

Multi-Site Domain for VXLAN BGP EVPN Fabrics

- Multi-Site Domain for VXLAN BGP EVPN Fabrics , on page 1
- MSD and Member Fabric Process Flow, on page 2
- Creating the VXLAN EVPN Multi-Site and Associating Member Fabrics, on page 5
- Creating and Deploying Networks and VRFs in an MSD Fabric, on page 11
- Moving a Standalone Fabric (With Existing Networks and VRFs) to an MSD Fabric , on page 12
- Support for CloudSec in Multi-Site Deployment, on page 13

Multi-Site Domain for VXLAN BGP EVPN Fabrics

A Multi-Site Domain (MSD) is a multifabric container that is created to manage multiple member fabrics. An MSD is a single point of control for definition of overlay networks and VRFs that are shared across member fabrics. When you move fabrics (that are designated to be part of the multifabric overlay network domain) under the MSD as member fabrics, the member fabrics share the networks and VRFs created at the MSD-level. This way, you can consistently provision network and VRFs for different fabrics, at one go. It significantly reduces the time and complexity involving multiple fabric provisionings.

As server networks and VRFs are shared across the member fabrics (as one stretched network), the new networks and VRFs provisioning function is provided at the MSD fabric level. Any new network and VRF creation is only allowed for the MSD. All member fabrics inherit any new network and VRF created for the MSD.

The topology view for the MSD fabric displays all member fabrics, and how they are connected to each other, in one view. You can deploy overlay networks (and VRFs) on member fabrics from a single topology deployment screen, instead of visiting each member fabric deployment screen separately and deploying.

Note

- The VXLAN OAM feature in Cisco NDFC is only supported on a single fabric or site.
- After you unpair a BGW vPC, perform a **Recalculate Config** and **Deploy Config** on the member fabric followed by a **Recalculate Config** and **Deploy Config** of the MSD fabric.

A few fabric-specific terms:

- **Standalone fabric** – A fabric that is not part of an MSD is referred as a standalone fabric from the MSD perspective. Before the MSD concept, all fabrics were considered standalone, though two or more such fabrics can be connected with each other.

- **Member fabrics** – Fabrics that are part of an MSD are called *member* fabrics or *members*. Create a standalone fabric (of the type *Data Center VXLAN EVPN*) first and then move it within an MSD as a member fabric.

When a standalone fabric is added to the MSD, the following actions take place:

- The standalone fabric's relevant attributes and the network and VRF definitions are checked against that of the MSD. If there is a *conflict*, then the standalone fabric addition to the MSD fails. If there are no conflicts, then the standalone fabric becomes a member fabric for the MSD. If there is a conflict, the exact conflicts are logged in the pending errors log for the MSD fabric. You can remedy the conflicts and then attempt to add the standalone fabric to the MSD again.
- All the VRFs and networks definitions from the standalone fabric that do not have presence in the MSD are copied over to the MSD and in turn inherited to each of its other existing member fabrics.
- The VRFs and networks (and their definitions) from the MSD (such as the MSD's VRF, and L2 and L3 VNI parameters that *do not* have presence in the standalone fabric) are inherited into the standalone fabric that just became a member.

Fabric and Switch Instance Variables

While the MSD provisions a global range of network and VRF values, some parameters are fabric-specific and some parameters are switch-specific. The parameters are called *fabric instance* and *switch instance* variables.

