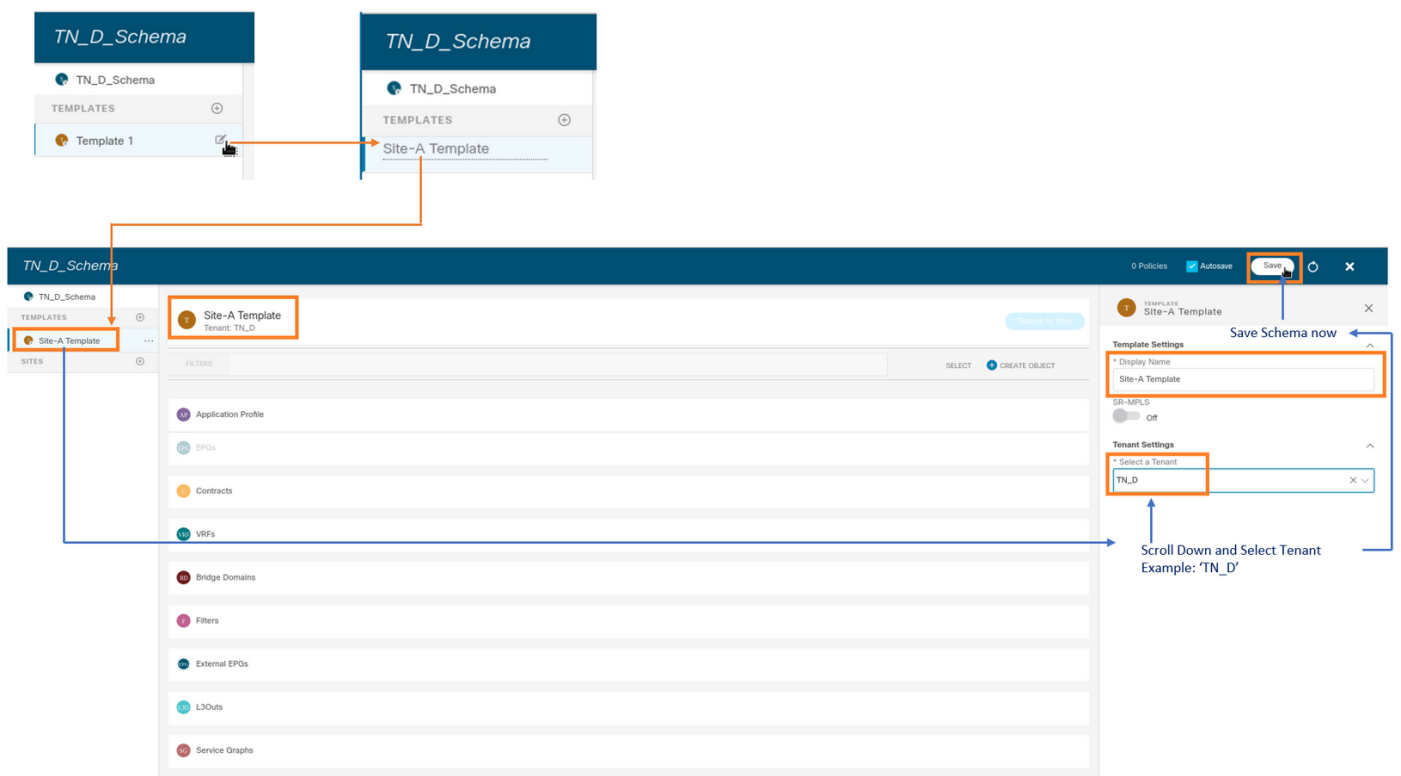
建立Site-A模板

步驟1.在架構中新增模板。

1. 要建立模板，請按一下已建立的架構下的**模板**。將顯示「選擇模板型別」對話方塊。
2. 選擇**ACI Multi-cloud**。
3. 按一下「**Add**」。



步驟2. 輸入模板的名稱。此模板特定於站點A，因此模板名稱為「站點A模板」。建立模板後，您可以將特定租戶附加到模板。在此示例中，將附加租戶「TN_D」。



配置模板

應用配置檔案配置

步驟1. 從建立的架構中選擇Site-A Template。按一下Add Application Profile。



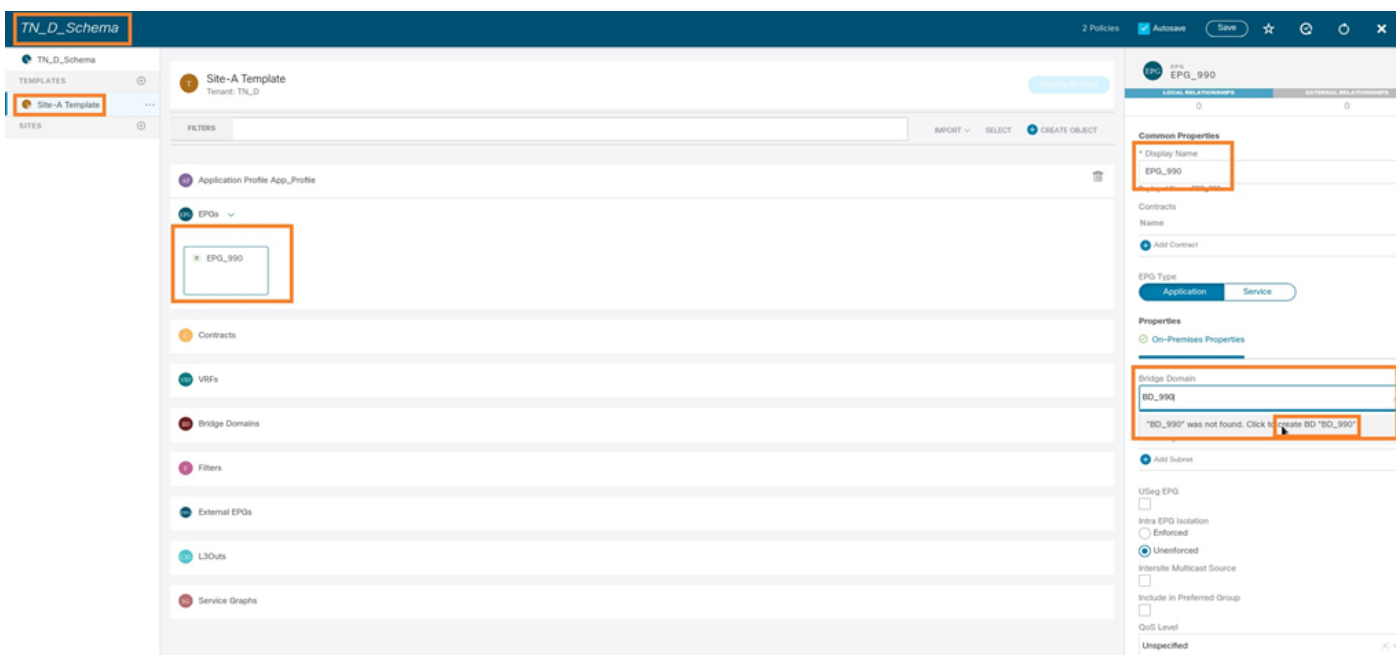
步驟2.在顯示名稱欄位中，輸入應用程式配置檔名稱App_Profile。



步驟3.下一步是建立EPG。要在應用配置檔案下新增EPG，請按一下Site-A模板下的Add EPG。您可以看到在EPG配置內建立了新的EPG。



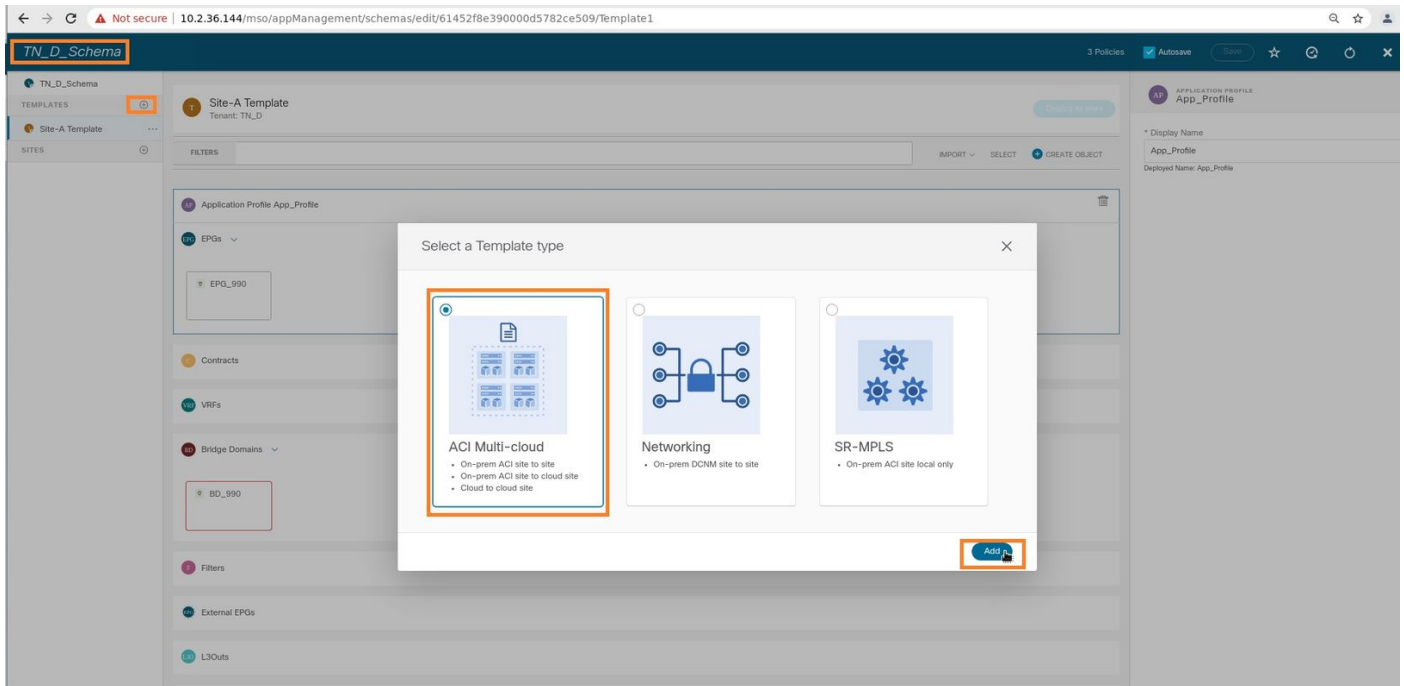
步驟4.要將EPG與BD和VRF連線，您必須將BD和VRF新增到EPG下。選擇Site-A Template。在Display Name欄位中，輸入EPG的名稱並附加一個新的BD（您可以建立一個新的BD或附加一個現有的BD）。



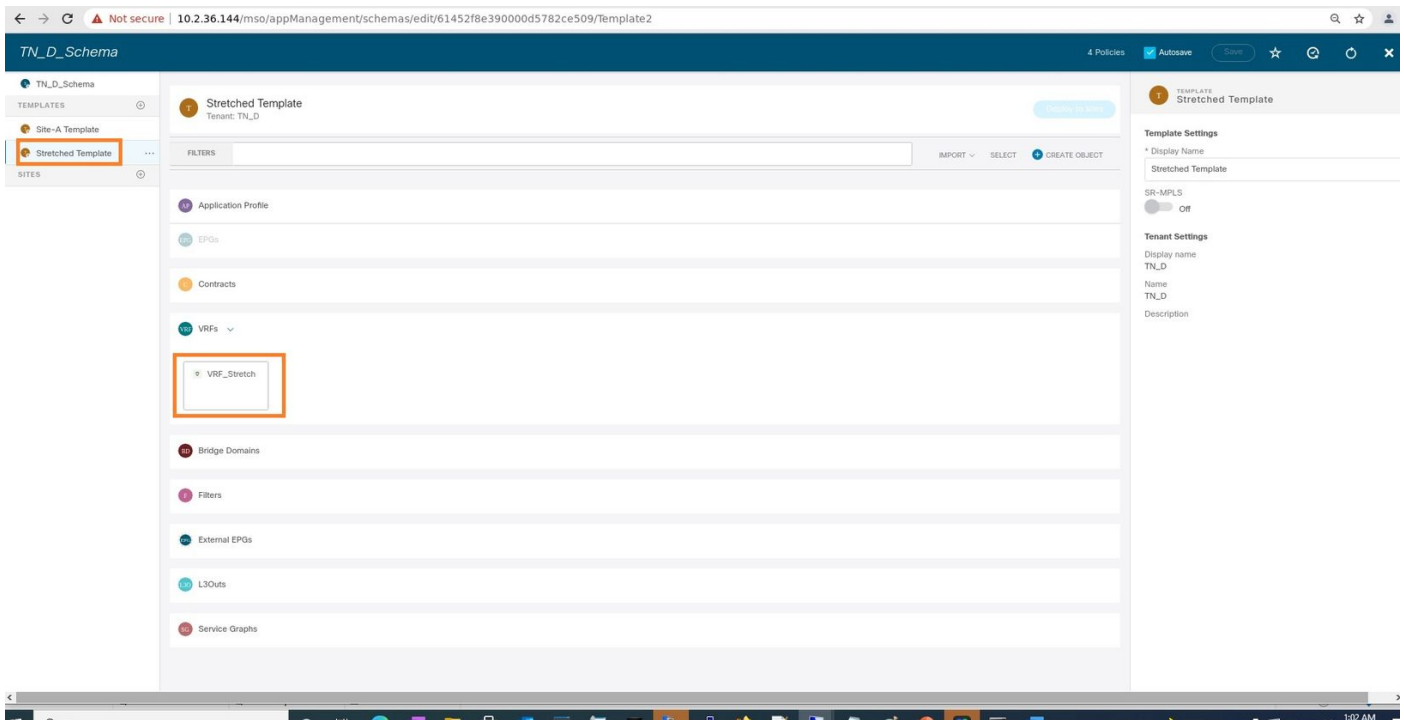
請注意，您必須將VRF連線到BD，但在此情況下會拉伸VRF。您可以使用延伸的VRF建立延伸模板，然後將該VRF附加到站點特定模板下的BD(在我們的情況下為Site-A Template)。

建立延伸模板

步驟1. 要建立拉伸模板，請在TN_D_Schema下按一下**Templates**。將顯示「選擇模板型別」對話方塊。選擇**ACI Multi-cloud**。按一下「Add」。輸入模板的**擴展模板**名稱。（可以輸入拉伸模板的任意名稱。）




步驟2. 選擇**拉伸模板**，然後建立名為**VRF_Stretch**的VRF。（您可以輸入VRF的任何名稱。）



BD是通過**Site-A Template**下的EPG建立建立的，但是沒有附加的VRF，因此您必須附加現在在**拉伸模板**中建立的**VRF**。

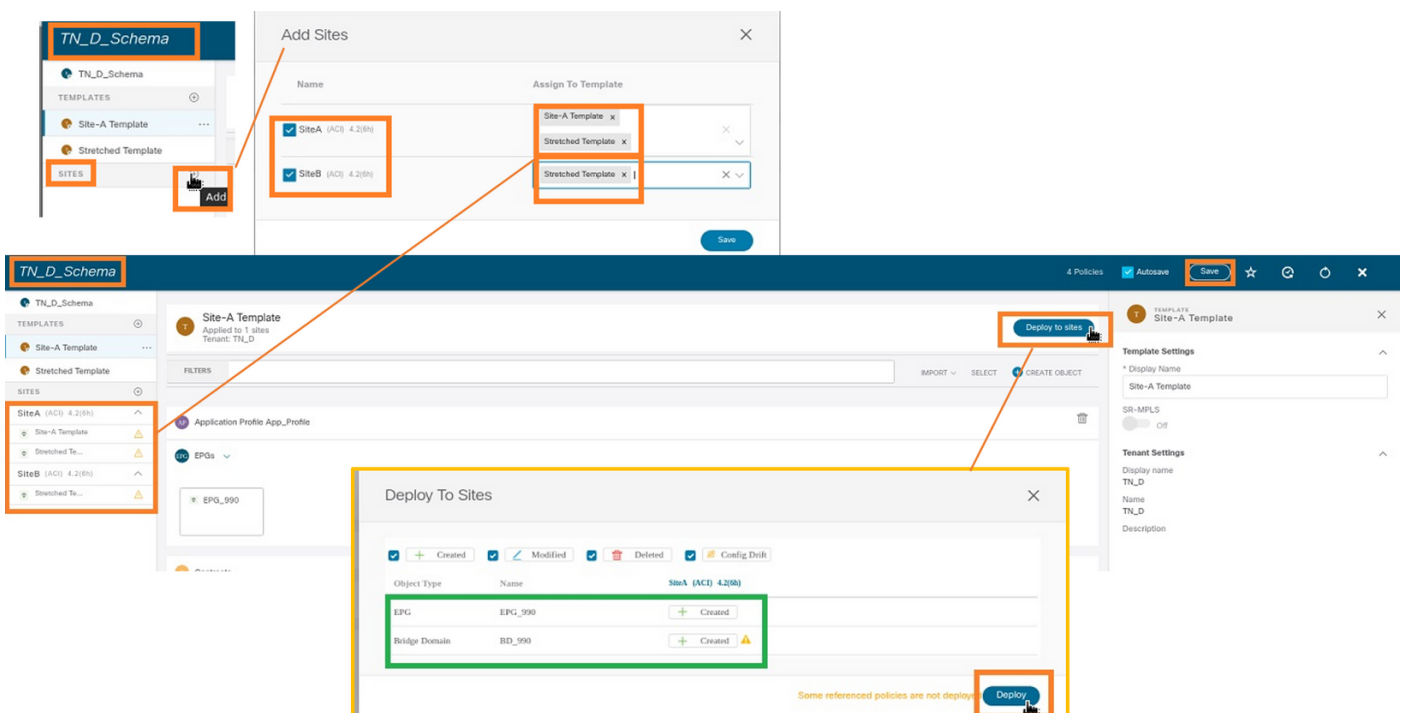
步驟3. 選擇**Site-A Template > BD_990**。在**Virtual Routing & Forwarding**下拉選單中，選擇**VRF_Stretch**。（您在本節的步驟2中建立的路徑。）





