Enhanced Classic LAN, Release 12.2.2/12.2.3

Table of Contents

New and Changed Information

The following table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version	Feature	Description
NDFC release 12.2.3	Support for Cisco encryption type 6 in **BGP Authentication Key Encryption Type** field	When creating an Enhanced Classic LAN fabric type, support is now available for Cisco encryption type 6 in the **BGP Authentication Key Encryption Type** field. See Creating an Enhanced Classic LAN Fabric for more information.
NDFC release 12.2.3	Introduction of **Freeform** tab	When creating an Enhanced Classic LAN fabric type, the freeform configuration fields are now consolidated under a single **Freeform** tab. See Creating an Enhanced Classic LAN Fabric for more information.
NDFC release 12.2.2	Support for connecting fabrics using inter-fabric links with MACsec using a QKD server or a preshared key	With this feature, you can connect two fabrics using inter-fabric links with Media Access Control Security (MACsec) using a quantum key distribution (QKD) server for secure exchange of encryption keys. Beginning with NDFC 12.2.2, NDFC added support for MACsec with QKD for inter-fabric links for the following fabric types: • Data Center VXLAN EVPN • Enhanced Classic LAN • External Connectivity Network Prior to NDFC 12.2.2, NDFC supported MACsec for intra-fabric links for the Data Center VXLAN EVPN fabric and the BGP fabric. With this feature, NDFC added a **Security** tab. For more information, see Security. For more information on configuring MACsec with or without QKD, see Connecting Two Fabrics with MACsec Using QKD.
NDFC release 12.2.2	Support for assigning a vPC/port-channel ID range and for specifying a custom vPC/PO ID	With this feature, you can assign one virtual port channel (vPC)/port-channel ID range for aggregation-access pairing and you can specify a custom vPC/PO ID in an Enhanced Classic LAN fabric. Beginning with NDFC 12.2.2, NDFC added an **Action > Edit Pairing** option on the **Access Pairing** page for editing access and aggregation vPC/port-channel IDs. For more information, see Configuring a Specific vPC/Port-Channel ID Range for Aggregation-Access Pairing.

Creating an Enhanced Classic LAN Fabric

This document describes how to create a new Enhanced Classic LAN fabric using the **Enhanced Classic LAN** fabric template.

Note that this document gives information specifically for the fields that you will see in the **Enhanced Classic LAN** fabric template. See the Managing Legacy/Classic Networks in Cisco Nexus Dashboard Controller document for detailed procedures around managing legacy/classic networks in NDFC using the **Enhanced Classic LAN** fabric template.

1. Navigate to the **LAN Fabrics** page:

   Manage > Fabrics

2. Click **Actions > Create Fabric**.

   The **Create Fabric** window appears.

3. Enter a unique name for the fabric in the **Fabric Name** field, then click **Choose Fabric**.

   A list of all available fabric templates are listed.

4. From the available list of fabric templates, choose the **Enhanced Classic LAN** template, then click **Select**.
5. Enter the necessary field values to create a fabric.

   The tabs and their fields in the screen are explained in the following sections. The fabric level parameters are included in these tabs.

   - General Parameters
   - Spanning Tree
   - vPC
   - Protocols
   - Security
   - Advanced
   - Freeform (NDFC 12.2.3 release only)
   - Resources
   - Manageability
   - Bootstrap
   - Configuration Backup
   - Flow Monitor
6. When you have completed the necessary configurations, click **Save**.
   - Click on the fabric to display a summary in the slide-in pane.
   - Click on the **Launch** icon to display the **Fabric Overview**.

General Parameters

The **General Parameters** tab is displayed by default. The fields in this tab are described in the following table.