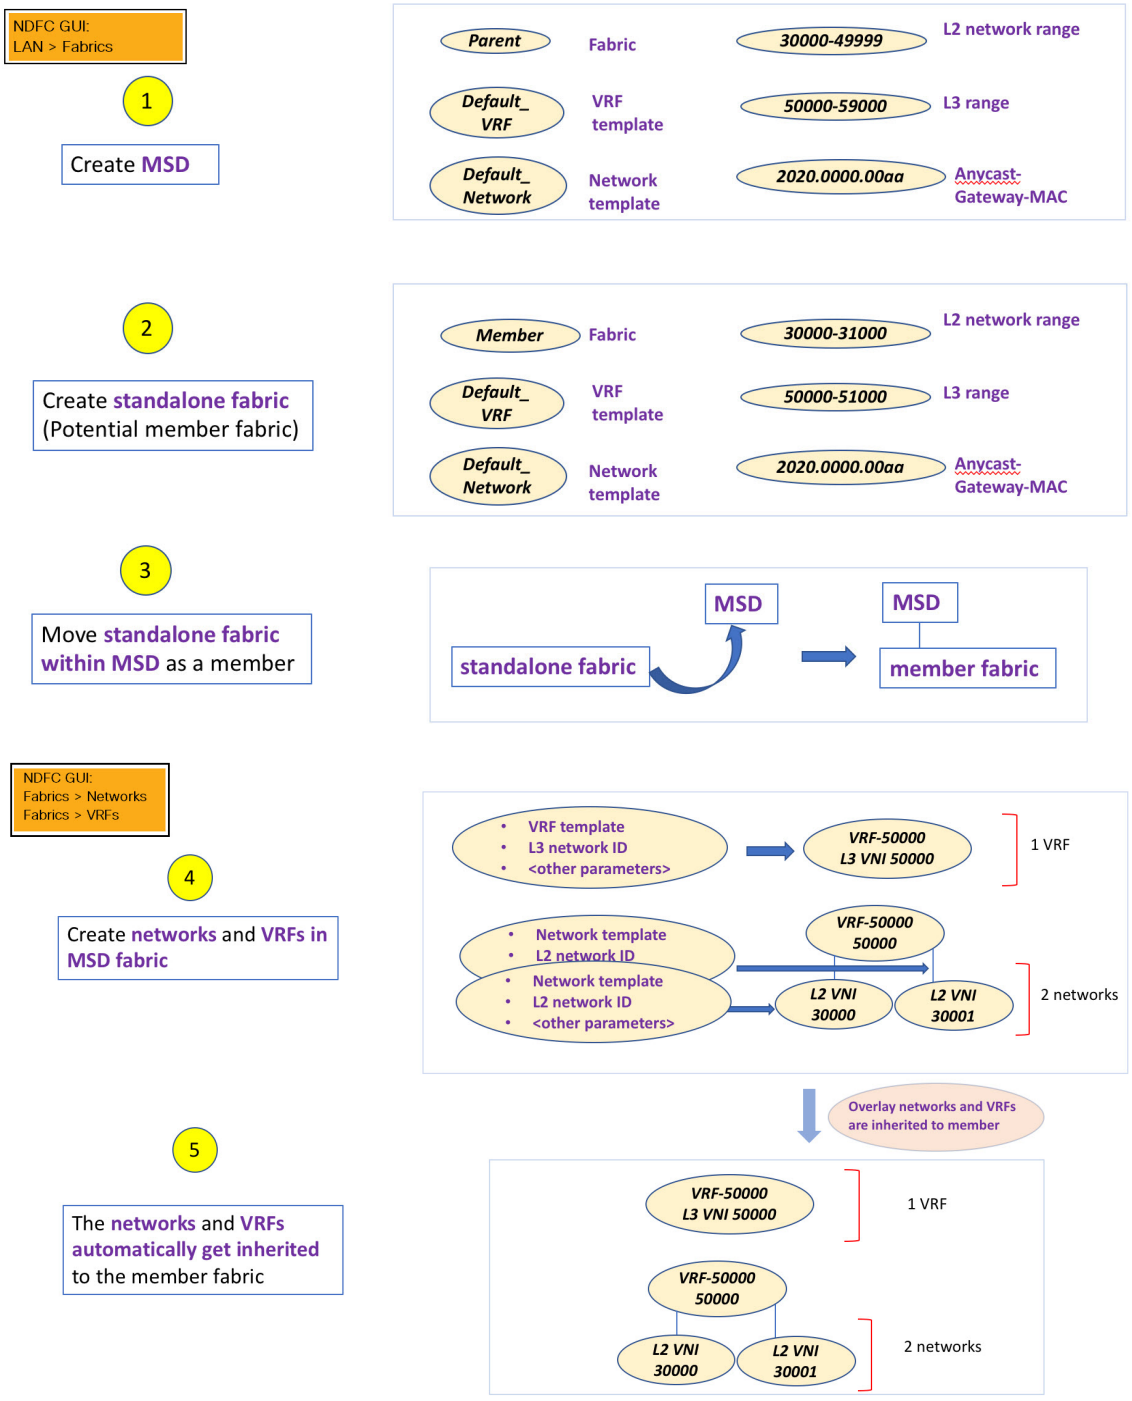
Fabric instance values can only be edited or updated in the fabric context from the VRFs and Networks window. Double click on the appropriate fabric to view **Fabric Overview** and choose **Networks** or **VRFs** tab. Some of the examples of fabric instance variables are BGP ASN, Multicast group per network or VRF, etc. For information about editing multicast group address, see Creating Networks in the MSD Fabric, on page 11.

Switch instance values can be edited on deployment of the network on the switch. For example, *VLAN ID*.

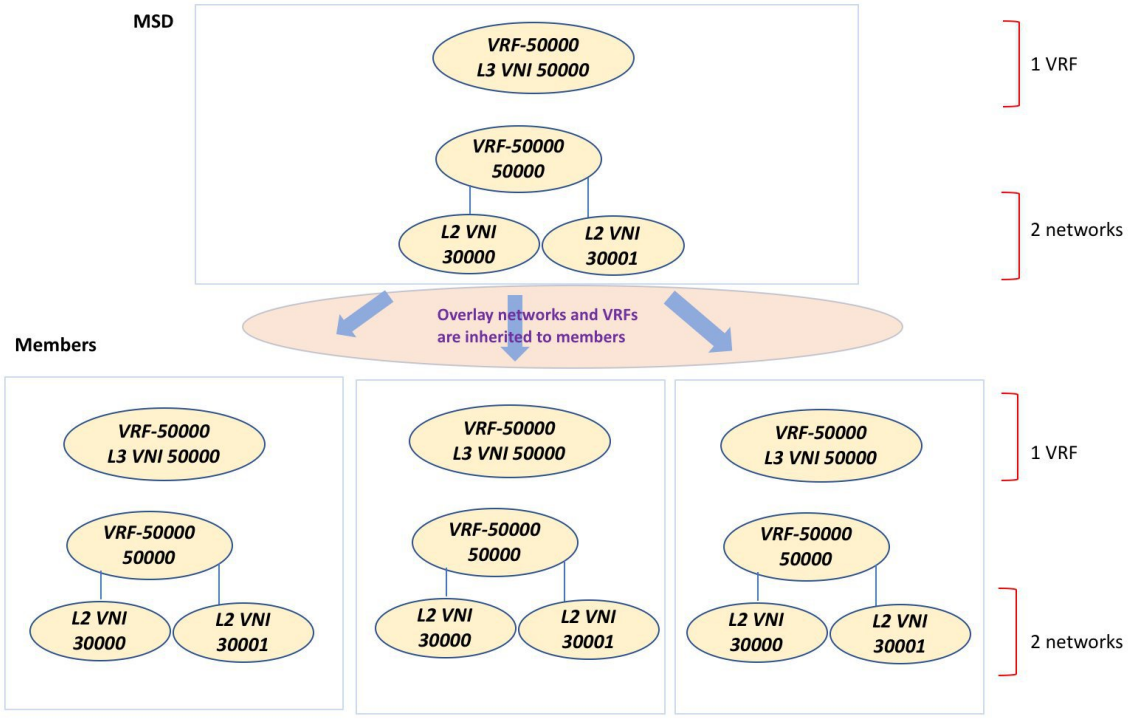
MSD and Member Fabric Process Flow

An MSD has multiple sites (and hence, multiple member fabrics under an MSD). VRFs and networks are created for the MSD and get inherited by the member fabrics. For example, VRF-50000 (and L3 network with ID 50000), and L2 networks with IDs 30000 and 30001 are created for the MSD, in one go.