附加模板

下一步是附加僅包含Site-A的Site-A模板，並且需要將延伸模板附加到兩個站點。按一下Deploy to site (部署到架構內的站點)，將模板部署到相應的站點。

步驟1. 按一下TN_D_Schema > SITES下的+符號將站點新增到模板。在「Assign to Template」下拉選單中，為相應的站點選擇相應的模板。



步驟2. 您可以看到站點A已建立EPG和BD，但站點B沒有建立相同的EPG/BD，因為這些配置僅適用於MSO的站點A。但是，您可以看到VRF是在延伸模板中建立的，因此它是在兩個站點中建立的。



步驟3.使用這些命令驗證配置。

```

APIC1# moquery -c fvAEPg -f 'fv.AEPg.name=="EPG_990"'
Total Objects shown: 1
# fv.AEPg
name                : EPG_990
annotation          : orchestrator:msc
childAction         :
configIssues       :
configSt           : applied
descr              :
dn                 : uni/tn-TN_D/ap-App_Profile/epg-EPG_990
exceptionTag       :
extMngdBy          :
floodOnEncap       : disabled
fwdCtrl            :
hasMcastSource     : no
isAttrBasedEPg    : no
isSharedSrvMsiteEPg : no
lcOwn              : local
matchT             : AtleastOne
modTs              : 2021-09-18T08:26:49.906+00:00
monPolDn           : uni/tn-common/monepg-default
nameAlias          :
pcEnfPref          : unenforced
pctag              : 32770
prefGrMemb         : exclude
prio               : unspecified
rn                 : epg-EPG_990
scope              : 2850817
shutdown           : no
status             :
triggerSt          : triggerable
txId               : 1152921504609182523
uid                : 0

```

```

APIC1# moquery -c fvBD -f 'fv.BD.name=="BD_990"'
Total Objects shown: 1
# fv.BD
name : BD_990
OptimizeWanBandwidth : yes
annotation : orchestrator:misc
arpFlood : yes
bcastP : 225.0.56.224
childAction :
configIssues :
descr :
dn : uni/tn-TN_D/BD-BD_990
epClear : no
epMoveDetectMode :
extMngdBy :
hostBasedRouting : no
intersiteBumTrafficAllow : yes
intersiteL2Stretch : yes
ipLearning : yes
ipv6McastAllow : no
lcOwn : local
limitIpLearnToSubnets : yes
llAddr : ::
mac : 00:22:BD:F8:19:FF
mcastAllow : no
modTs : 2021-09-18T08:26:49.906+00:00
monPolDn : uni/tn-common/monepg-default
mtu : inherit
multiDstPktAct : bd-flood
nameAlias :
ownerKey :
ownerTag :
pcTag : 16387
rn : BD-BD_990
scope : 2850817
seg : 16580488
status :
type : regular
uid : 0
unicastRoute : yes
unkMacUcastAct : proxy
unkMcastAct : flood
v6unkMcastAct : flood
vmac : not-applicable
: 0

```

```

APIC1# moquery -c fvCtx -f 'fv.Ctx.name=="VRF_Stretch"'
Total Objects shown: 1
# fv.Ctx
name : VRF_Stretch
annotation : orchestrator:misc
bdEnforcedEnable : no
childAction :
descr :
dn : uni/tn-TN_D/ctx-VRF_Stretch
extMngdBy :
ipDataPlaneLearning : enabled
knwMcastAct : permit
lcOwn : local
modTs : 2021-09-18T08:26:58.185+00:00

```



```

monPolDn          : uni/tn-common/monepg-default
nameAlias         :
ownerKey          :
ownerTag          :
pcEnfDir          : ingress
pcEnfDirUpdated  : yes
pcEnfPref         : enforced
pcTag             : 16386
rn                : ctx-VRF_Stretch
scope             : 2850817
seg               : 2850817
status            :
uid               : 0

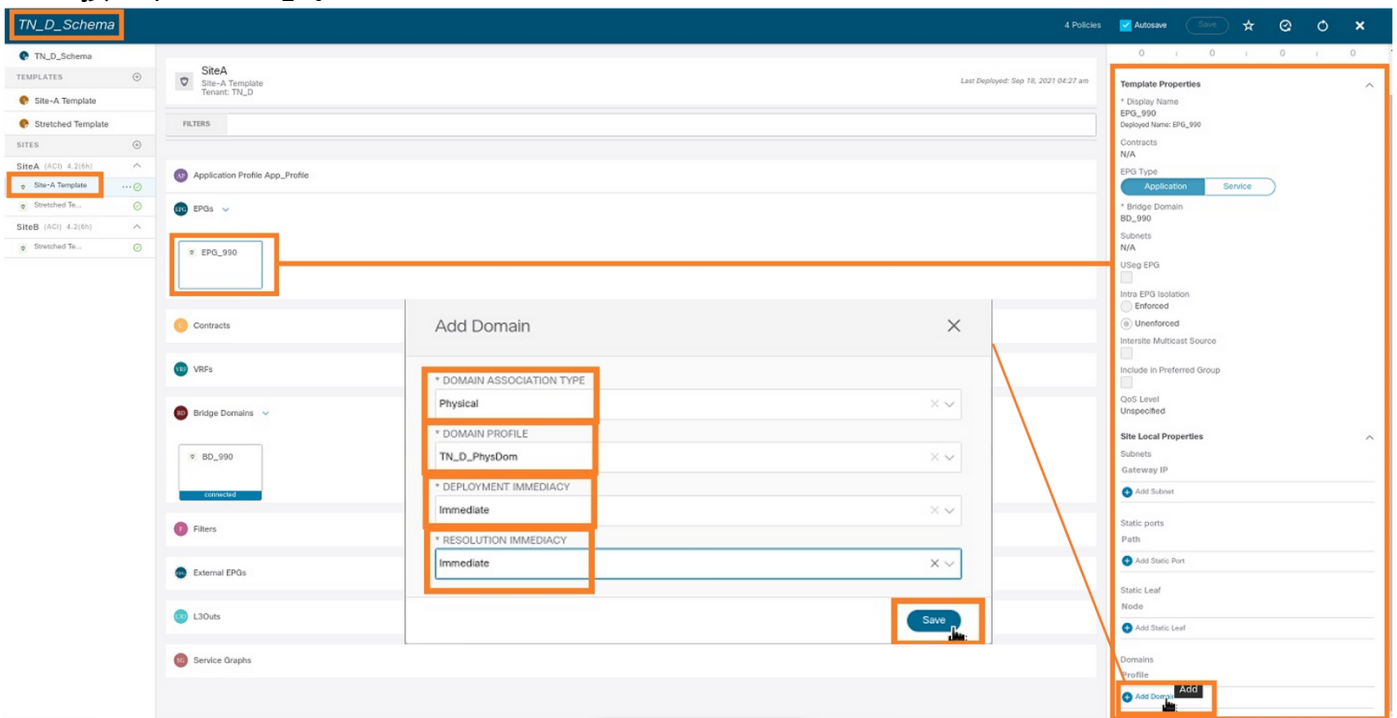
```

配置靜態埠繫結

您現在可以在EPG "EPG_990"下配置靜態埠繫結，還可以使用VRF HOST_A配置N9K (基本上是模擬HOST_A)。ACI端靜態埠繫結配置將首先完成。

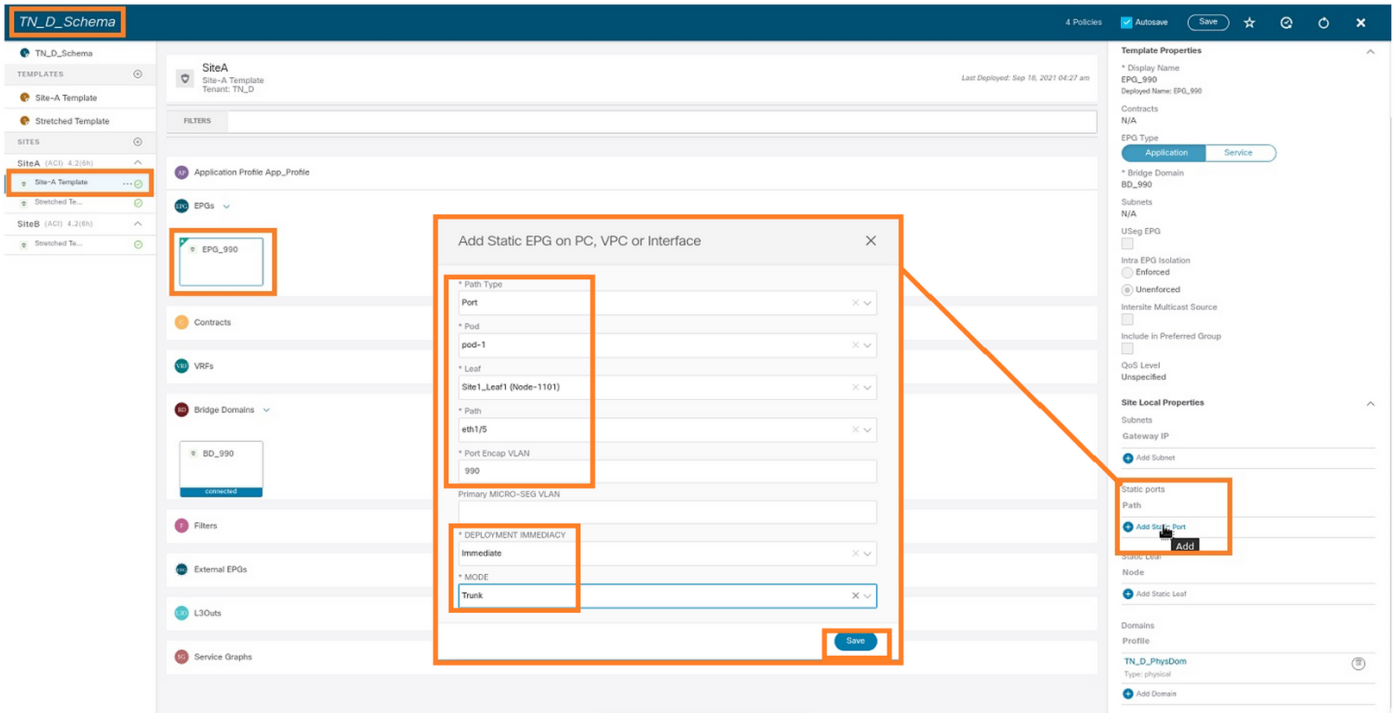
步驟1.在EPG_990下新增物理域。

1. 從建立的架構中選擇Site-A Template > EPG_990。
2. 在Template Properties框中，按一下Add Domain。
3. 在Add Domain對話方塊中，從下拉選單中選擇以下選項：域關聯型別 — 物理域配置檔案-TN_D_PhysDom部署即時性 — 即時解決方案即時 — 即時
4. 按一下「Save」。




步驟2.新增靜態埠(Site1_Leaf1 eth1/5)。

1. 從您建立的架構中選擇Site-A Template > EPG_990。
2. 在Template Properties框中，按一下Add Static Port。
3. 在Add Static EPG on PC, VPC or Interface對話方塊中，選擇Node-101 eth1/5 並分配VLAN 990。



步驟3.確保在EPG_990下新增靜態埠和物理域。



使用以下命令驗證靜態路徑繫結：

```

APIC1# moquery -c fvStPathAtt -f 'fv.StPathAtt.pathName=="eth1/5"' | grep EPG_990 -A 10 -B 5
# fv.StPathAtt
pathName      : eth1/5
childAction   :
descr         :
dn            : uni/epp/fv-[uni/tn-TN_D/ap-App_Profile/epg-EPG_990]/node-1101/stpathatt-[eth1/5]
lcOwn         : local
modTs         : 2021-09-19T06:16:46.226+00:00
monPolDn      : uni/tn-common/monepg-default
name          :
  
```

```

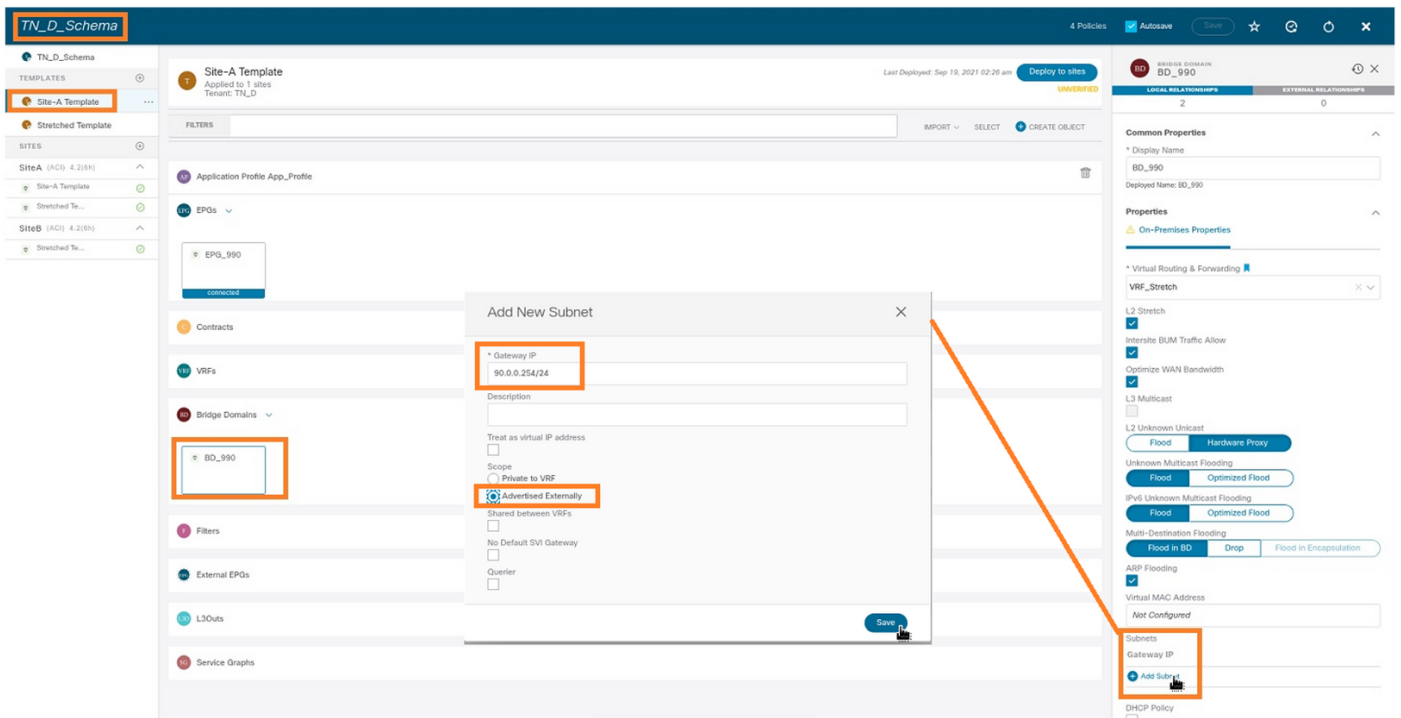
nameAlias      :
ownerKey       :
ownerTag       :
rn             : stpathatt-[eth1/5]
status         :

```

配置BD

步驟1.在BD下新增子網/IP (HOST_A使用BD IP作為網關)。

1. 從建立的架構中選擇Site-A Template > BD_990。
2. 按一下「Add Subnet」。
3. 在Add New Subnet對話方塊中，輸入Gateway IP address，然後按一下Advertised External單選按鈕。



步驟2.使用此命令驗證是否已將子網新增到APIC1站點A中。

```

APIC1# moquery -c fvSubnet -f 'fv.Subnet.ip=="90.0.0.254/24"'
Total Objects shown: 1

# fv.Subnet
ip           : 90.0.0.254/24
annotation   : orchestrator:misc
childAction  :
ctrl         : nd
descr        :
dn           : uni/tn-TN_D/BD-BD_990/subnet-[90.0.0.254/24]
extMngdBy   :
lcOwn        : local
modTs        : 2021-09-19T06:33:19.943+00:00
monPolDn     : uni/tn-common/monepg-default
name         :
nameAlias    :
preferred    : no
rn           : subnet-[90.0.0.254/24]
scope        : public