Field	Description
First Hop Redundancy Protocol	Specifies the FHRP protocol. Options are: · **none**: Select this option if you want Layer 2 only. · **hsrp** · **vrrp** · **vrrpv3**
Routing Protocol	Specifies the VRF-Lite Agg-Core/Edge or Collapsed Core-WAN peering protocol options. Options are: · **ebgp** · **ospf** · **none**: NDFC does not configure the peering protocol if the **none** option is selected. You must manually configure the peering protocol with this option, if necessary.
BGP ASN	This field becomes editable if you selected **ebgp** in the **Routing Protocol** field. Enter the BGP AS number the fabric is associated with. This must be same as existing fabric.
Enable Performance Monitoring	Check the check box to enable performance monitoring. Ensure that you do not clear interface counters from the Command Line Interface of the switches. Clearing interface counters can cause the Performance Monitor to display incorrect data for traffic utilization. If you must clear the counters and the switch has both clear counters and clear counters snmp commands (not all switches have the clear counters snmp command), ensure that you run both the main and the SNMP commands simultaneously. For example, you must run the clear counters interface ethernet slot/port command followed by the clear counters interface ethernet slot/port snmp command. This can lead to a one time spike.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Spanning Tree

The fields in the **Spanning Tree** tab are described in the following table. All of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Spanning Tree Root Bridge Protocol	Specify the protocol to be used for configuring Root Bridge: Options are: • **rpvst+**: Rapid Per-VLAN Spanning Tree • **mst**: Multiple Spanning Tree • **unmanaged (default)**: STP Root not managed by NDFC Spanning Tree settings and bridge configurations are applicable at the Aggregation layer only.
Spanning Tree VLAN Range	Specify the VLAN range. For example: 1, 3-5, 7, 9-11 The default value is 1-3967. Applicable only for Aggregation devices.
MST Instance Range	Specify the MST instance range. For example: 0-3,5,7-9 The default value is 0. Applicable only for Aggregation devices.
Spanning Tree Bridge Priority	Specify the bridge priority for the spanning tree in increments of 4096. Applicable only for Aggregation devices.
Spanning Tree Hello Interval	Set the number of seconds between the generation of config spanning-tree Bridge Protocol Data Unit (BPDU). The default value is 2. Applicable only for Aggregation devices.
Spanning Tree Forward Delay	Set the number of seconds for the forward delay timer. The default value is 15. Applicable only for Aggregation devices.
Spanning Tree Max Age Interval	Set the maximum number of seconds the information in a spanning-tree Bridge Protocol Data Unit (BPDU) is valid. The default value is 20. Applicable only for Aggregation devices.
Spanning Tree Pathcost Method	Options are: • **short**: (default): Use 16-bit based values for default port path costs • **long**: Use 32-bit based values for default port path costs Applicable only for Aggregation devices.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

vPC

The fields in the **vPC** tab are described in the following table. All of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
vPC Auto Recovery Time	Specifies the vPC auto recovery time-out period in seconds. • Minimum value: 240 • Maximum value: 3600
vPC Delay Restore Time	Specifies the vPC delay restore period in seconds. • Minimum value: 1 • Maximum value: 3600
vPC Peer Link Port Channel ID	Specifies the Port Channel ID for a vPC Peer Link. Default value in this field is 500. • Minimum value: 1 • Maximum value: 4096
vPC IPv6 ND Synchronize	Enables IPv6 Neighbor Discovery synchronization between vPC switches. The check box is enabled by default. Uncheck the check box to disable the function.
vPC Domain Id Range	Specifies the vPC Domain Id range to use for new pairings.
vPC Layer-3 Peer-Router Option	Enables the Layer 3 device to form peering adjacency with both the peers.  Configure this command in both the peers. If you configure this command only on one of the peers or you disable it on one peer, the operational state of layer 3 peer-router gets disabled. You get a notification when there is a change in the operational state.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Protocols

The fields in the **Protocols** tab are described in the following table. Most of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
OSPF Process Tag	This field becomes editable if you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. The OSPF Routing Process Tag. Maximum size is 20.