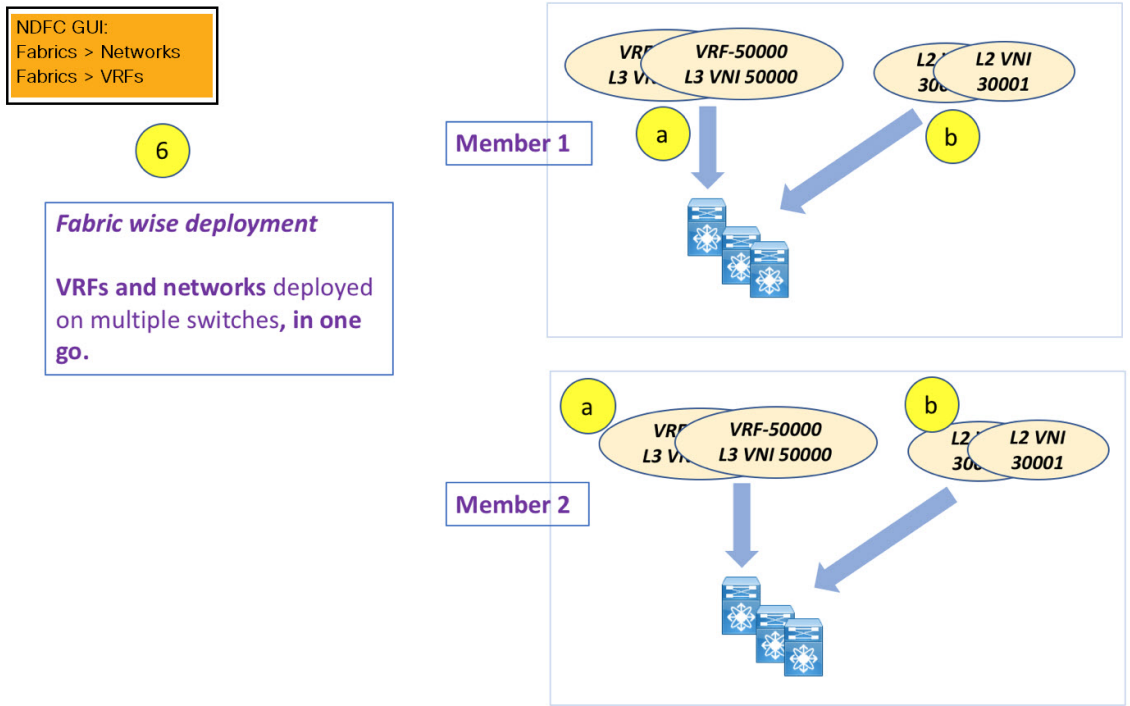
A high-level flow chart of the MSD and member fabric creation and MSD-to-member fabric inheritance process:



The sample flow explained the inheritance from the MSD to one member. An MSD has multiple sites (and hence, multiple member fabrics under an MSD). A sample flow from an MSD to multiple members:



In this example, VRF-50000 (and L3 network with ID 50000), and L2 networks with IDs 30000 and 30001 are created in one go. Networks and VRFs are deployed on the member fabric switches, one after another, as depicted in the image.



You can provision overlay networks through a single MSD deployment screen.



Note If you move a standalone fabric with existing networks and VRFs to an MSD, NDFC does appropriate validation. This is explained in detail in an upcoming section.

Upcoming sections in the document explain the following:

- Creation of an MSD fabric.
- Creation of a standalone fabric (as a potential member) and its movement under the MSD as a member.
- Creation of networks and VRFs in the MSD and their inheritance to the member fabrics.
- Deployment of networks and VRFs from the MSD and member fabric topology views.
- Other scenarios for fabric movement:
  - Standalone fabric with existing networks and VRFs to an MSD fabric.
  - Member fabric from one MSD to another.

Creating the VXLAN EVPN Multi-Site and Associating Member Fabrics

The process is explained in two steps:

1. Create the VXLAN EVPN Multi-Site fabric.
2. Create a new standalone fabric and move it under the VXLAN EVPN Multi-Site fabric as a member fabric.

Creating the VXLAN EVPN Multi-Site fabric

1. From **Actions** drop-down list, choose **Create Fabric**.

   The **Create Fabric** window appears.

2. Enter a unique name for the Fabric.

   Click on **Choose Template** to pick a template for your fabric.

   A list of all available fabric templates are listed.

3. From the available list of Fabric templates, choose the **VXLAN EVPN Multi-Site** fabric template.

   Click **Select**.

   Enter the necessary field values to create a Fabric.

   The tabs and their fields in the screen are explained in the subsequent points. The overlay and underlay network parameters are included in these tabs.

4. In the **General Parameters** tab, all fields are autopopulated with data. The fields consist of the Layer 2 and Layer 3 VXLAN segment identifier range, the default network and VRF templates, and the anycast gateway MAC address. Update the relevant fields as needed.

   Layer 2 VXLAN VNI Range - Layer 2 VXLAN segment identifier range.

   Layer 3 VXLAN VNI Range - Layer 3 VXLAN segment identifier range.

   VRF Template - Default VRF template.

   Network Template - Default network template.

   VRF Extension Template - Default VRF extension template.

   Network Extension Template - Default network extension template.

   Anycast-Gateway-MAC - Anycast gateway MAC address.

   Multisite Routing Loopback Id – The multisite routing loopback ID is populated in this field.

   ToR Auto-deploy Flag - Select this check box to enable automatic deployment of the networks and VRFs in the Easy Fabric to the ToR switches in the External Fabric when you click **Recalculate and Deploy** in the VXLAN EVPN Multi-Site fabric.

5. Click the **DCI** tab.

   The fields are:

   Multi-Site Overlay IFC Deploy Method – Choose how you will connect the data centers through the BGW, manually, in a back-to-back fashion or through a route server.

   Multi-Site Route Server List – Specify the IP addresses of the route server. If you specify more than one, separate the IP addresses by a comma.

   Multi-Site Route Server BGP ASN List – Specify the BGP AS Number of the router server. If you specify more than one route server, separate the AS Numbers by a comma.

   Multi-Site Underlay IFC Auto Deployment Flag - Check the check box to enable auto configuration. Uncheck the check box for manual configuration.