```

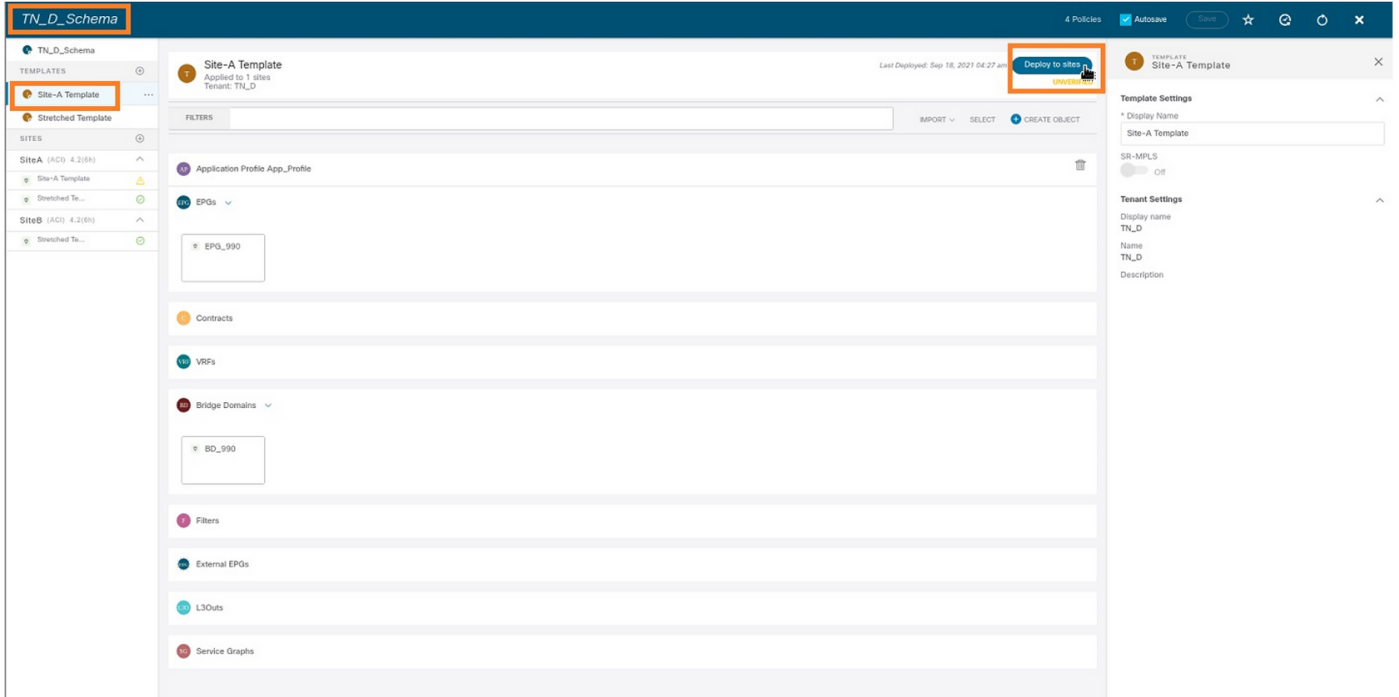
```

status      :
uid         : 0
virtual    : no

```

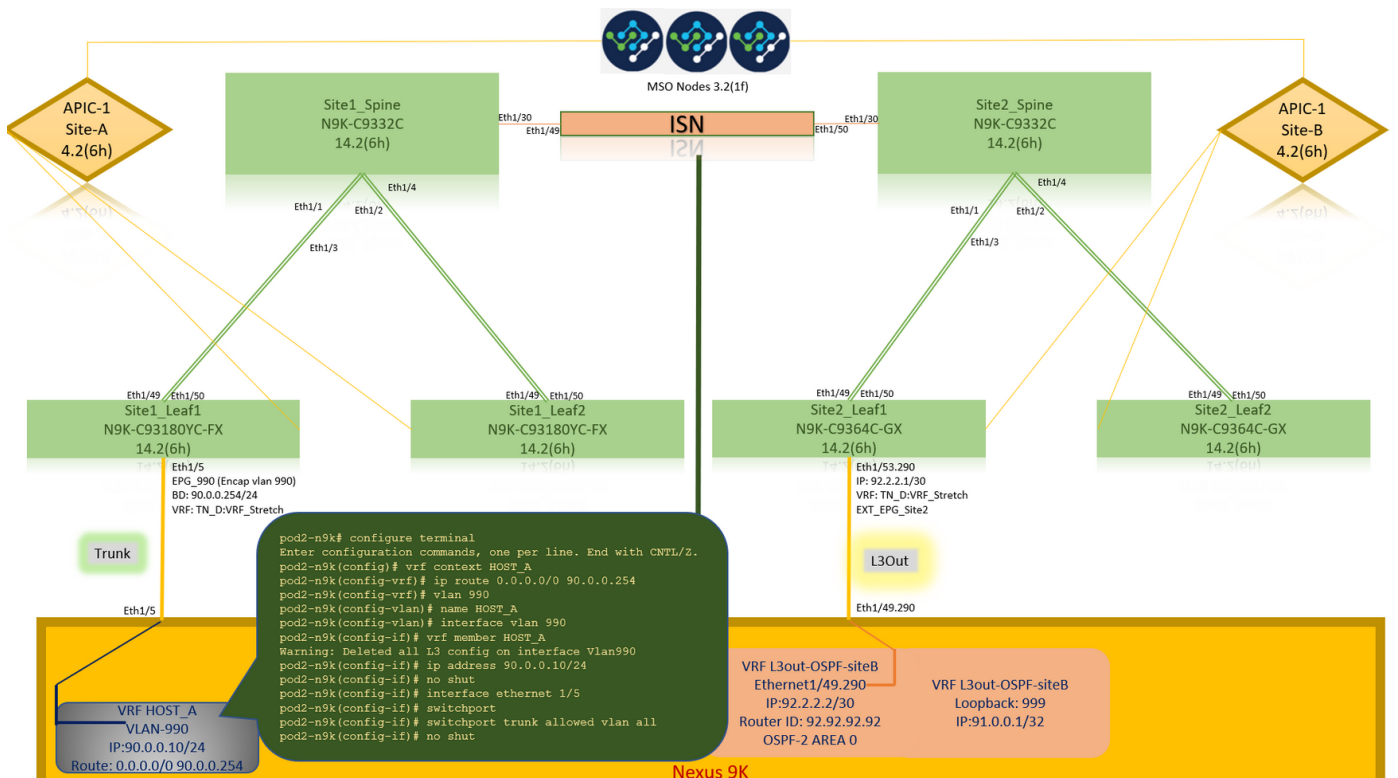
步驟3. 部署Site-A模板。

1. 從建立的架構中選擇Site-A Template。
2. 按一下Deploy to sites。



配置主機A(N9K)

使用VRF HOST_A配置N9K裝置。在N9K配置完成後，可以看到ACI枝葉BD任播地址 (HOST_A的網關) 現在可通過ICMP(ping)訪問。



在ACI操作頁籤中，您可以看到90.0.0.10 (HOST_A IP地址) 已獲取。

The screenshot displays the Cisco APIC interface for the TN_D tenant. The left sidebar shows the navigation tree with 'EPG_990' selected under 'Application EPGs'. The main area shows the configuration for 'EPG - EPG_990' in the 'Operational' tab, with the 'Client End-Points' sub-tab active. A table lists the end-points, showing one with IP 90.0.0.10 on interface Pod-1/Node-1101/eth1/5. Below the table, a terminal window shows the output of a ping command from pod2-n9k to 90.0.0.254 in vrf HOST_A, resulting in a 20.00% packet loss.

End Point	MAC	IP	Learning Source	Hosting Server	Reporting Controller Name	Interface	Multicast Address	Encap
EP-C0:14.FE.5E:1...	C0:14.FE.5E:14:07	90.0.0.10	learned	---	---	Pod-1/Node-1101/eth1/5 (learned)	---	vlan-990

```

pod2-n9k# ping 90.0.0.254 vrf HOST_A
PING 90.0.0.254 (90.0.0.254): 56 data bytes
36 bytes from 90.0.0.10: Destination Host Unreachable
Request 0 timed out
64 bytes from 90.0.0.254: icmp_seq=1 ttl=63 time=0.902 ms
64 bytes from 90.0.0.254: icmp_seq=2 ttl=63 time=0.576 ms
64 bytes from 90.0.0.254: icmp_seq=3 ttl=63 time=0.708 ms
64 bytes from 90.0.0.254: icmp_seq=4 ttl=63 time=0.659 ms

--- 90.0.0.254 ping statistics ---
5 packets transmitted, 4 packets received, 20.00% packet loss
round-trip min/avg/max = 0.576/0.711/0.902 ms
pod2-n9k#
    
```

建立Site-B模板

步驟1.從建立的架構中選擇TEMPLATES。按一下+並建立一個名為Site-B Template的模板。

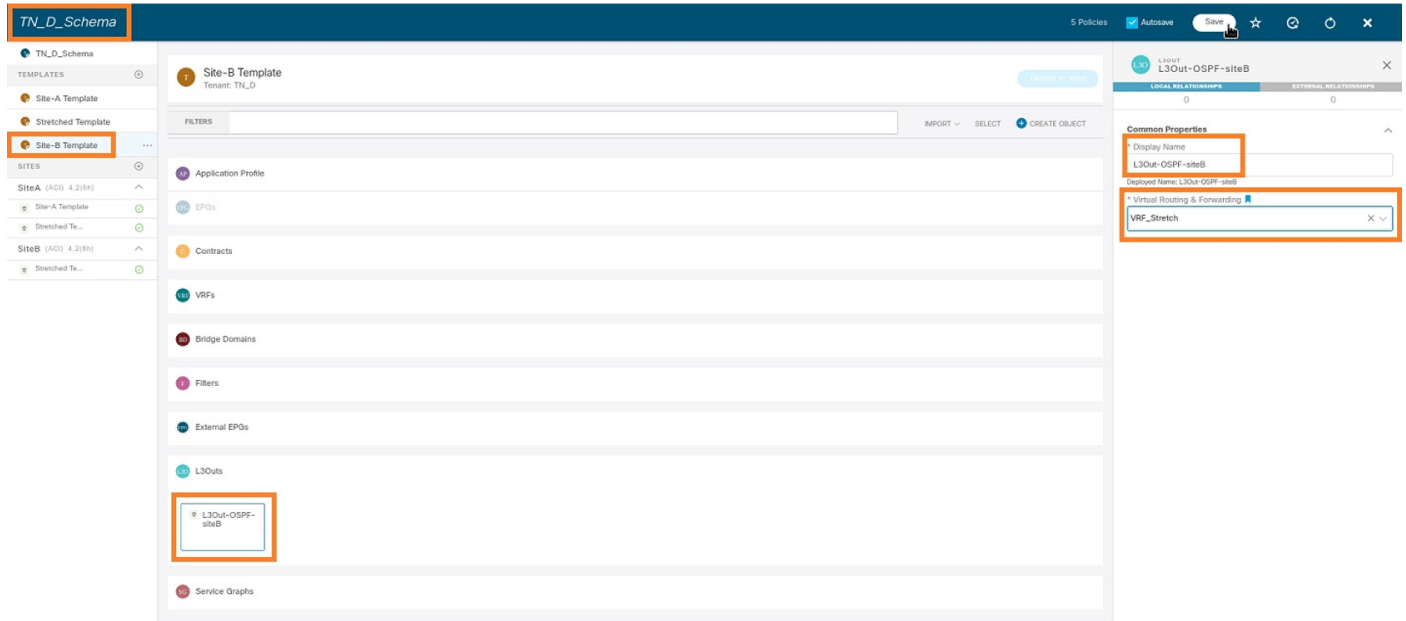


The screenshot shows the 'Site-B Template' configuration page in the APIC. The left sidebar lists the templates, with 'Site-B Template' selected. The main area shows the configuration details for the template, including the display name 'Site-B Template' and tenant settings for 'TN_D'. A notification at the top indicates that the schema was saved.

配置站點B L3out

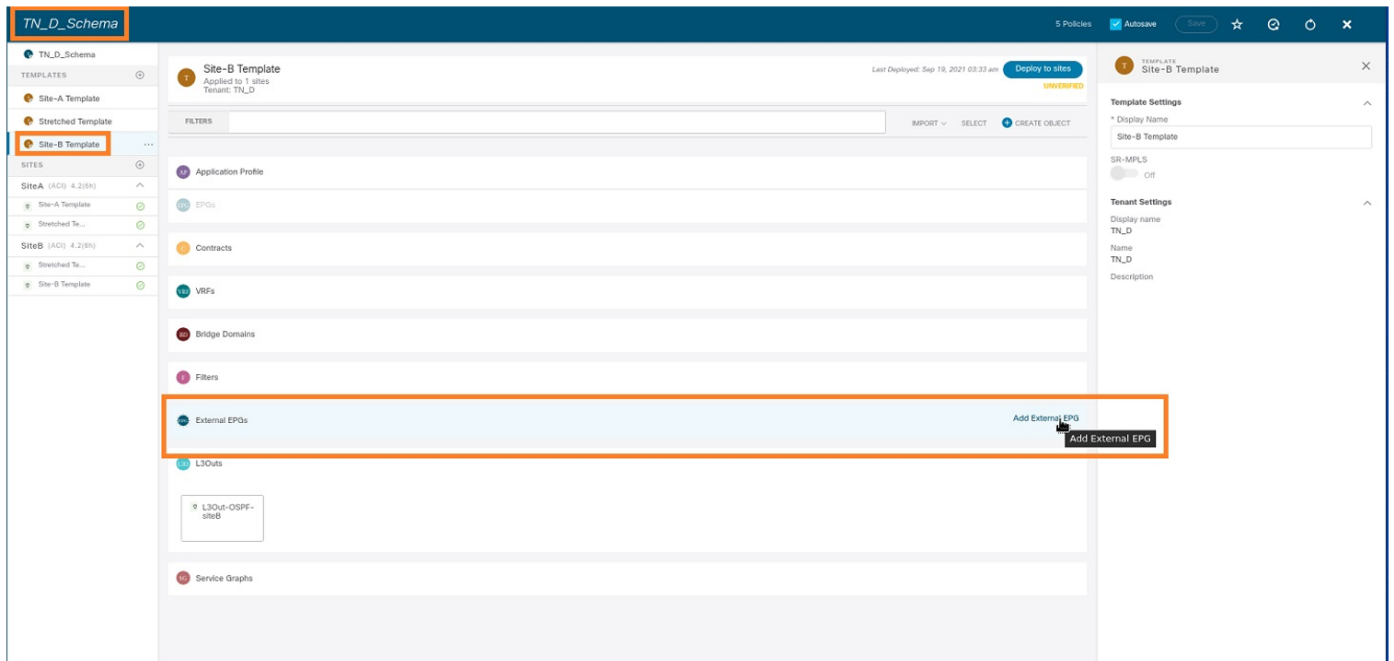
建立L3out並連線VRF_Stretch。您必須從MSO建立L3out對象，而其餘的L3out配置需要從APIC完成（因為L3out引數在MSO中不可用）。此外，從MSO建立外部EPG（僅在Site-B模板中，因為外部EPG未延伸）。

步驟1.從建立的架構中選擇Site-B Template。在「Display Name」欄位中，輸入L3out_OSPF_siteB。在「Virtual Routing & Forwarding」下拉選單中，選擇「VRF_Stretch」。