   Delay Restore Time - Specifies the Multi-Site underlay and overlay control planes convergence time. The minimum value is 30 seconds and the maximum value is 1000 seconds.

   Multi-Site CloudSec – Enables CloudSec configurations on border gateways. If you enable this field, the remaining three fields for CloudSec are editable. For more information, see Support for CloudSec in Multi-Site Deployment, on page 13.

   Enable Multi-Site eBGP Password - Enables eBGP password for Multi-Site underlay/overlay IFCs.

   eBGP Password - Specifies the encrypted eBGP Password Hex String.

   eBGP Authentication Key Encryption Type - Specifies the BGP key encryption type. It is **3** for 3DES and **7** for Cisco.

6. Click the **Resources** tab.

   MultiSite Routing Loopback IP Range – Specify the Multi-Site loopback IP address range used for the EVPN Multi-Site function.

   A unique loopback IP address is assigned from this range to each member fabric because each member site must have a Multi-site Routing Loopback IP address assigned for overlay network reachability. The per-fabric loopback IP address is assigned on all the BGWs in a specific member fabric.

DCI Subnet IP Range and **Subnet Target Mask** – Specify the Data Center Interconnect (DCI) subnet IP address and mask.

7. Click the **Configuration Backup** tab.

   Scheduled Fabric Backup: Check the check box to enable a daily backup. This backup tracks changes in running configurations on the fabric devices that are not tracked by configuration compliance.

   Scheduled Time: Specify the scheduled backup time in a 24-hour format. This field is enabled if you check the Scheduled Fabric Backup check box.

   Select both the check boxes to enable both back up processes.

   The backup process is initiated after you click Save.

8. Click **Save**.

   A message appears briefly at the bottom right part of the screen, indicating that you have created a new VXLAN EVPN Multi-Site fabric. After fabric creation, the fabric page comes up.

   When a new VXLAN EVPN Multi-Site fabric is created, the newly created VXLAN EVPN Multi-Site fabric instance appears in the Fabrics table.

   The VXLAN EVPN Multi-Site fabric is displayed, and it contains the member fabric names as a branch. When no member fabric is created, it is displayed as a standalone fabric.

The steps for creation of the VXLAN EVPN Multi-Site fabric and moving member fabrics under it are:

1. Create the VXLAN EVPN Multi-Site fabric.
2. **Create a new standalone fabric and move it under the VXLAN EVPN Multi-Site fabric as a member fabric.**

Step 1 is completed. Step 2 is explained in the next section.

Creating and Moving a New Fabric Under the VXLAN EVPN Multi-Site Fabric as a Member

A new fabric is created as a standalone fabric. After you create a new fabric, you can move it under the VXLAN EVPN Multi-Site fabric as a member. As a best practice, when you create a new fabric that is a potential member fabric (of the VXLAN EVPN Multi-Site fabric), do not add networks and VRFs to the fabric. Move the fabric under the VXLAN EVPN Multi-Site fabric and then add networks and VRFs for the VXLAN EVPN Multi-Site fabric. That way, there will not be any need for validation (or conflict resolution) between the member and VXLAN EVPN Multi-Site fabric network and VRF parameters.

New fabric creation is explained in the Easy Fabric creation process. In the VXLAN EVPN Multi-Site fabric document, fabric movement is covered. However, some pointers about a standalone (potential member) fabric:

The values under the **Resources** tab are automatically generated. The VXLAN VNI ID ranges (in the L2 Segment ID Range and L3 Partition ID Range fields) allocated for new network and VRF creation are values from the VXLAN EVPN Multi-Site fabric segment ID range. If you want to update the VXLAN VNI ranges or the VRF and Network VLAN ranges, ensure the following:

- If you update a range of values, ensure that it does not overlap with other ranges.
- You must update one range of values at a time. If you want to update more than one range of values, do it in separate instances. For example, if you want to update L2 and L3 ranges, you should do the following:

  1. Update the L2 range and click **Save**.

2. Click the **Edit Fabric** option again, update the L3 range and click **Save**.

Ensure that the **Anycast Gateway MAC**, the **Network Template** and the **VRF Template** field values are the same as the VXLAN EVPN Multi-Site fabric. Else, member fabric movement to the VXLAN EVPN Multi-Site fabric fail.

Other pointers:

- The member fabric should have a Site ID configured and the Site ID must be unique among the members.
- The BGP AS number should be unique for a member fabric.
- The underlay subnet range for loopback0 should be unique.
- The underlay subnet range for loopback1 should be unique.

After you click **Save**, a note appears at the bottom right part of the screen indicating that the fabric is created. When a fabric is created, the fabric page comes up. The fabric name appears in the list of fabrics.

Moving the Member1 Fabric Under VXLAN EVPN Multi-Site-Parent-Fabric

You should go to the VXLAN EVPN Multi-Site fabric Overview to associate a member fabric under it.

1. Double click on the VXLAN EVPN Multi-Site fabric to view the **Fabric Overview** screen.
2. On the **Child Fabrics** tab, click .

   You can also click on **Fabric Overview > Actions > Add Child Fabrics** to add member fabrics to the VXLAN EVPN Multi-Site fabric.

   A list of child fabrics that are not part of any VXLAN EVPN Multi-Site fabric appears. Member fabrics of other VXLAN EVPN Multi-Site fabric container fabrics are not displayed here.
3. As *Member1* fabric is to be associated with the VXLAN EVPN Multi-Site fabric, select the **Member1** fabric and click **Select**.
4. Select the Fabric and click **Select**.