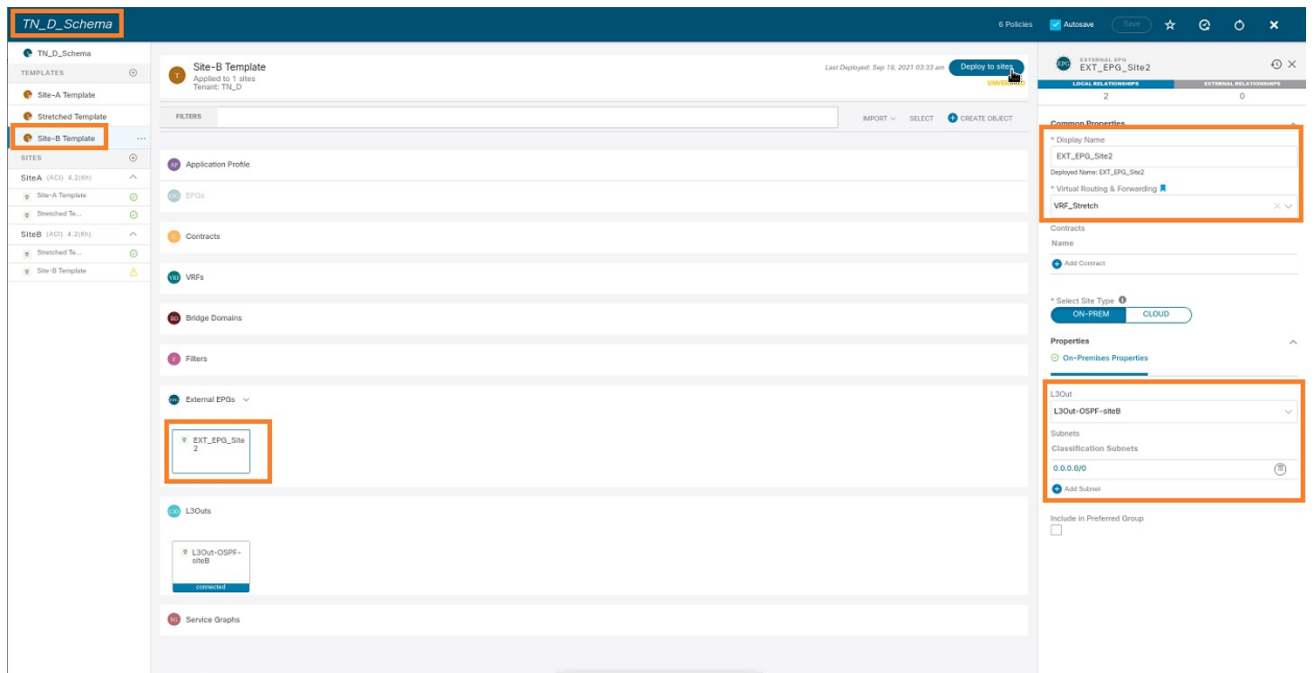
建立外部EPG

步驟1.從建立的架構中選擇Site-B Template。按一下Add External EPG。



步驟2.使用外部EPG連線L3out。

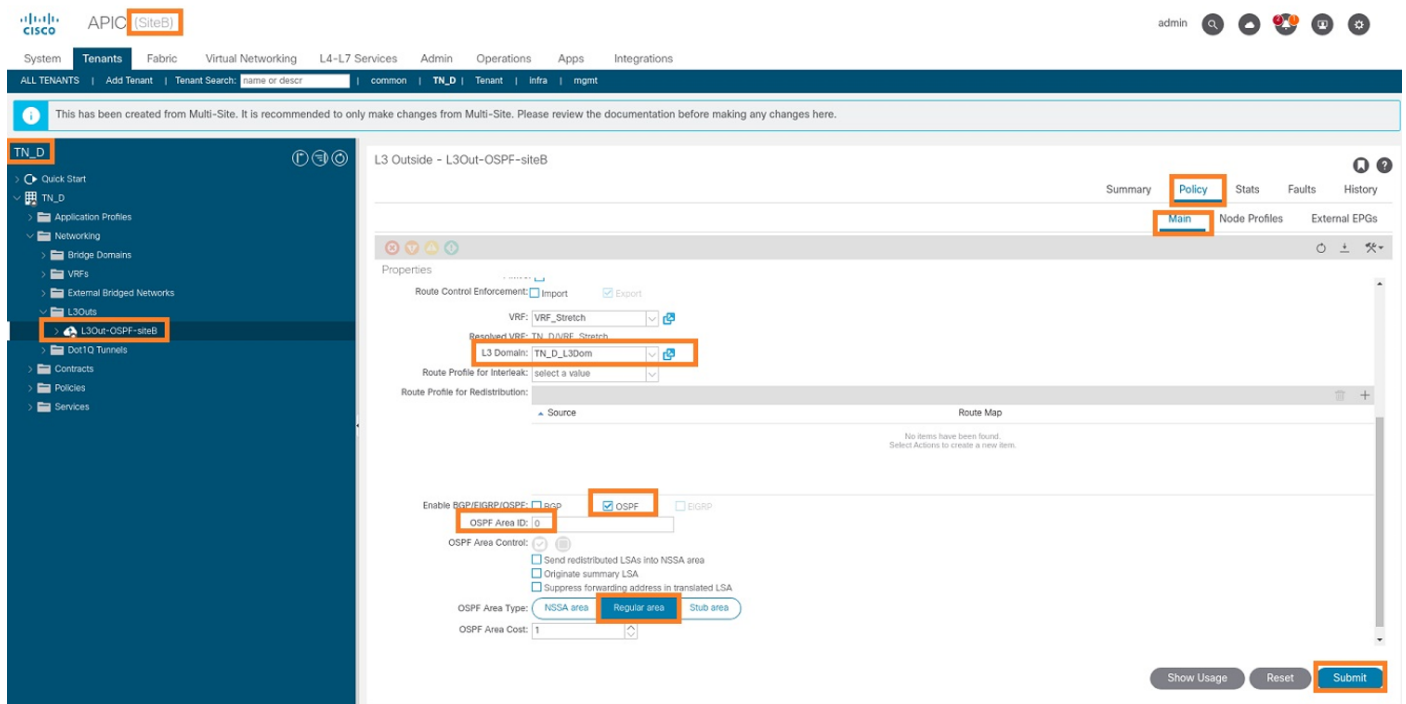
1. 從建立的架構中選擇Site-B Template。
2. 在「Display Name」欄位中，輸入EXT_EPG_Site2。
3. 在Classification Subnets欄位中，為外部EPG的外部子網輸入0.0.0.0/0。



L3out配置的其餘部分從APIC (站點B) 完成。

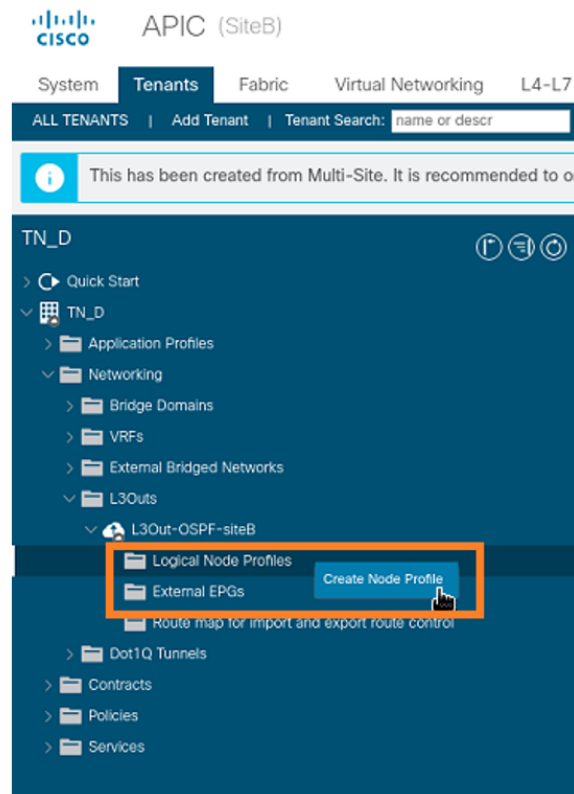
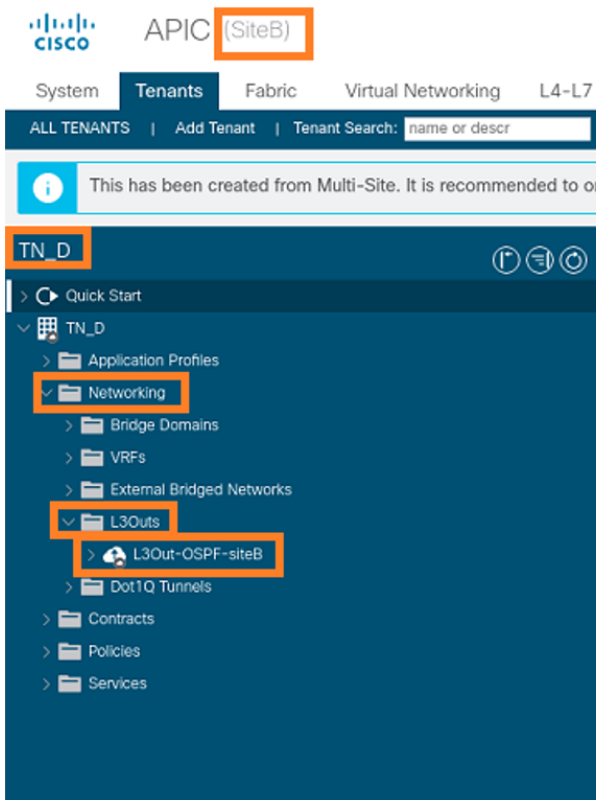
步驟3.新增L3域，啟用OSPF協定，並使用常規區域0配置OSPF。

1. 從Site-B的APIC-1中選擇TN_D > Networking > L3out-OSPF-siteB > Policy > Main。
2. 在「L3 Domain」下拉選單中，選擇TN_D_L3Dom。
3. 選中啟用BGP/EIGRP/OSPF的OSPF覈取方塊。
4. 在OSPF Area ID欄位中輸入0。
5. 在OSPF Area Type中選擇Regular area。
6. 按一下「Submit」。



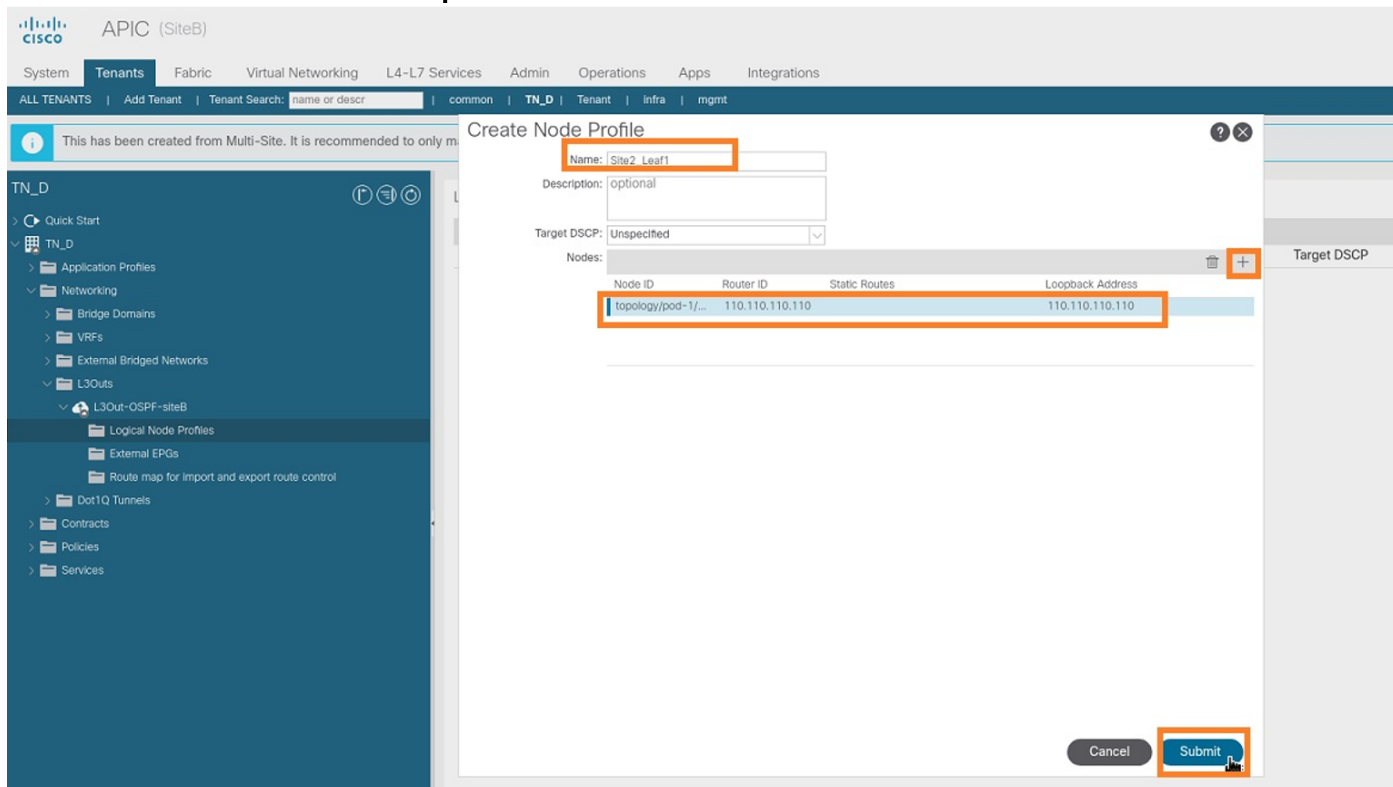
步驟4.建立節點配置檔案。

1. 從Site-B的APIC-1中選擇TN_D > Networking > L3Outs > L3Out-OSPF-siteB > Logical Node Profiles。
2. 按一下建立節點配置檔案。




步驟5.選擇交換機Site2_Leaf1作為站點B的節點。

1. 從Site-B的APIC-1中選擇TN_D > Networking > L3Outs > L3Out-OSPF-siteB > Logical Node Profiles > Create Node Profile。
2. 在「Name」欄位中，輸入Site2_Leaf1。
3. 按一下+符號新增節點。
4. 使用路由器ID IP地址新增pod-2 node-101。



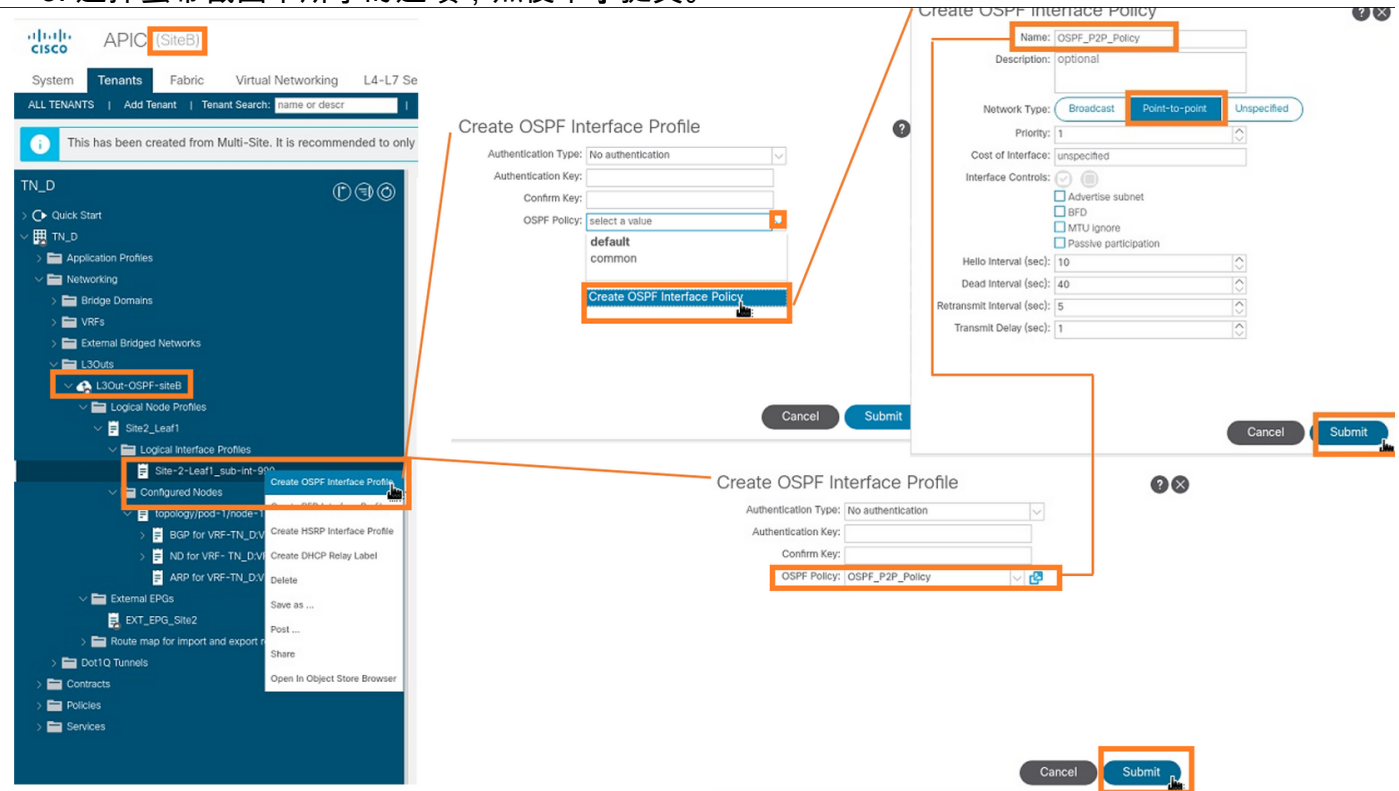
步驟6.新增介面配置檔案(外部VLAN為920 (SVI建立))。

1. 從Site-B的APIC-1中選擇TN_D > Networking > L3Outs > L3out-OSPF-SiteB > Logical Interface Profiles。
2. 按一下右鍵並新增介面配置檔案。
3. 選擇Routed Sub-Interfaces。
4. 配置IP地址、MTU和VLAN-920。



步驟7. 建立OSPF策略 (點對點網路)。

1. 從Site-B的APIC-1中選擇TN_D > Networking > L3Outs > L3Out-OSPF-siteB > Logical Interface Profiles。
2. 按一下右鍵並選擇建立OSPF介面配置檔案。
3. 選擇螢幕截圖中所示的選項，然後單擊提交。



步驟8. 驗證在TN_D > Networking > L3Outs > L3Out-OSPF-siteB > Logical Interface Profiles > (interface profile) > OSPF Interface Profile 下連線的OSPF介面配置檔案策略。

System | **Tenants** | Fabric | Virtual Networking | L4-L7 Services | Admin | Operations | Apps | Integrations

ALL TENANTS | Add Tenant | Tenant Search: name or descr | common | **TN_D** | Tenant | infra | mgmt

This has been created from Multi-Site. It is recommended to only make changes from Multi-Site. Please review the documentation before making any changes here.

TN_D

- Quick Start
- TN_D
 - Application Profiles
 - Networking
 - Bridge Domains
 - VRFs
 - External Bridged Networks
 - L3Outs
 - L3Out-OSPF-siteB
 - Logical Node Profiles
 - Site2_Leaf1
 - Logical Interface Profiles
 - OSPF Interface Profile**
 - Configured Nodes
 - topology/pod-1/node-1101
 - ARP for VRF-TN_D:VRF_Stretch
 - BGP for VRF-TN_D:VRF_Stretch
 - ND for VRF-TN_D:VRF_Stretch
 - OSPF for VRF-TN_D:VRF_Stretch
 - External EPGs
 - EXT_EPG_Site2

Interface Profile - OSPF Interface Profile

Properties

Name:

Description: optional

Authentication Key:

Confirm Authentication Key:

Authentication Key ID: 1

Authentication Type: **MD5 authentication** | No authentication | Simple authentication

Associated OSPF Interface Policy Name: **OSPF_P2P_Policy**

步驟9.驗證外部EPG「EXT_EPG_Site2」是否由MSO建立。從Site-B的APIC-1中選擇TN_D > L3Outs > L3Out-OSPF-siteB > External EPGs > EXT_EPG_Site2。

System | **Tenants** | Fabric | Virtual Networking | L4-L7 Services | Admin | Operations | Apps | Integrations

ALL TENANTS | Add Tenant | Tenant Search: name or descr | common | **TN_D** | Tenant | infra | mgmt

This has been created from Multi-Site. It is recommended to only make changes from Multi-Site. Please review the documentation before making any changes here.