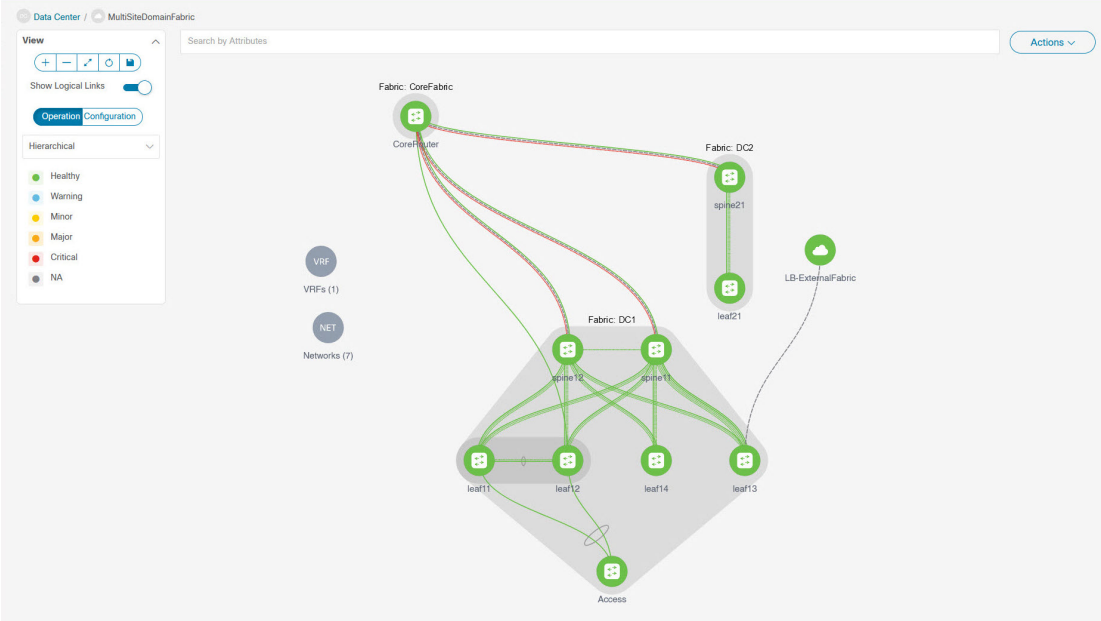
   You can see that *Member1* is now added to VXLAN EVPN Multi-Site fabric and is displayed in the **Child Fabrics** in the Fabrics list table.

VXLAN EVPN Multi-Site Fabric Topology View Pointers

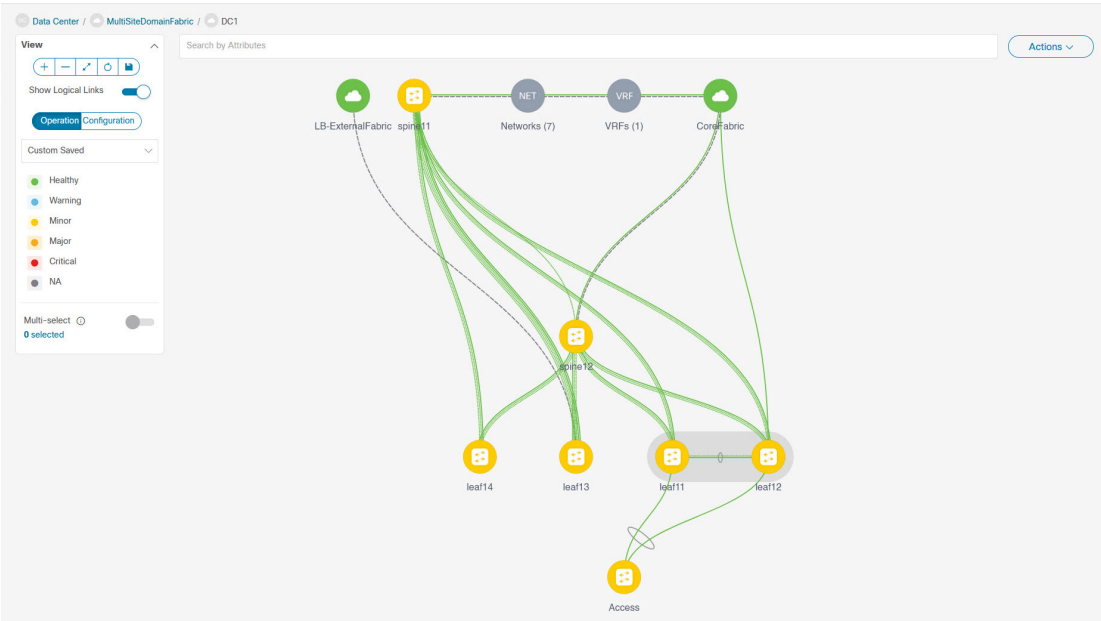
The Topology tab displays the configured VXLAN EVPN Multi-Site fabrics and its child fabrics.

- **VXLAN EVPN Multi-Site fabric topology view** – VXLAN EVPN Multi-Site fabric and their member fabrics displayed. A boundary defines each member fabric. All fabric devices of the fabric are confined to the boundary.

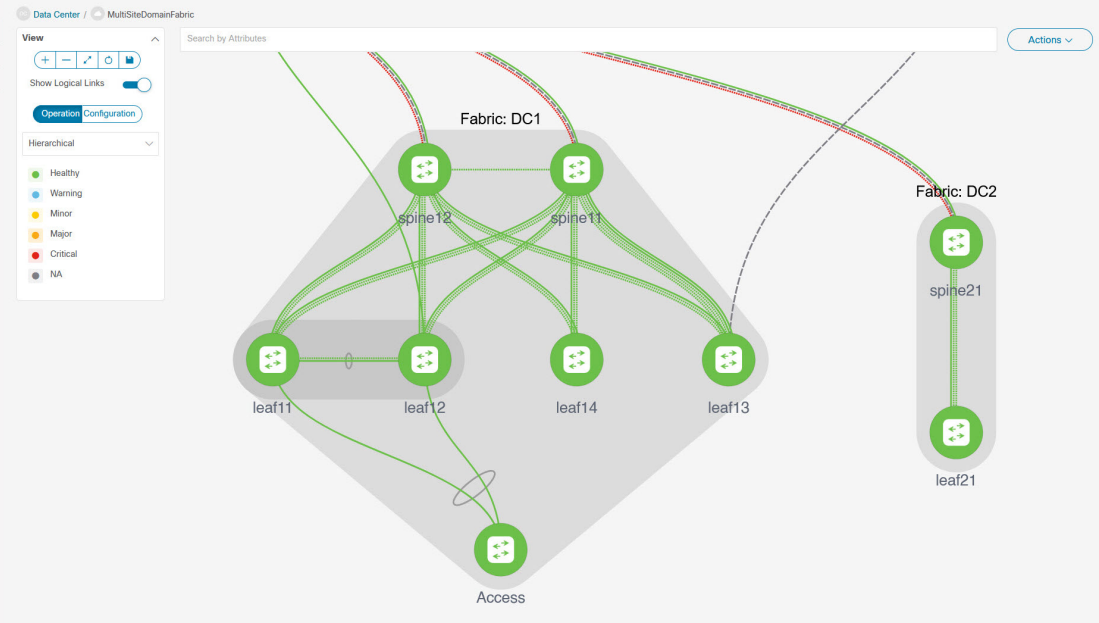
  Double click on the member fabric to view further elements.