TN_D

- Quick Start
- TN_D
 - Application Profiles
 - Networking
 - Bridge Domains
 - VRFs
 - External Bridged Networks
 - L3Outs
 - L3Out-OSPF-siteB
 - Logical Node Profiles
 - External EPGs
 - EXT_EPG_Site2**
 - Route map for import and export route control
 - Dot1Q Tunnels
 - Contracts
 - Policies
 - Services

External EPG Instance Profile - EXT_EPG_Site2

Policy | Operational | Stats | Health | Faults | History

General | Contracts | Inherited Contracts

Properties

Name: EXT_EPG_Site2

Alias:

Tags:

Global Alias:

Description: optional

pcTag: 32770

Contract Exception Tag:

Configured VRF Name: VRF_Stretch

Resolved VRF: uri/tn-TN_D/ctx-VRF_Stretch

QoS Class: Unspecified

Target DSCP: Unspecified

Configuration Status: applied


Configuration Issues:

Preferred Group Member: **Exclude** | Include

| Subnets | IP Address | Scope | Name | Aggregate | Route Control Profile | Route Summarization Policy |
|---------|------------|-------|------------------------------------|-----------|-----------------------|----------------------------|
| | 0.0.0.0/0 | | External Subnets for the Extern... | | | |

Show Usage | Reset | Submit

配置外部N9K(Site-B)



在N9K配置(VRF L3out-OSPF-siteB)之後，我們可以看到N9K與ACI枝葉 (位於Site-B) 之間已建立 OSPF 鄰居關係。

檢驗OSPF鄰居關係是否已建立並處於UP (完全狀態)。

從站點B的APIC-1中選擇TN_D > Networking > L3Outs > L3Out-OSPF-siteB > Logical Node Profiles > Logical Interface Profiles > Configured Nodes > topology/pod01/node-1101 > OSPF for VRF-TN_DVRF_Switch > Neighbor ID state > Full。

| Neighbor Id | State | Peer Ip | Interface |
|-------------|-------|----------|------------|
| 92.92.92.92 | Full | 92.2.2.2 | eth1/53.25 |

```

VRF L3out-OSPF-siteB
Ethernet1/49.290
IP:92.2.2.2/30
Router ID: 92.92.92.92
OSPF-2 AREA 0
  
```

```

VRF L3out-OSPF-siteB
Loopback: 999
IP:91.0.0.1/32
  
```

您還可以檢查N9K中的OSPF鄰居關係。此外，您還可以對ACI枝葉IP(Site-B)執行ping操作。

```

pod2-n9k(config-if)# ping 92.2.2.1 vrf L3out-OSPF-siteB
PING 92.2.2.1 (92.2.2.1): 56 data bytes
64 bytes from 92.2.2.1: icmp_seq=0 ttl=63 time=0.734 ms
64 bytes from 92.2.2.1: icmp_seq=1 ttl=63 time=0.591 ms
64 bytes from 92.2.2.1: icmp_seq=2 ttl=63 time=0.631 ms
64 bytes from 92.2.2.1: icmp_seq=3 ttl=63 time=0.588 ms
64 bytes from 92.2.2.1: icmp_seq=4 ttl=63 time=0.654 ms

--- 92.2.2.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.588/0.639/0.734 ms

```

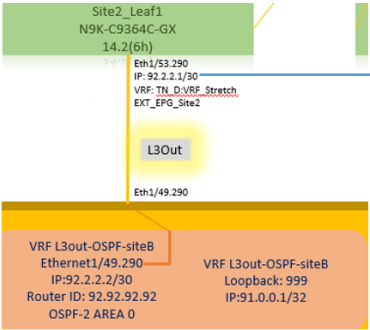
```

pod2-n9k(config-if)# show ip ospf neighbors vrf L3out-OSPF-siteB
OSPF Process ID TN_D_OSPF VRF L3out-OSPF-siteB
Total number of neighbors: 1
Neighbor ID      Pri State           Up Time  Address      Interface
110.110.110.110  1 FULL/ -          00:06:47 92.2.2.1    Eth1/49.290

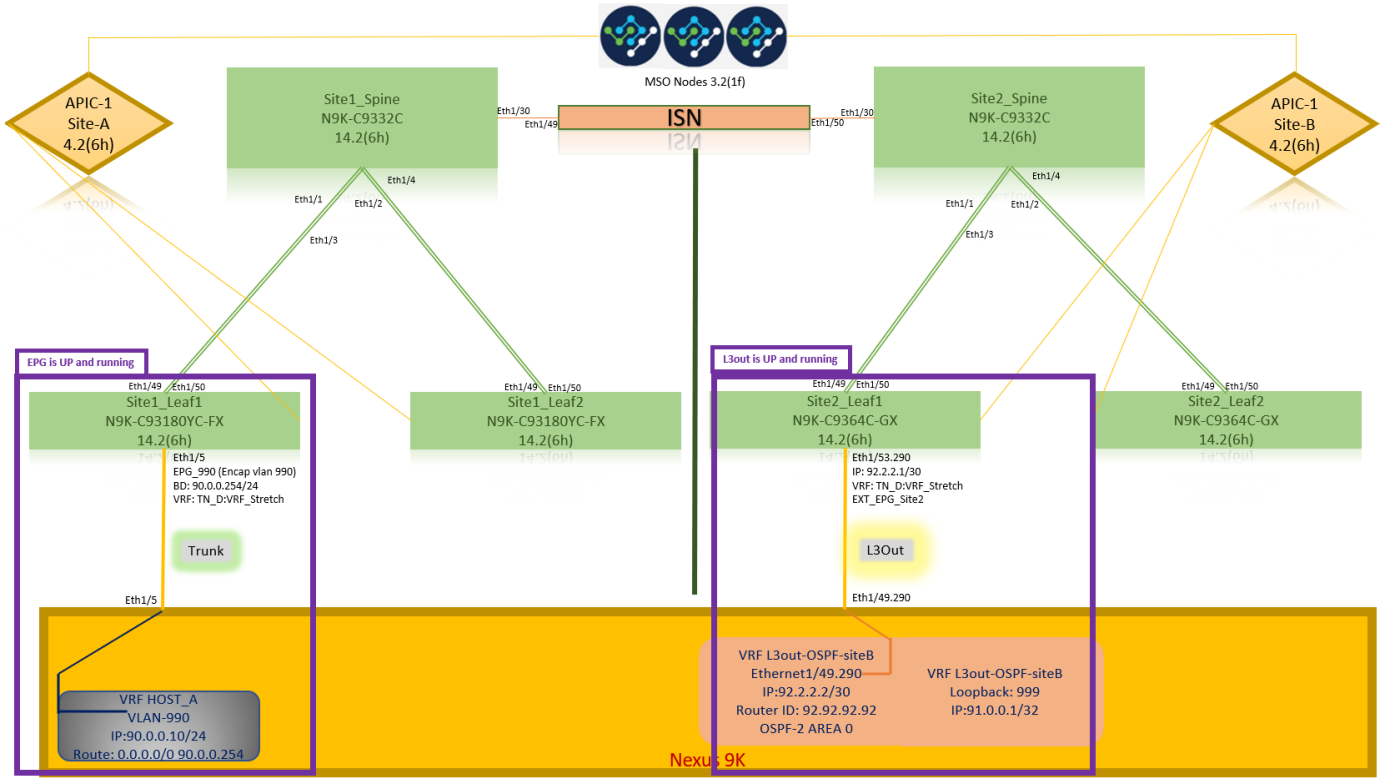
pod2-n9k(config-if)# show ip route vrf L3out-OSPF-siteB
IP Route Table for VRF "L3out-OSPF-siteB"
'*' denotes best ucast next-hop
'***' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

92.2.2.0/30, ubest/mbest: 1/0, attached
  *via 92.2.2.2, Eth1/49.290, [0/0], 00:19:38, direct
92.2.2.2/32, ubest/mbest: 1/0, attached
  *via 92.2.2.2, Eth1/49.290, [0/0], 00:19:38, local
110.110.110.110/32, ubest/mbest: 1/0
  *via 92.2.2.1, Eth1/49.290, [110/2], 00:06:48, ospf-TN_D_OSPF, intra

```



此時，站點A的Host_A配置和站點B的L3out配置已完成。

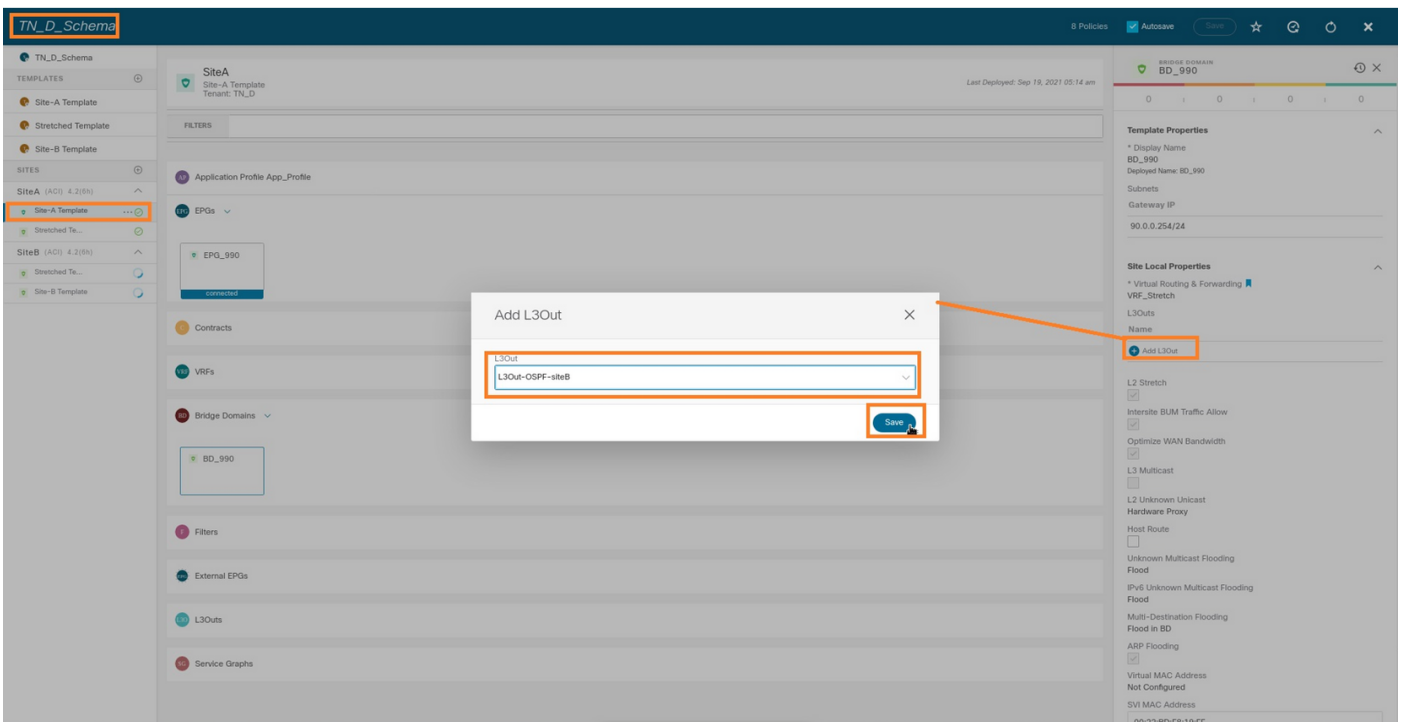



將站點B L3out連線到站點A EPG(BD)

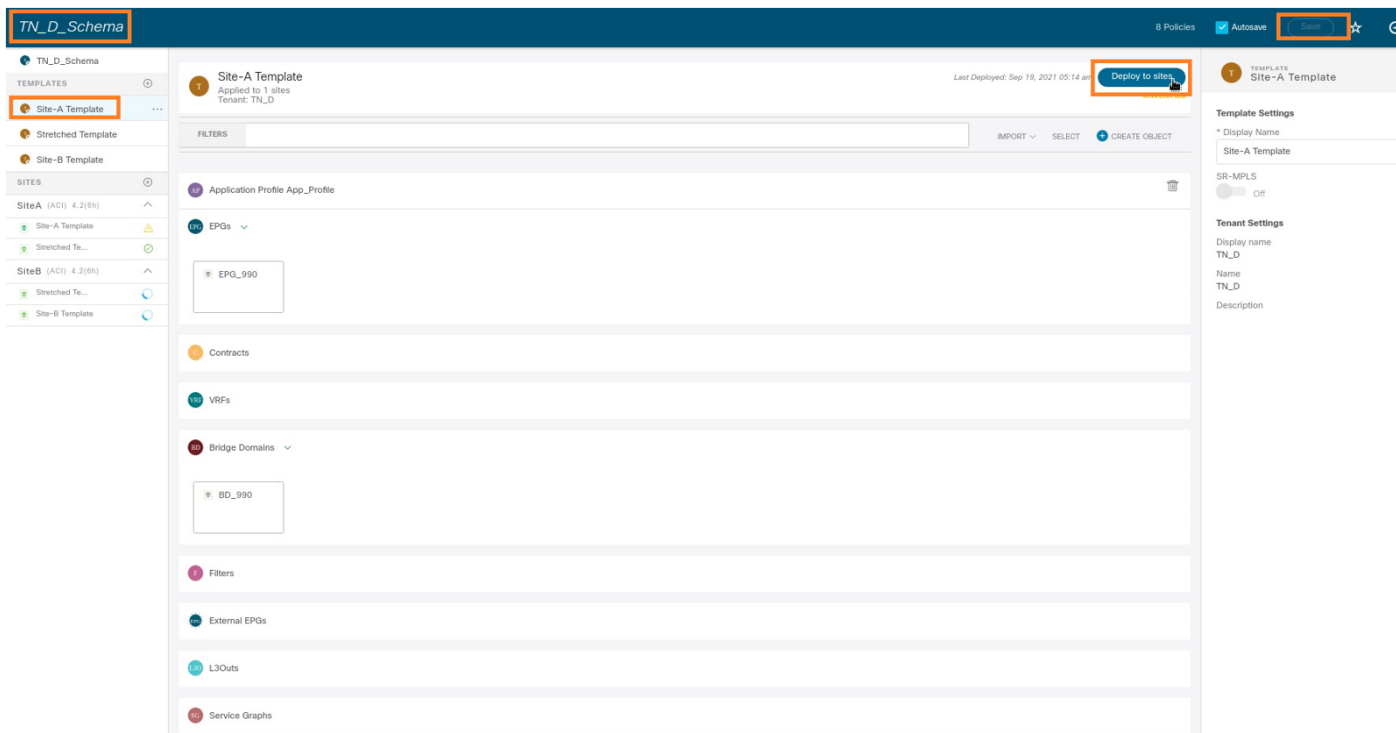
接下來，您可以將Site-B L3out連線到MSO的Site-A BD-990。請注意，左側列有兩個部分：1)模板和2)站點。

步驟1。在「站點」的第二部分，您可以看到每個站點附帶的模板。將L3out附加到「站點A模板」時，您基本上是從「站點」部分中已附加的模板進行附加。

但是，部署模板時，請從模板>站點 — A模板部分進行部署，然後選擇儲存/部署到站點。



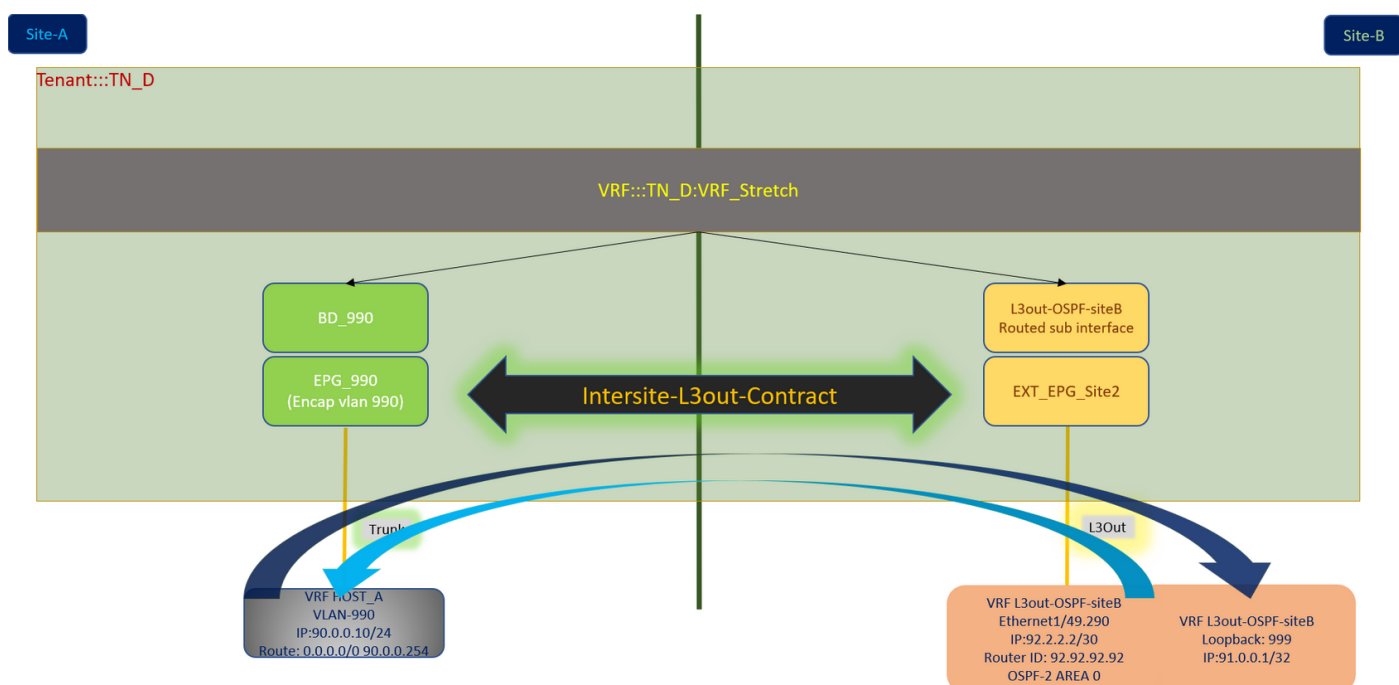
步驟2.從第一部分「模板」中的「站點A模板」主模板進行部署。



設定合約

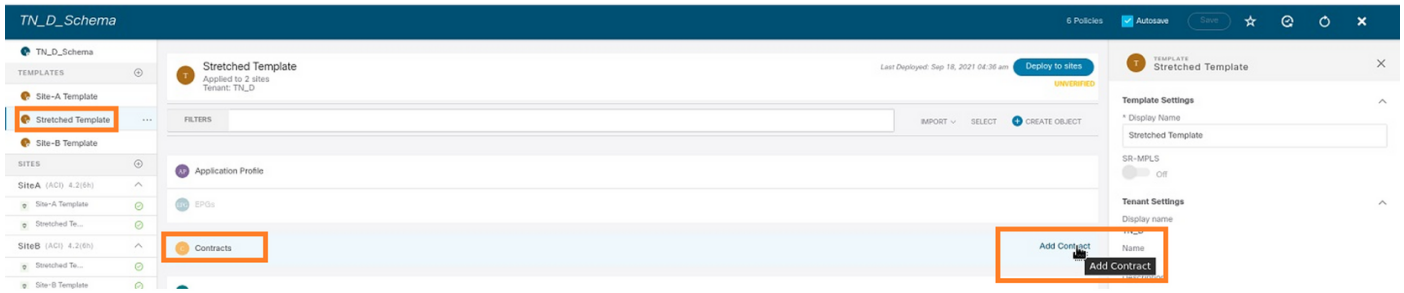
您需要在站點B的外部EPG和站點A的內部EPG_990之間簽訂合約。因此，您可以先從MSO建立合約，並將其附加到兩個EPG。