- **Member fabric topology view** – A member fabric and its switches are displayed. In addition, the connected external fabric is displayed.



- A boundary defines a standalone VXLAN fabric, and each member fabric in the VXLAN EVPN Multi-Site fabric. A fabric's devices are confined to the fabric boundary. You can move a switch icon by dragging it. For a better user experience, in addition to switches, NDFC allows you to move an entire fabric. To move a fabric, place the cursor within the fabric boundary (but not on a switch icon), and drag it in the desired direction.



Configuring External Fabric Routers for VXLAN EVPN Child Fabric Connectivity

To successfully pair external fabric routers with VXLAN EVPN child fabrics using the **Fabric-Group Settings** for auto-deploy configurations, the following prerequisites and configurations must be applied to the external fabric routers:

- **Role of the device:** The device in the external fabric must be set as a core-router.
- **Loopback0 Interface Configuration:** The loopback0 interface must be configured with an IP address and tag:

  ```
  interface loopback0

  ip address <ip>/<mask> tag <tag>
  ```

- **BGP Configuration:**BGP must be configured on the external fabric router:

  ```
  feature bgp

  router bgp <ASN>
  router-id <ip>
  address-family ipv4 unicast
  ```

  You can use the ASN, along with the router ID, when defining the auto-deploy configuration in the **Fabric-Group Settings** within NDFC.

  These configurations are crucial for establishing the necessary control plane and routing capabilities, enabling seamless integration and auto-deployment of IFCs between the external fabric and the VXLAN EVPN child fabrics.

Adding and Editing Links

To add a link, choose **Actions** > **More** > **Add Link**. To edit a link, choose **Actions** > **More** > **Edit Link**.

To know how to add links between border switches of different fabrics (inter-fabric links) or between switches in the same fabric (intra-fabric links), refer the **Fabric Links** topic.

Creating and Deploying Networks and VRFs in an MSD Fabric

In standalone fabrics, networks and VRFs are created for each fabric. In an MSD fabric, networks and VRFs should be created at the MSD fabric level. The networks and VRFs are inherited by all the member networks. You cannot create or delete networks and VRFs for member fabrics. However, you can edit them.

For example, consider an MSD fabric with two member fabrics. If you create three networks in the MSD fabric, then all three networks will automatically be available for deployment in both the member fabrics.

Though member fabrics inherit the MSD fabric's networks and VRFs, you have to deploy the networks and VRFs distinctly, for each fabric.

A deployment view is introduced for the MSD, in addition to the per-fabric deployment view. In this view, you can view and provision overlay networks for all member fabrics within the MSD, at once. However, you still have to apply and save network and VRF configurations distinctly, for each fabric.



Note Networks and VRFs are the common identifiers (represented across member fabrics) that servers (or end hosts) are grouped under so that traffic can be sent between the end hosts based on the network and VRF IDs, whether they reside in the same or different fabrics. Since they have common representation across member fabrics, networks and VRFs can be provisioned at one go. As the switches in different fabrics are physically and logically distinct, you have to deploy the same networks and VRFs separately for each fabric.

For example, if you create networks 30000 and 30001 for an MSD that contains two member fabrics, the networks are automatically created for the member fabrics and are available for deployment.

You can deploy 30000 and 30001 on the border devices of all member fabrics through a single (MSD fabric) deployment screen. Prior to this, you had to access the first member fabric deployment screen, deploy 30000 and 300001 on the fabric's border devices, and then access the second member fabric deployment screen and deploy again.

Networks and VRFs are created in the MSD and deployed in the member fabrics. The steps are explained below:

1. Create networks and VRFs in the MSD fabric.
2. Deploy the networks and VRFs in the member fabric devices.

Creating Networks in the MSD Fabric

Some guidelines and pointers:

- In the MSD fabric level, if the **Enable L3 Gateway on Border** check box is selected and you upgrade the NDFC service, then it is automatically removed from the MSD fabric level during upgrade.
- You can edit the **Network Profile** part (**General** and **Advanced** tabs) of the MSD fabric network.
- An MSD can contain multiple fabrics. These fabrics forward BUM traffic via Multicast or Ingress replication. Even if all the fabrics use multicast for BUM traffic, the multicast groups within these fabrics need not be the same.

- When you create a network in MSD, it is inherited by all the member fabrics. However, the multicast group address is a fabric instance variable. To edit the multicast group address, you need to navigate to the member fabric and edit the network. For more information about the **Multicast Group Address** field, see *Creating Networks for the Standalone Fabric*.
- You can only delete networks from the MSD fabric, and not member fabrics. You must undeploy the networks on the respective fabric devices before deletion.
- When you delete networks from the MSD fabric, the networks are automatically removed from the member fabrics too.

See Creating Networks for the Standalone Fabric.

Creating VRFs in the MSD Fabric

You cannot delete VRFs at the member fabric level. Delete VRFs in the MSD fabric. The deleted VRFs are automatically removed from all member fabrics.

See Creating VRF.

Deleting Networks and VRFs in the MSD and Member Fabrics

You can only delete networks from the MSD fabric, and not member fabrics. To delete networks and corresponding VRFs in the MSD fabric, follow this order:

1. Undeploy the networks on the respective fabric devices before deletion.
2. Delete the networks from the MSD fabric.
3. Undeploy the VRFs on the respective fabric devices before deletion.
4. Delete the VRFs from the MSD fabric. You can delete multiple VRF instances at once.