[思科以應用為中心的基礎設施 — 思科ACI合同指南](#)可幫助瞭解該合約。通常，內部EPG配置為提供者，外部EPG配置為消費者。



建立合約

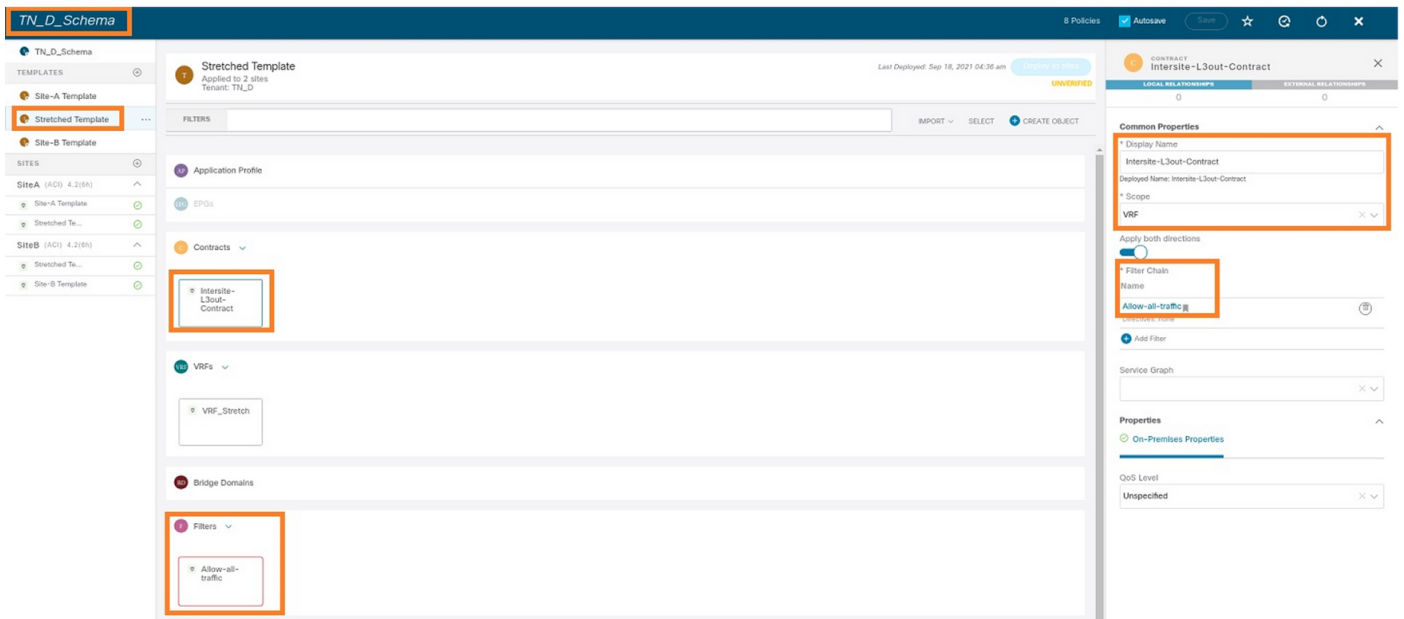
步驟1.從TN_D_Schema中選擇**拉伸模板**>**合約**。按一下 **新增合約**。



步驟2.新增過濾器以允許所有流量。

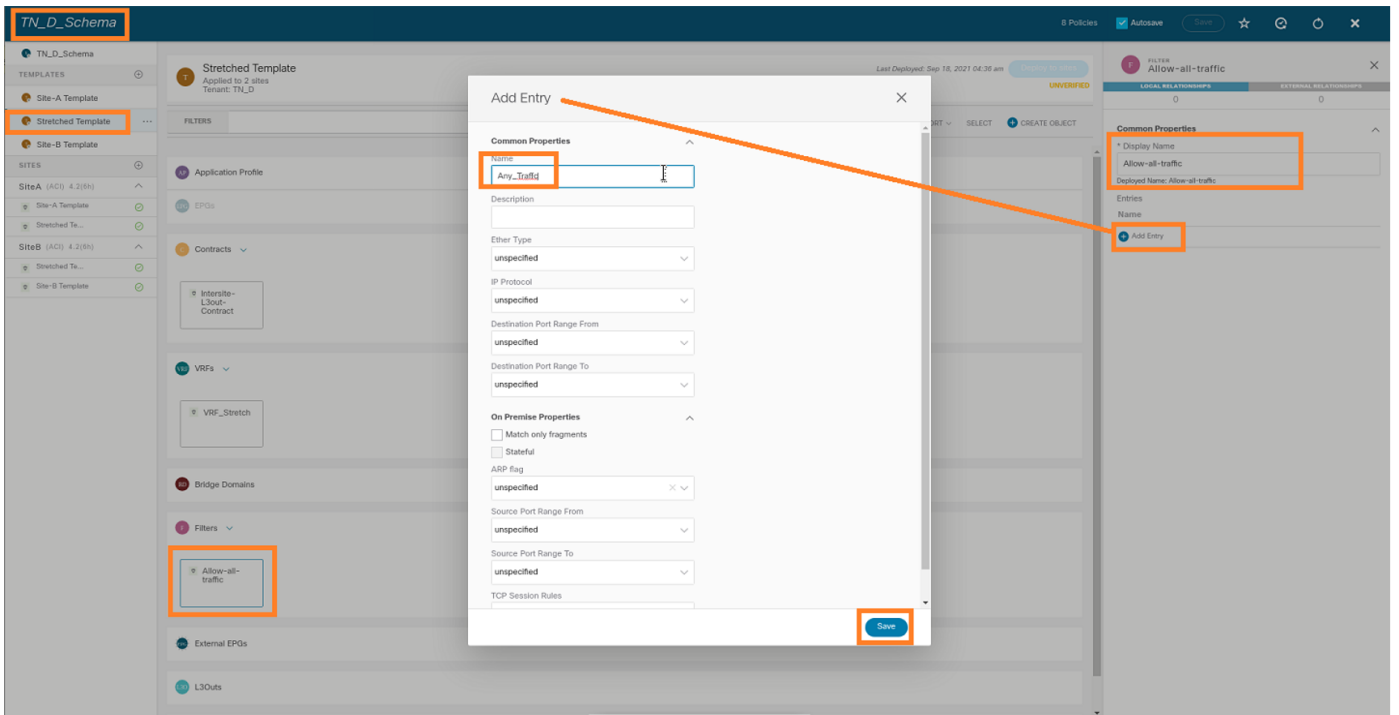
1. 從TN_D_Schema中，選擇拉伸模板>合約。
2. 新增合約：

- 顯示名稱：站點間 — L3out-Contract
- 範圍：VRF



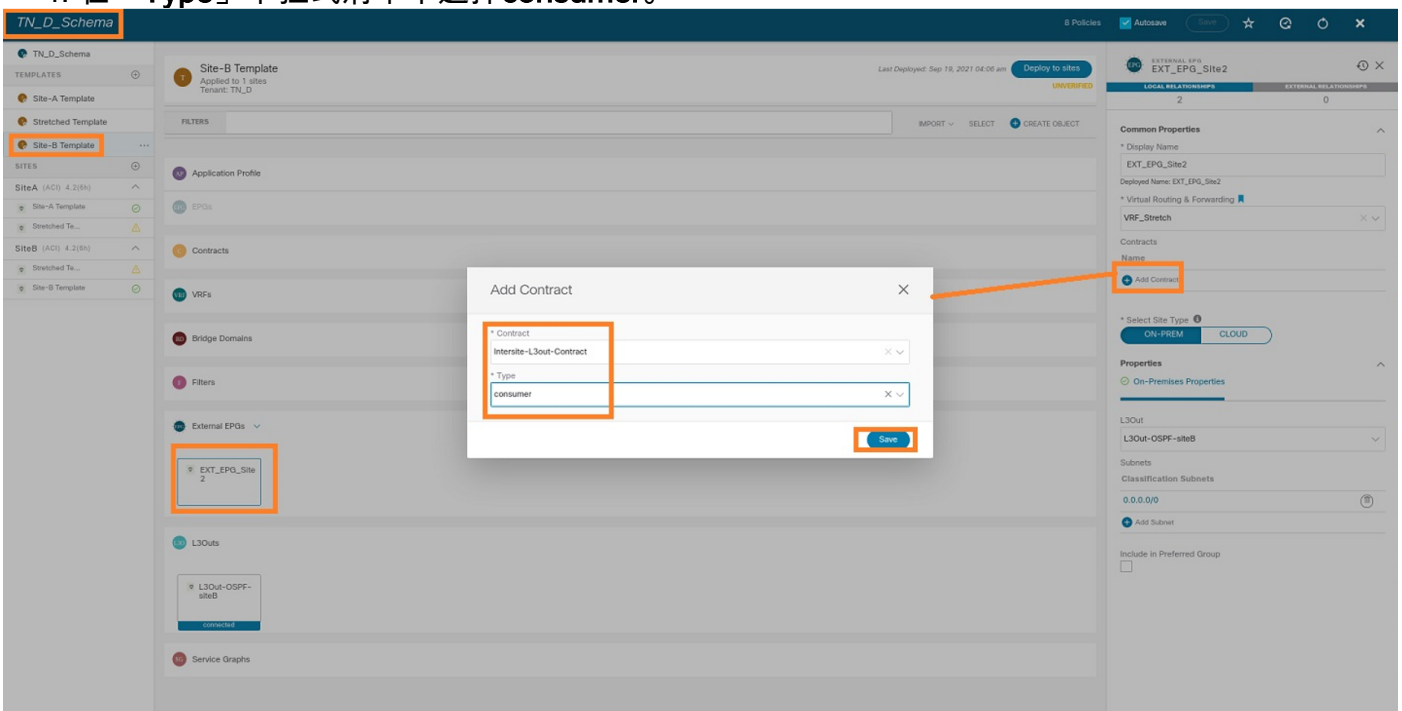
步驟3.

1. 在TN_D_Schema中，選擇拉伸模板>篩選器。
2. 在「Display Name」欄位中，輸入Allow-all-traffic。
3. 按一下「Add Entry」。將顯示「新增條目」對話方塊。
4. 在「Name」欄位中，輸入Any_Traffic。
5. 在「Ether Type」下拉選單中，選擇「unspecified」以允許所有流量。
6. 按一下Save。




步驟4.將合約作為「消費者」新增到外部EPG（在站點B模板中）（部署到站點）。

1. 從TN_D_Schema中選擇Site-B Template > EXT_EPG_Site2。
2. 按一下「Add Contract」。系統隨即會顯示「新增合約」對話方塊。
3. 在「Contract」欄位中，輸入Intersite-L3out-Contract。
4. 在「Type」下拉式清單中選擇consumer。



步驟5.將合約作為「Provider」（在站點A模板中）新增到內部EPG「EPG_990」（部署到站點）。

1. 從TN_D_Schema中選擇Site-A Template > EPG_990。
2. 按一下「Add Contract」。系統隨即會顯示「新增合約」對話方塊。
3. 在「Contract」欄位中，輸入Intersite-L3out-Contract。
4. 在「Type」下拉式清單中選擇provider。



新增合約後，您就可以看到站點A上建立的「影子L3out /外部EPG」。



APIC (SiteA)

System

Tenants

Fabric

Virtual Networking

L4-L7

ALL TENANTS

| Add Tenant

| Tenant Search:

name or descr



This has been created from Multi-Site. It is recommended to or

TN_D



> Quick Start

▼ TN_D

> Application Profiles

▼ Networking

> Bridge Domains

> VRFs

> External Bridged Networks

▼ L3Outs

▼ L3Out-OSPF-siteB

Shadow L3out site-B

Logical Node Profiles

▼ External EPGs

EXT_EPG_Site2

Shadow Ext EPG

> Route map for import and export route control

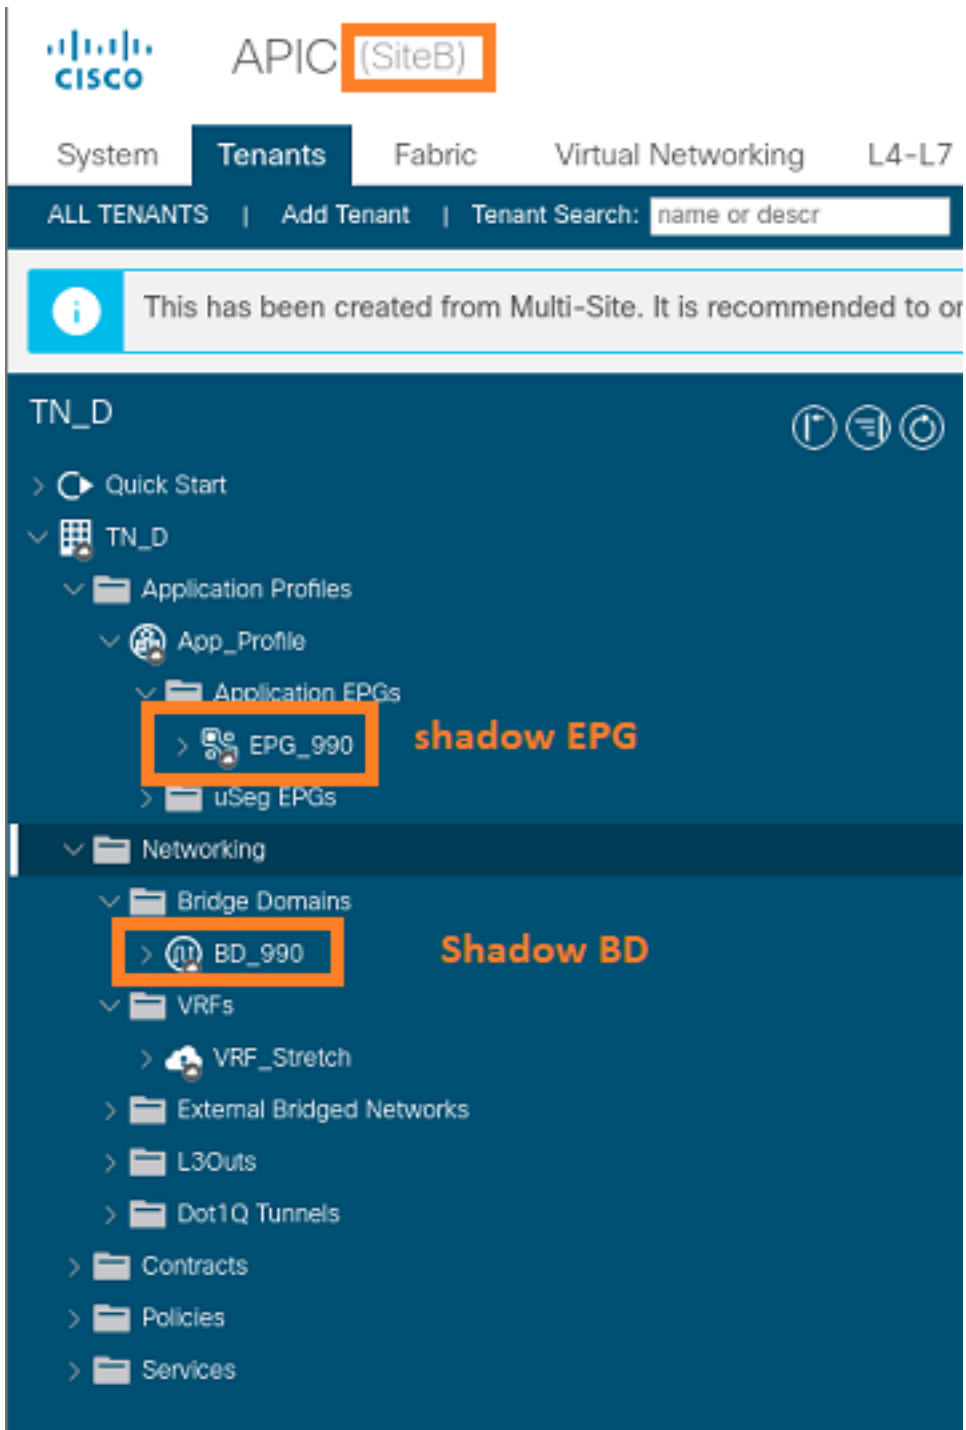
> Dot1Q Tunnels

> Contracts

> Policies

> Services

您還可以看到，站點B上還建立了「陰影EPG_990和BD_990」。



步驟6.輸入以下命令以驗證Site-B APIC。

```
apic1# moquery -c fvAEPg -f 'fv.AEPg.name=="EPG_990"'
Total Objects shown: 1
# fv.AEPg
name                : EPG_990
annotation          : orchestrator:msc
childAction         :
configIssues        :
configSt            : applied
descr               :
dn                  : uni/tn-TN_D/ap-App_Profile/epg-EPG_990
exceptionTag        :
extMngdBy           :
floodOnEncap        : disabled
fwdCtrl             :
```

```

hasMcastSource      : no
isAttrBasedEPg     : no
isSharedSrvMsiteEPg : no
lcOwn               : local
matchT              : AtleastOne
modTs               : 2021-09-19T18:47:53.374+00:00
monPolDn            : uni/tn-common/monepg-default
nameAlias           :
pcEnfPref           : unenforced
pcTag              : 49153          <<< Note that pcTag is different for shadow EPG.
prefGrMemb          : exclude
prio                : unspecified
rn                  : epg-EPG_990
scope               : 2686978
shutdown            : no
status              :
triggerSt           : triggerable
txId                : 1152921504609244629
uid                 : 0

```

```
apic1# moquery -c fvBD -f 'fv.BD.name=="BD_990\"'
```

```
Total Objects shown: 1
```

```
# fv.BD
```


```

name              : BD_990
OptimizeWanBandwidth : yes
annotation            : orchestrator:misc
arpFlood              : yes
bcastP                : 225.0.181.192
childAction           :
configIssues          :
descr                 :
dn                : uni/tn-TN_D/BD-BD_990
epClear               : no
epMoveDetectMode     :
extMngdBy             :
hostBasedRouting     : no
intersiteBumTrafficAllow : yes
intersiteL2Stretch   : yes
ipLearning            : yes
ipv6McastAllow       : no
lcOwn                 : local
limitIpLearnToSubnets : yes
llAddr                : ::
mac                   : 00:22:BD:F8:19:FF
mcastAllow            : no
modTs                 : 2021-09-19T18:47:53.374+00:00
monPolDn              : uni/tn-common/monepg-default
mtu                   : inherit
multiDstPktAct       : bd-flood
nameAlias             :
ownerKey              :
ownerTag              :
pcTag                 : 32771
rn                    : BD-BD_990
scope                 : 2686978
seg                   : 15957972
status                :
type                  : regular
uid                   : 0
unicastRoute          : yes
unkMacUcastAct    : proxy
unkMcastAct      : flood

```

```
v6unkMcastAct      : flood
vmac                : not-applicable
```

步驟7.檢查並驗證外部裝置N9K配置。



驗證

使用本節內容，確認您的組態是否正常運作。

終端學習

驗證是否已將Site-A端點獲知為Site1_Leaf1中的端點。

```
Site1_Leaf1# show endpoint interface ethernet 1/5
```

Legend:

```

s - arp          H - vtep          V - vpc-attached    p - peer-aged
R - peer-attached-rl B - bounce      S - static          M - span
D - bounce-to-proxy O - peer-attached a - local-aged    m - svc-mgr
L - local        E - shared-service
  
```

| VLAN/
Interface
Domain | Encap
VLAN | MAC Address
IP Address | MAC Info/
IP Info |
|------------------------------|---------------|---------------------------|----------------------|
| 18 | vlan-990 | c014.fe5e.1407 | L |
| eth1/5 | | | |
| TN_D:VRF_Stretch | vlan-990 | 90.0.0.10 | L |
| | | | eth1/5 |

ETEP/RTEP驗證

站點(_A)。

Site1_Leaf1# show ip interface brief vrf overlay-1

IP Interface Status for VRF "overlay-1"(4)

| Interface | Address | Interface Status |
|------------|---------------------------|---|
| eth1/49 | unassigned | protocol-up/link-up/admin-up |
| eth1/49.7 | unnumbered
(lo0) | protocol-up/link-up/admin-up |
| eth1/50 | unassigned | protocol-up/link-up/admin-up |
| eth1/50.8 | unnumbered
(lo0) | protocol-up/link-up/admin-up |
| eth1/51 | unassigned | protocol-down/link-down/admin-up |
| eth1/52 | unassigned | protocol-down/link-down/admin-up |
| eth1/53 | unassigned | protocol-down/link-down/admin-up |
| eth1/54 | unassigned | protocol-down/link-down/admin-up |
| vlan9 | 10.0.0.30/27 | protocol-up/link-up/admin-up |
| lo0 | 10.0.80.64/32 | protocol-up/link-up/admin-up |
| lo1 | 10.0.8.67/32 | protocol-up/link-up/admin-up |
| lo8 | 192.168.200.225/32 | protocol-up/link-up/admin-up <<<<< IP from ETEP site-A |
| lo1023 | 10.0.0.32/32 | protocol-up/link-up/admin-up |


Site2_Leaf1# show ip interface brief vrf overlay-1

IP Interface Status for VRF "overlay-1"(4)

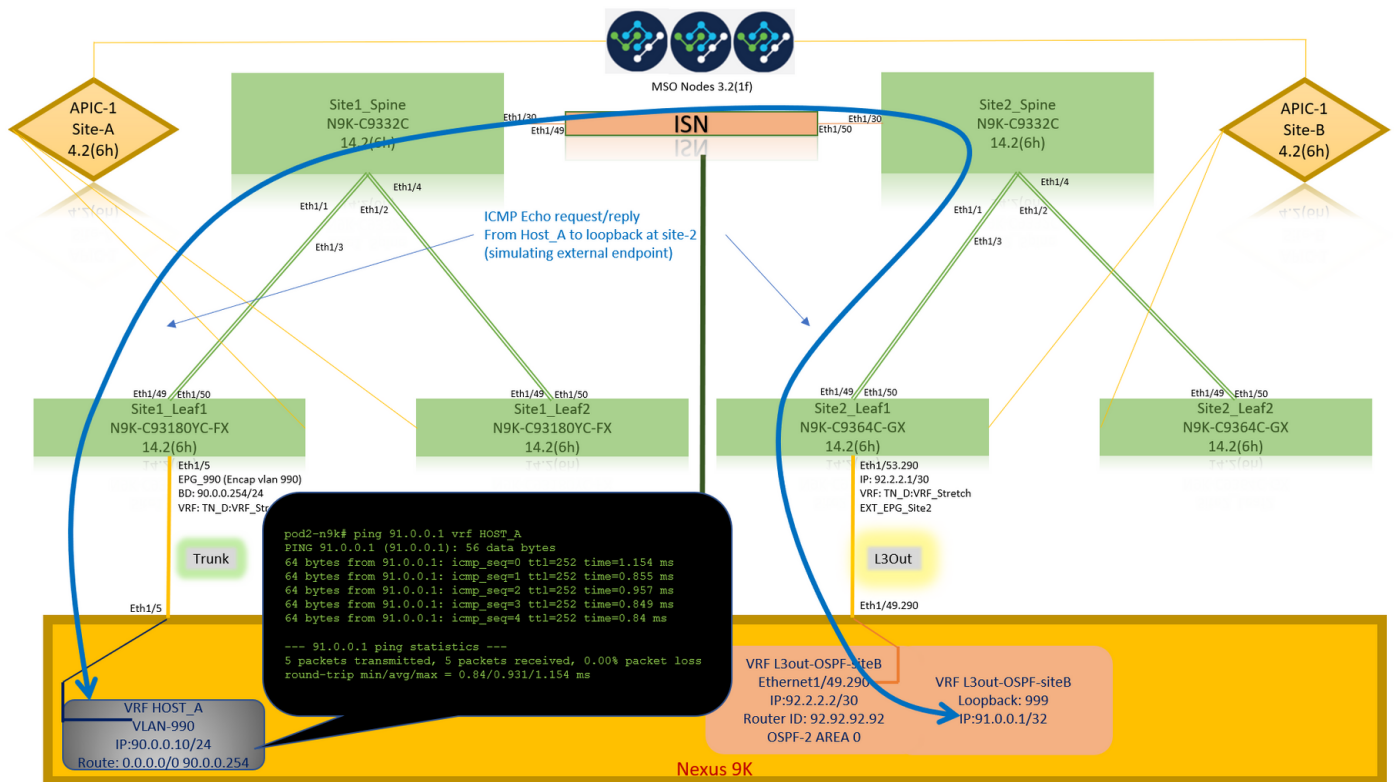
| Interface | Address | Interface Status |
|------------|---------------------------|---|
| eth1/49 | unassigned | protocol-up/link-up/admin-up |
| eth1/49.16 | unnumbered
(lo0) | protocol-up/link-up/admin-up |
| eth1/50 | unassigned | protocol-up/link-up/admin-up |
| eth1/50.17 | unnumbered
(lo0) | protocol-up/link-up/admin-up |
| eth1/51 | unassigned | protocol-down/link-down/admin-up |
| eth1/52 | unassigned | protocol-down/link-down/admin-up |
| eth1/54 | unassigned | protocol-down/link-down/admin-up |
| eth1/55 | unassigned | protocol-down/link-down/admin-up |
| eth1/56 | unassigned | protocol-down/link-down/admin-up |
| eth1/57 | unassigned | protocol-down/link-down/admin-up |
| eth1/58 | unassigned | protocol-down/link-down/admin-up |
| eth1/59 | unassigned | protocol-down/link-down/admin-up |
| eth1/60 | unassigned | protocol-down/link-down/admin-up |
| eth1/61 | unassigned | protocol-down/link-down/admin-up |
| eth1/62 | unassigned | protocol-down/link-down/admin-up |
| eth1/63 | unassigned | protocol-down/link-down/admin-up |
| eth1/64 | unassigned | protocol-down/link-down/admin-up |
| vlan18 | 10.0.0.30/27 | protocol-up/link-up/admin-up |
| lo0 | 10.0.72.64/32 | protocol-up/link-up/admin-up |
| lo1 | 10.0.80.67/32 | protocol-up/link-up/admin-up |
| lo6 | 192.168.100.225/32 | protocol-up/link-up/admin-up <<<<< IP from ETEP site-B |
| lo1023 | 10.0.0.32/32 | protocol-up/link-up/admin-up |

ICMP可達性

從HOST_A ping外部裝置WAN IP地址。



Ping外部裝置環回地址。



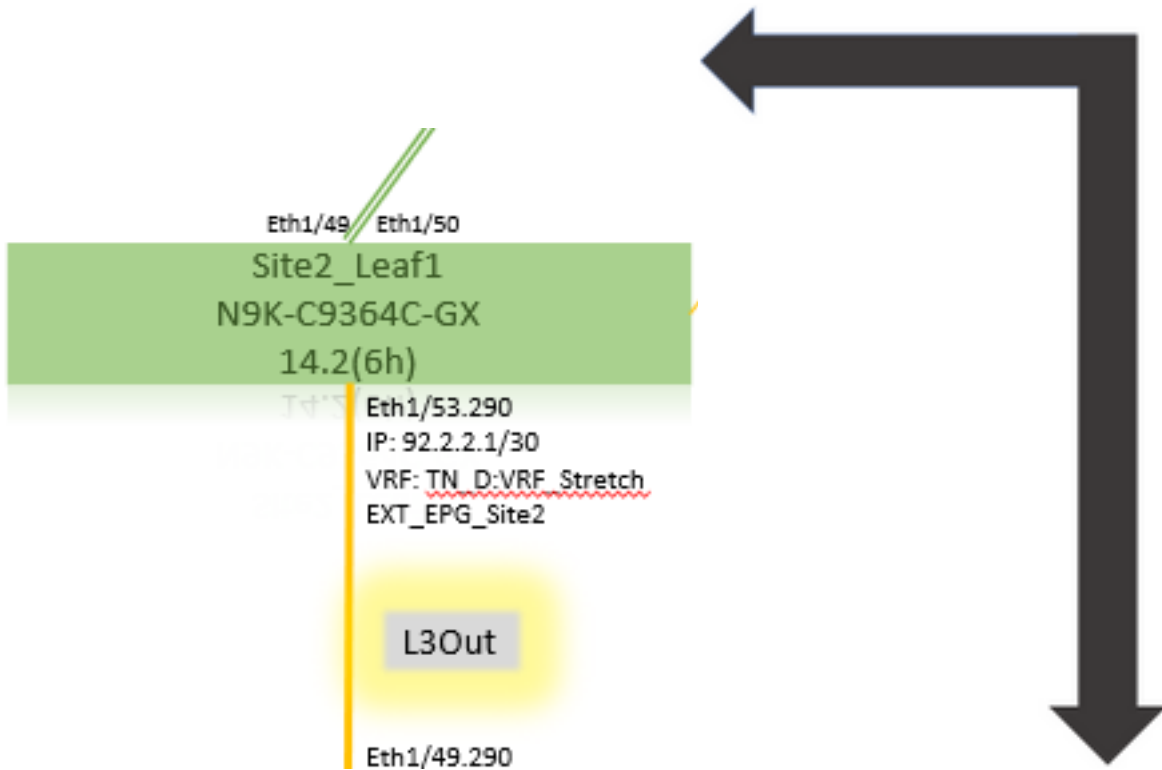
路由驗證

驗證路由表中是否存在外部裝置WAN IP地址或環回子網路。當您在「Site1_Leaf1」中檢查外部裝置子網的下一跳時，它是枝葉「Site2-Leaf1」的外部TEP IP。

| VRF | Type | VRF vnid | Context ID | Status | Endpoint Count |
|------------------|--------|----------|------------|--------|----------------|
| TN_D:VRF_Stretch | Tenant | 2686978 | 46 | Up | 1 |

Site2_Leaf1# show vrf TN_D:VRF_Stretch detail

VRF-Name: TN_D:VRF_Stretch, VRF-ID: 46, State: Up
 VPNID: unknown
RD: 1101:2686978
 Max Routes: 0 Mid-Threshold: 0
 Table-ID: 0x8000002e, AF: IPv6, Fwd-ID: 0x8000002e, State: Up
 Table-ID: 0x0000002e, AF: IPv4, Fwd-ID: 0x0000002e, State: Up



Site2_Leaf1# vsh

Site2_Leaf1# show bgp vpnv4 unicast 91.0.0.1 vrf TN_D:VRF_Stretch

BGP routing table information for VRF overlay-1, address family VPNv4 Unicast
Route Distinguisher: 1101:2686978 (VRF TN_D:VRF_Stretch)
BGP routing table entry for 91.0.0.1/32, version 12 dest ptr 0xae6da350
 Paths: (1 available, best #1)
 Flags: (0x80c0002 00000000) on xmit-list, is not in urib, exported
 vpn: version 346, (0x100002) on xmit-list
 Multipath: eBGP iBGP

Advertised path-id 1, VPN AF advertised path-id 1
 Path type: redistrib 0x408 0x1 ref 0 adv path ref 2, path is valid, is best path
 AS-Path: NONE, path locally originated

0.0.0.0 (metric 0) from 0.0.0.0 (10.0.72.64)
 Origin incomplete, MED 2, localpref 100, weight 32768
 Extcommunity:

RT:65001:2686978
VNID:2686978
 COST:pre-bestpath:162:110

VRF advertise information:
 Path-id 1 not advertised to any peer
 VPN AF advertise information:
 Path-id 1 advertised to peers:

10.0.72.65

```

apic1# acidiag fnvread ID Pod ID Name Serial Number IP Address Role State LastUpdMsgId -----
-----
----- 101 1

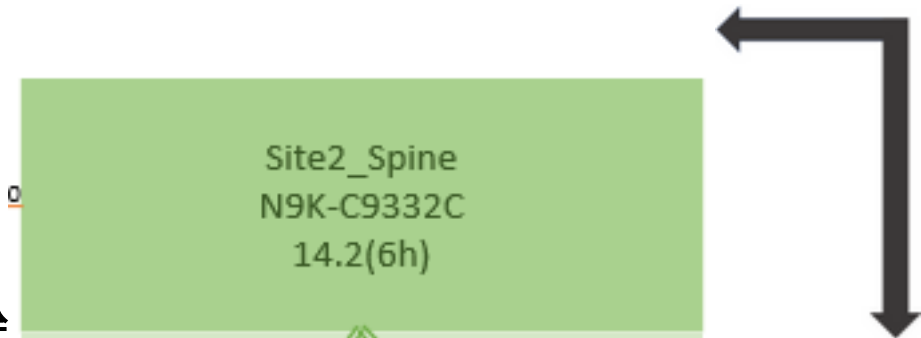
```

Site2_Spine FDO243207JH

```

10.0.72.65/32 spine active 0 102 1 Site2_Leaf2 FDO24260FCH 10.0.72.66/32 leaf active 0 1101
1 Site2_Leaf1 FDO24260ECW 10.0.72.64/32 leaf active 0

```




站點2_骨幹

```

Site2_Spine# vsh
Site2_Spine# show bgp vpnv4 unicast 91.0.0.1 vrf overlay-1
BGP routing table information for VRF overlay-1, address family VPNv4 Unicast
<-----26bits----->
Route Distinguisher: 1101:2686978 <<<<<2686978 <--
Binary--> 000010100100000000000000010
BGP routing table entry for 91.0.0.1/32, version 717 dest ptr 0xae643d0c
Paths: (1 available, best #1)
Flags: (0x000002 00000000) on xmit-list, is not in urib, is not in HW
Multipath: eBGP iBGP
  Advertised path-id 1
  Path type: internal 0x40000018 0x800040 ref 0 adv path ref 1, path is valid, is best path
  AS-Path: NONE, path sourced internal to AS
  10.0.72.64 (metric 2) from 10.0.72.64 (10.0.72.64) <<< Site2_leaf1 IP
  Origin incomplete, MED 2, localpref 100, weight 0
  Received label 0
  Received path-id 1
  Extcommunity:
    RT:65001:2686978
    COST:pre-bestpath:168:3221225472
    VNID:2686978
    COST:pre-bestpath:162:110
  Path-id 1 advertised to peers:
    192.168.10.13 <<<< Site1_Spine mscp-etest IP.
Site1_Spine# show ip interface vrf overlay-1
<snip...>
lo12, Interface status: protocol-up/link-up/admin-up, iod: 89, mode: mscp-etest
IP address: 192.168.10.13, IP subnet: 192.168.10.13/32 <<