Note When you delete VRFs from the MSD fabric, they are automatically removed from the member fabrics too.

Moving a Standalone Fabric (With Existing Networks and VRFs) to an MSD Fabric

If you move a standalone fabric with existing networks and VRFs to an MSD fabric as a member, ensure that common networks (that is, L2 VNI and L3 VNI information), anycast gateway MAC, and VRF and network templates are the same across the fabric and the MSD. NDFC validates the standalone fabric (network and VRF information) against the (network and VRF information) of the MSD fabric to avoid conflict entries. An example of conflict entries is two common network names with a different network ID. After validation and there is no conflict, the standalone fabric is moved to the MSD fabric as a member fabric. Details:

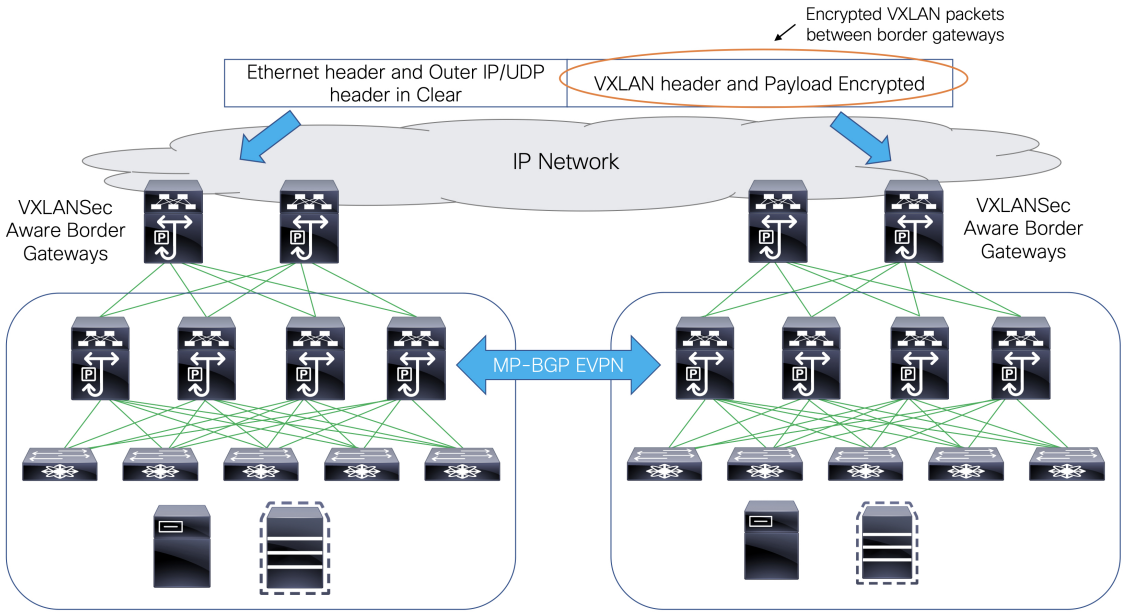
- The MSD fabric inherits the networks and VRFs of the standalone fabric that do not exist in the MSD fabric. These networks and VRFs are in turn inherited by the member fabrics.
- The newly created member fabric inherits the networks and VRFs of the MSD fabric (that do not exist in the newly created member fabric).

- If there are conflicts between the standalone and MSD fabrics, validation ensures that an error message is displayed. You can move the standalone fabric to MSD again after updating. If the move is successful, a message comes up at the top of the page indicating that the move is successful.

If you move back a member fabric to standalone status, then the networks and VRFs remain as they are, but they remain relevant as in an independent fabric, outside the purview of an MSD fabric.

Support for CloudSec in Multi-Site Deployment

CloudSec feature allows secured data center interconnect in a multi-site deployment by supporting source-to-destination packet encryption between border gateway devices in different fabrics.



CloudSec feature is supported on Cisco Nexus 9000 Series FX2 platform with Cisco NX-OS Release 9.3(5) or later. The border gateways, border gateway spines, and border gateway superspines that are FX2 platforms, and run Cisco NX-OS Release 9.3(5) or later are referred as CloudSec capable switches.

You can enable CloudSec while creating an MSD fabric.

Note The CloudSec session is point to point over DCI between border gateways (BGWs) on two different sites. All communication between sites uses Multi-Site PIP instead of VIP. Enabling CloudSec requires a switch from VIP to PIP, which could cause traffic disruption for data flowing between sites. Therefore, it is recommended to enable or disable CloudSec during a maintenance window.

Enabling CloudSec in MSD

On the NDFC Web UI, choose **LAN > Fabircs**. You can either create a new MSD fabric by clicking **Create Fabric** or edit the existing MSD fabric by clicking **Edit Fabric**.

Under the **DCI** tab, you can specify the CloudSec configuration details.

Multi-Site CloudSec – Enables CloudSec configurations on border gateways. If you enable this field, the remaining three fields for CloudSec are editable.

When Cloudsec is enabled at MSD level, NDFC also enables **dci-advertise-pip** under **evpn multisite border-gateway** and **tunnel-encryption** on the uplinks for all Cloudsec capable gateways.

When you click **Recalculate & Deploy**, you can verify theses configs in the **Preview Config** window for the border gateway switches.



Note CloudSec isn't supported if the border gateway has vPC or TRM is enabled on it, that is, TRM enabled on multisite overlay IFC. If CloudSec is enabled in this scenario, appropriate warning or error messages are generated.

CloudSec Key String – Specifies the hex key string. Enter a 66 hexadecimal string if you choose **AES_128_CMAC** or enter a 130 hexadecimal string if you choose **AES_256_CMAC**.

CloudSec Cryptographic Algorithm – Choose **AES_128_CMAC** or **AES_256_CMAC**.

CloudSec Enforcement – Specifies whether the CloudSec enforcement should be strict or loose.

strict – Deploys the CloudSec configuration to all the border gateways in fabrics in MSD. If there are any border gateways that don't support CloudSec, then an error message is generated, and the configuration isn't pushed to any switch.