```



站點1_骨幹

```
Site1_Spine# vsh
Site1_Spine# show bgp vpnv4 unicast 91.0.0.1 vrf overlay-1
BGP routing table information for VRF overlay-1, address family VPNv4 Unicast
<-----26Bits----->
Route Distinguisher: 1101:36241410
<<<<<36241410<--binary-->100010100100000000000000010
BGP routing table entry for 91.0.0.1/32, version 533 dest ptr 0xae643dd4
Paths: (1 available, best #1)
Flags: (0x000002 00000000) on xmit-list, is not in urib, is not in HW
Multipath: eBGP iBGP
  Advertised path-id 1
    Path type: internal 0x40000018 0x880000 ref 0 adv path ref 1, path is valid, is best path,
remote site path
  AS-Path: NONE, path sourced internal to AS
    192.168.100.225 (metric 20) from 192.168.11.13 (192.168.11.13) <<< Site2_Leaf1 ETEP IP
learn via Site2_Spine mscsp-etest address.
  Origin incomplete, MED 2, localpref 100, weight 0
  Received label 0
  Extcommunity:
    RT:65001:36241410
    SOO:65001:50331631
    COST:pre-bestpath:166:2684354560
    COST:pre-bestpath:168:3221225472
    VNID:2686978
    COST:pre-bestpath:162:110
  Originator: 10.0.72.64 Cluster list: 192.168.11.13 <<< Originator Site2_Leaf1 and
Site2_Spine ips are listed here...
  Path-id 1 advertised to peers:
    10.0.80.64 <<<< Site1_Leaf1 ip
```

```
Site2_Spine# show ip interface vrf overlay-1
<snip..>
lo13, Interface status: protocol-up/link-up/admin-up, iod: 92, mode: mscsp-etest IP address:
192.168.11.13, IP subnet: 192.168.11.13/32
  IP broadcast address: 255.255.255.255
  IP primary address route-preference: 0, tag: 0
<snip..>
```

```
Site-B apic1# acidiag fvnread
```

| ID | Pod ID | Name | Serial Number | IP Address | Role | State |
|------|--------|-------------|---------------|---------------|-------|----------|
| 101 | 1 | Site2_Spine | FDO243207JH | 10.0.72.65/32 | spine | active 0 |
| 102 | 1 | Site2_Leaf2 | FDO24260FCH | 10.0.72.66/32 | leaf | active 0 |
| 1101 | 1 | Site2_Leaf1 | FDO24260ECW | 10.0.72.64/32 | leaf | active 0 |

檢驗站點間標誌。

```
Site1_Spine# moquery -c bgpPeer -f 'bgp.Peer.addr*"192.168.11.13"'
```

```

Total Objects shown: 1
# bgp.Peer
addr           : 192.168.11.13/32
activePfxPeers : 0
adminSt        : enabled
asn            : 65001
bgpCfgFailedBmp :
bgpCfgFailedTs : 00:00:00:00.000
bgpCfgState    : 0
childAction    :
ctrl           :
curPfxPeers    : 0
dn             : sys/bgp/inst/dom-overlay-1/peer-[192.168.11.13/32]
lcOwn          : local
maxCurPeers   : 0
maxPfxPeers    : 0
modTs          : 2021-09-13T11:58:26.395+00:00
monPolDn       :
name           :
passwdSet      : disabled
password       :
peerRole       : msite-speaker
privateASctrl  :
rn             : peer-[192.168.11.13/32] <<

```

<<

瞭解路由識別符號條目當設定站點間標誌時，本地站點主幹可以從第25位開始在路由目標中設定本地站點ID。當Site1獲取在RT中設定了此位的BGP路徑時，它知道這是遠端站點路徑。

```

Site2_Leaf1# vsh
Site2_Leaf1# show bgp vpnv4 unicast 91.0.0.1 vrf TN_D:VRF_Stretch
BGP routing table information for VRF overlay-1, address family VPNv4 Unicast
<-----26Bits----->
Route Distinguisher: 1101:2686978      (VRF TN_D:VRF_Stretch)                <<<<<2686978
<--Binary--> 00001010010000000000000010
BGP routing table entry for 91.0.0.1/32, version 12 dest ptr 0xae6da350

Site1_Spine# vsh
Site1_Spine# show bgp vpnv4 unicast 91.0.0.1 vrf overlay-1

<-----26Bits----->
Route Distinguisher: 1101:36241410
<<<<<36241410<--binary-->10001010010000000000000010

^^---26th bit set to 1 and with 25th bit value it become 10.

```

請注意，除了第26位設定為1外，Site1的RT二進位制值完全相同。它有十進位制值（標籤為藍色）。1101:36241410是您期望在Site1中看到的內容，以及必須匯入Site1的內部枝葉。

Site1_Leaf1
N9K-C93180YC-FX
14.2(6h)

站點1_葉1

```
Site1_Leaf1# show vrf TN_D:VRF_Stretch detail
VRF-Name: TN_D:VRF_Stretch, VRF-ID: 46, State: Up
VPNID: unknown
RD: 1101:2850817
Max Routes: 0 Mid-Threshold: 0
Table-ID: 0x8000002e, AF: IPv6, Fwd-ID: 0x8000002e, State: Up
Table-ID: 0x0000002e, AF: IPv4, Fwd-ID: 0x0000002e, State: Up
```

```
Site1_Leaf1# show bgp vpnv4 unicast 91.0.0.1 vrf overlay-1
BGP routing table information for VRF overlay-1, address family VPNv4 Unicast
Route Distinguisher: 1101:2850817 (VRF TN_D:VRF_Stretch)
BGP routing table entry for 91.0.0.1/32, version 17 dest ptr 0xadeda550
Paths: (1 available, best #1)
Flags: (0x08001a 00000000) on xmit-list, is in urib, is best urib route, is in HW
      vpn: version 357, (0x100002) on xmit-list
Multipath: eBGP iBGP
  Advertised path-id 1, VPN AF advertised path-id 1
  Path type: internal 0xc0000018 0x80040 ref 56506 adv path ref 2, path is valid, is best path,
remote site path
      Imported from 1101:36241410:91.0.0.1/32
AS-Path: NONE, path sourced internal to AS
  192.168.100.225 (metric 64) from 10.0.80.65 (192.168.10.13)
  Origin incomplete, MED 2, localpref 100, weight 0
  Received label 0
  Received path-id 1
  Extcommunity:
    RT:65001:36241410
    SOO:65001:50331631
    COST:pre-bestpath:166:2684354560
    COST:pre-bestpath:168:3221225472
    VNID:2686978
    COST:pre-bestpath:162:110
  Originator: 10.0.72.64 Cluster list: 192.168.10.13192.168.11.13 <<<<
'10.0.72.64'='Site2_Leaf1' , '192.168.10.13'='Site1_Spine' , '192.168.11.13'='Site2_Spine'
  VRF advertise information:
  Path-id 1 not advertised to any peer
  VPN AF advertise information:
  Path-id 1 not advertised to any peer
<snip..>
```

```
Site1_Leaf1# show bgp vpnv4 unicast 91.0.0.1 vrf TN_D:VRF_Stretch
BGP routing table information for VRF overlay-1, address family VPNv4 Unicast
Route Distinguisher: 1101:2850817 (VRF TN_D:VRF_Stretch)
BGP routing table entry for 91.0.0.1/32, version 17 dest ptr 0xadeda550
Paths: (1 available, best #1)
Flags: (0x08001a 00000000) on xmit-list, is in urib, is best urib route, is in HW
      vpn: version 357, (0x100002) on xmit-listMultipath: eBGP iBGP
  Advertised path-id 1, VPN AF advertised path-id 1
  Path type: internal 0xc0000018 0x80040 ref 56506 adv path ref 2, path is valid, is best path,
remote site path
      Imported from 1101:36241410:91.0.0.1/32
AS-Path: NONE, path sourced internal to AS
  192.168.100.225 (metric 64) from 10.0.80.65 (192.168.10.13)
  Origin incomplete, MED 2, localpref 100, weight 0
```

```

Received label 0
Received path-id 1
Extcommunity:
  RT:65001:36241410
  SOO:65001:50331631
  COST:pre-bestpath:166:2684354560
  COST:pre-bestpath:168:3221225472
  VNID:2686978
  COST:pre-bestpath:162:110
Originator: 10.0.72.64 Cluster list: 192.168.10.13 192.168.11.13
VRF advertise information:
Path-id 1 not advertised to any peer
VPN AF advertise information:
Path-id 1 not advertised to any peer

```

因此，「Site1_Leaf1」具有子網91.0.0.1/32的路由條目，帶有下一跳「Site2_Leaf1」ETEP地址192.168.100.225。

```

Site1_Leaf1# show ip route 91.0.0.1 vrf TN_D:VRF_Stretch
IP Route Table for VRF "TN_D:VRF_Stretch"
 '*' denotes best ucast next-hop
  '*' denotes best mcast next-hop
  '[x/y]' denotes [preference/metric]
  '%' in via output denotes VRF
91.0.0.1/32, ubest/mbest: 1/0
  *via 192.168.100.225%overlay-1, [200/2], 5d23h, bgp-65001, internal, tag 65001 <<<< Note
that next hope is External TEP pool (ETEP) ip address of Site-B.
  recursive next hop: 192.168.100.225/32%overlay-1

```

站點A主幹向「Site2_Spine」mosp-ETEP的BGP鄰居IP地址新增路由對映。因此，如果您考慮流量傳輸，當站點A端點與外部IP位址通話時，封包可以使用來源封裝為「Site1_Leaf1」TEP位址，且目的地為「Site2_Leaf」IP位址192.168.100.225的ETEP位址。驗證ELAM(Site1_Spine)

```

Site1_Spine# vsh_lc
module-1# debug platform internal roc elam asic 0
module-1(DBG-elam)# trigger reset
module-1(DBG-elam)# trigger init in-select 14 out-select 1
module-1(DBG-elam-insel14)# set inner ipv4 src_ip 90.0.0.10 dst_ip 91.0.0.1 next-protocol 1
module-1(DBG-elam-insel14)# start
module-1(DBG-elam-insel14)# status
  ELAM STATUS
  =====
  Asic 0 Slice 0 Status Armed
  Asic 0 Slice 1 Status Armed
  Asic 0 Slice 2 Status Armed
  Asic 0 Slice 3 Status Armed

```

```

pod2-n9k# ping 91.0.0.1 vrf HOST_A source 90.0.0.10
PING 91.0.0.1 (91.0.0.1) from 90.0.0.10: 56 data bytes
64 bytes from 91.0.0.1: icmp_seq=0 ttl=252 time=1.015 ms
64 bytes from 91.0.0.1: icmp_seq=1 ttl=252 time=0.852 ms
64 bytes from 91.0.0.1: icmp_seq=2 ttl=252 time=0.859 ms
64 bytes from 91.0.0.1: icmp_seq=3 ttl=252 time=0.818 ms
64 bytes from 91.0.0.1: icmp_seq=4 ttl=252 time=0.778 ms
--- 91.0.0.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.778/0.864/1.015 ms

```

Site1_Spine ELAM被觸發。Ereport確認資料包封裝了站點A枝葉TEP IP地址和目的地的TEP地址，指向Site2_Leaf1 ETEP地址。

```

module-1(DBG-elam-insel14)# status
  ELAM STATUS
  =====

```

```
Asic 0 Slice 0 Status Armed
Asic 0 Slice 1 Status Armed
Asic 0 Slice 2 Status Triggered
Asic 0 Slice 3 Status Armed
module-1(DBG-elam-insell14)# ereport
Python available. Continue ELAM decode with LC Pkg
ELAM REPORT
```

```
-----
-----
Outer L3 Header
-----
```

```
-----
L3 Type           : IPv4
DSCP              : 0
Don't Fragment Bit : 0x0
TTL              : 32
IP Protocol Number : UDP
Destination IP    : 192.168.100.225    <<<'Site2_Leaf1' ETEP address
Source IP        : 10.0.80.64         <<<'Site1_Leaf1' TEP address
-----
```

```
-----
-----
Inner L3 Header
-----
```

```
-----
L3 Type           : IPv4
DSCP              : 0
Don't Fragment Bit : 0x0
TTL              : 254
IP Protocol Number : ICMP
Destination IP    : 91.0.0.1
Source IP        : 90.0.0.10
-----
```

站點1_骨幹驗證路由對映當站點A主幹收到資料包時，它可以重定向到「站點2_Leaf1」ETEP地址，而不是查詢coop或route條目。（當您在Site-B上有intersite-L3out時，Site-A主幹會建立一個名為「infra-intersite-l3out」的路由對映，將流量重定向到Site2_Leaf1的ETEP並從L3out退出。）

```
Site1_Spine# show bgp vpnv4 unicast neighbors 192.168.11.13 vrf overlay-1
BGP neighbor is 192.168.11.13, remote AS 65001, ibgp link, Peer index 4
  BGP version 4, remote router ID 192.168.11.13
  BGP state = Established, up for 10w4d
  Using loopback12 as update source for this peer
  Last read 00:00:03, hold time = 180, keepalive interval is 60 seconds
  Last written 00:00:03, keepalive timer expiry due 00:00:56
  Received 109631 messages, 0 notifications, 0 bytes in queue
  Sent 109278 messages, 0 notifications, 0 bytes in queue
  Connections established 1, dropped 0
  Last reset by us never, due to No error
  Last reset by peer never, due to No error
  Neighbor capabilities:
  Dynamic capability: advertised (mp, refresh, gr) received (mp, refresh, gr)
  Dynamic capability (old): advertised received
  Route refresh capability (new): advertised received
  Route refresh capability (old): advertised received
  4-Byte AS capability: advertised received
  Address family VPNv4 Unicast: advertised received
  Address family VPNv6 Unicast: advertised received
  Address family L2VPN EVPN: advertised received
  Graceful Restart capability: advertised (GR helper) received (GR helper)
  Graceful Restart Parameters:
  Address families advertised to peer:
```

```

Address families received from peer:
Forwarding state preserved by peer for:
Restart time advertised by peer: 0 seconds
Additional Paths capability: advertised received
Additional Paths Capability Parameters:
Send capability advertised to Peer for AF:
  L2VPN EVPN
Receive capability advertised to Peer for AF:
  L2VPN EVPN
Send capability received from Peer for AF:
  L2VPN EVPN
Receive capability received from Peer for AF:
  L2VPN EVPN
Additional Paths Capability Parameters for next session:
[E] - Enable [D] - Disable
Send Capability state for AF:
  VPNv4 Unicast[E] VPNv6 Unicast[E]
Receive Capability state for AF:
  VPNv4 Unicast[E] VPNv6 Unicast[E]
Extended Next Hop Encoding Capability: advertised received
Receive IPv6 next hop encoding Capability for AF:
  IPv4 Unicast
Message statistics:

```

| | Sent | Rcvd |
|-----------------|---------|---------|
| Opens: | 1 | 1 |
| Notifications: | 0 | 0 |
| Updates: | 1960 | 2317 |
| Keepalives: | 107108 | 107088 |
| Route Refresh: | 105 | 123 |
| Capability: | 104 | 102 |
| Total: | 109278 | 109631 |
| Total bytes: | 2230365 | 2260031 |
| Bytes in queue: | 0 | 0 |

```

For address family: VPNv4 Unicast
BGP table version 533, neighbor version 533
3 accepted paths consume 360 bytes of memory
3 sent paths
0 denied paths
Community attribute sent to this neighbor
Extended community attribute sent to this neighbor
Third-party Nexthop will not be computed.
Outbound route-map configured is infra-intersite-l3out, handle obtained <<<< route-map to
redirect traffic from Site-A to Site-B 'Site2_Leaf1' L3out
For address family: VPNv6 Unicast
BGP table version 241, neighbor version 241
0 accepted paths consume 0 bytes of memory
0 sent paths
0 denied paths
Community attribute sent to this neighbor
Extended community attribute sent to this neighbor
Third-party Nexthop will not be computed.
Outbound route-map configured is infra-intersite-l3out, handle obtained
<snip...> Site1_Spine# show route-map infra-intersite-l3out
route-map infra-intersite-l3out, permit, sequence 1
Match clauses:
  ip next-hop prefix-lists: IPv4-Node-entry-102
  ipv6 next-hop prefix-lists: IPv6-Node-entry-102
Set clauses:
  ip next-hop 192.168.200.226
route-map infra-intersite-l3out, permit, sequence 2 <<<< This route-map match if destination
IP of packet 'Site1_Spine' TEP address then send to 'Site2_Leaf1' ETEP address.
Match clauses:
  ip next-hop prefix-lists: IPv4-Node-entry-1101
  ipv6 next-hop prefix-lists: IPv6-Node-entry-1101

```



```
Set clauses:
  ip next-hop 192.168.200.225
route-map infra-intersite-l3out, deny, sequence 999
Match clauses:
  ip next-hop prefix-lists: infra_prefix_local_pteps_inexact
Set clauses:
route-map infra-intersite-l3out, permit, sequence 1000
Match clauses:
Set clauses:
  ip next-hop unchanged

Site1_Spine# show ip prefix-list IPv4-Node-entry-1101
ip prefix-list IPv4-Node-entry-1101: 1 entries
seq 1 permit 10.0.80.64/32 <<
```