If **strict** is chosen, the **tunnel-encryption must-secure** CLI is pushed to the CloudSec enabled gateways within MSD.

loose – Deploys the CloudSec configuration to all the border gateways in fabrics in MSD. If there are any border gateways that don't support CloudSec, then a warning message is generated. In this case, the CloudSec config is only deployed to the switches that support CloudSec. If **loose** is chosen, the **tunnel-encryption must-secure** CLI is removed if it exists.



Note There should be at least two fabrics in MSD with border gateways that support CloudSec. If there's only one fabric with a CloudSec capable device, then the following error message is generated:

CloudSec needs to have at least 2 sites that can support CloudSec.

To remove this error, meet the criteria of having at least two sites that can support CloudSec or disable CloudSec.

CloudSec Status Report Timer – Specifies the CloudSec Operational Status periodic report timer in minutes. This value specifies how often the NDFC polls the CloudSec status data from the switch. The default value is 5 minutes and the range is from 5 to 60 minutes.

Using the CloudSec feature in NDFC, you can have all the gateways within the MSD to use the same keychain (and have only one key string) and policy. You can provide one key chain string for NDFC to form the key chain policy. NDFC forms the encryption-policy by taking all default values. NDFC pushes the same key chain policy, the same encryption-policy, and encryption-peer policies to each CloudSec capable gateways. On each gateway, there's one encryption-peer policy for each remote gateway that is CloudSec capable, using the same keychain and same key policy.

If you don't want to use the same key for the whole MSD fabric or want to enable CloudSec on a subset of all sites, you can use **switch_freeform** to manually push the CloudSec config to the switches.

Capture all the CloudSec config in **switch_freeform**.

For example, the below config is included in the **switch_freeform** policy:

```
feature tunnel-encryption
evpn multisite border-gateway 600
  dci-advertise-pip
tunnel-encryption must-secure-policy
tunnel-encryption policy CloudSec_Policy1
tunnel-encryption source-interface loopback20
key chain CloudSec_Key_Chain1 tunnel-encryption
  key 1000
    key-octet-string 7 075e731f1a5c4f524f43595f507f7d737062677147524054590 70b0b0701585440
    cryptographic-algorithm AES_128_CMA
tunnel-encryption peer-ip 192.168.0.6
  keychain CloudSec_Key_Chain1 policy CloudSec_Policy1
```

Add **tunnel-encryption** in the Freeform Config of the uplink interface policy which will generate config like the following:

```
interface ethernet1/13
  no switchport
  ip address 192.168.1.14/24 tag 54321
  evpn multisite dci-tracking
  tunnel-encryption
  mtu 9216
  no shutdown
```

For more information, see Enabling Freeform Configurations on Fabric Switches.

When CloudSec configuration is added to or removed from the switch, the DCI uplinks will flap, which will trigger multisite BGP session flapping. For multisite with existing cross site traffic, there will be traffic disruption during this transition. Therefore, it is recommended to make the transition during a maintenance window.

If you're migrating an MSD fabric with the CloudSec configuration into NDFC, the Cloudsec related configuration is captured in **switch_freeform** and interface freeform config. You do not need to turn on Multi-Site Cloudsec in the MSD fabric setting. If you want to add more fabrics and establish CloudSec tunnels which share the same CloudSec policy including key as the existing one, then you can enable the CloudSec config in the MSD fabric settings. The CloudSec parameters in the MSD fabric setting need to match the existing CloudSec configuration on the switch. The CloudSec configuration is already captured in the freeform config, and enabling CloudSec in MSD will also generate config intents. Therefore, there's a double intent. For example, if you want to change the CloudSec key in the MSD settings, you need to remove the CloudSec freeform config because NDFC won't modify config in **switch_freeform**. Otherwise, the key in the MSD fabric settings is a conflict with the key in the freeform config.

Viewing CloudSec Operational State

You can use **CloudSec Operational View** to check the operational status of the CloudSec sessions if CloudSec is enabled on the MSD fabric.

Procedure

Step 1 Choose an MSD fabric.

The fabric topology window appears.

Step 2 Select **Actions** > **Detailed View**.

Step 3 Click the **Link** tab and choose **CloudSec Operational View** tab on the left.

Step 4 If CloudSec is disabled, the **CloudSec Operational View** isn't displayed.

The **Operational View** has the following fields and descriptions.

Fields	Description
Fabric Name	Specifies the fabrics that have a CloudSec session.
Session	Specifies the fabrics and border gateway switches involved in the CloudSec session.
Link State	Specifies the status of the CloudSec session. It can be in one of the following states: • Up: The CloudSec session is successfully established between the switches. • Down: The CloudSec session isn't operational.
Uptime	Specifies the duration of the uptime for the CloudSec session. Specifically, it's the uptime since the last Rx and Tx sessions flapped, and the smaller value among the 2 sessions is displayed.
Oper Reason	Specifies the down reason for the CloudSec session state.

Note
After CloudSec is enabled on a fabric, the operational status may not be available until after sessions are created, and the next status poll occurs.

Troubleshooting a CloudSec Session

If a CloudSec session is down, you can find more information about it using Programmable Report.

Procedure

Step 1 On the NDFC Web UI, choose **Operations > Programmable Reports**.

Step 2 Click **Create Report**.

Step 3 Specify a unique name for the report in the **Report Name** field.

Step 4 From the **Select Template** drop-down list, select **fabric_cloudsec_oper_status**.

Step 5 Click **Next** to view the **Source & Recurrence** tab.

Step 6 In the **Recurrence** field, choose the frequency at which the report job should be run.

Step 7 In the **Email Report To** field, enter an email ID or mailer ID if you want the report in an email.

You must configure SMTP settings in **Settings > Server Settings > SMTP** tab. If the Data service IP address is in private subnet, the static management route for SMTP server must be added in Cisco Nexus Dashboard cluster configuration.

Step 8 In the **Select fabric(s)** table, select the MSD fabric on which the report job should be run.

Step 9 Click **Save**.

The report job will be executed at the configured interval.