Field	Description
OSPF Area ID	This field becomes editable in these conditions: - If you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. - If you enter a value in the **OSPF Process Tag** field above. The **OSPF Area ID** in an IP address format.
OSPFv3 Process Tag	This field becomes editable if you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. The **OSPFv3 Routing Process Tag**. Maximum size is 20.
OSPFv3 Area ID	This field becomes editable in these conditions: - If you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. - If you enter a value in the **OSPFv3 Process Tag** field above. The **OSPFv3 Area ID** in an IP address format.
Enable BGP Authentication	This field becomes editable if you selected **ebgp** in the **Routing Protocol** field under the General Parameters tab. Select the check box to enable BGP authentication. Deselect the check box to disable it. If you enable this field, the **BGP Password Key Encryption Type** and **BGP Neighbor Password** fields are enabled.
BGP Password Key Encryption Type	This field becomes editable in these conditions: - If you selected **ebgp** in the **Routing Protocol** field under the General Parameters tab. - If you enabled the **Enable BGP Authentication** field above. Choose the encryption type: - 3 for the 3DES encryption type - 6 for the Cisco encryption type 6 (available only in NDFC 12.2.3 and later) - 7 for the Cisco encryption type 7
BGP Neighbor Password	This field becomes editable in these conditions: - If you selected **ebgp** in the **Routing Protocol** field under the General Parameters tab. - If you enabled the **Enable BGP Authentication** field above. Enter the VRF Lite BGP neighbor password as a hex string.

Field	Description
Enable OSPF Authentication	This field becomes editable if you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. Select the check box to enable OSPF authentication. Deselect the check box to disable it. If you enable this field, the **OSPF Authentication Key ID** and **OSPF Authentication Key** fields get enabled.
OSPF Authentication Key ID	This field becomes editable in these conditions: · If you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. · If you enabled the **Enable OSPF Authentication** field above. The **Key ID** is populated.
OSPF Authentication Key	This field becomes editable in these conditions: · If you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. · If you enabled the **Enable OSPF Authentication** field above. The OSPF authentication key must be the 3DES key from the switch. NOTE: Plain text passwords are not supported. Log in to the switch, retrieve the encrypted key and enter it in this field. For more information, see the *Retrieving the Authentication Key* section for details.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Security

The fields on the **Security** tab are described in the following table.

For more information on configuring data center interconnect (DCI) MACsec, see Connecting Two Fabrics with MACsec Using QKD.

Field	Description
Enable DCI MACsec	Check the check box to enable MACsec on DCI links.
Enable QKD	Check the check box to enable the QKD server for generating quantum keys for encryption. If you choose to not enable the **Enable QKD** option, NDFC uses preshared keys provided by the user instead of using the QKD server to generate the keys. If you disable the **Enable QKD** option, all the fields pertaining to QKD are grayed out.

Field	Description
DCI MACsec Cipher Suite	Choose one of the following DCI MACsec cipher suites for the DCI MACsec policy: • **GCM-AES-128** • **GCM-AES-256** • **GCM-AES-XPN-128** • **GCM-AES-XPN-256** The default value is **GCM-AES-XPN-256**.
DCI MACsec Primary Key String	Specify a Cisco Type 7 encrypted octet string that is used for establishing the primary DCI MACsec session. For **AES_256_CMAC**, the key string length must be 130 and for **AES_128_CMAC**, the key string length must be 66. If these values are not specified correctly, an error is displayed when you save the fabric.  The default key lifetime is infinite.
DCI MACsec Primary Cryptographic Algorithm	Choose the cryptographic algorithm used for the primary key string. It can be **AES_128_CMAC** or **AES_256_CMAC**. The default value is **AES_128_CMAC**. You can configure a fallback key on the device to initiate a backup session if the primary session fails.
DCI MACsec Fallback Key String	Specify a Cisco Type 7 encrypted octet string that is used for establishing a fallback MACsec session. For **AES_256_CMAC**, the key string length must be 130 and for **AES_128_CMAC**, the key string length must be 66. If these values are not specified correctly, an error is displayed when you save the fabric.  This parameter is mandatory if the **Enable QKD** option is not selected.
DCI MACsec Fallback Cryptographic Algorithm	Choose the cryptographic algorithm used for the fallback key string. It can be **AES_128_CMAC** or **AES_256_CMAC**. The default value is **AES_128_CMAC**.
QKD Profile Name	Specify the crypto profile name. The maximum size is 63.
KME Server IP	Specify the IPv4 address for the Key Management Entity (KME) server.
KME Server Port Number	Specify the port number for the KME server.
Trustpoint Label	Specify the authentication type trustpoint label. The maximum size is 64.
Ignore Certificate	Enable this check box to skip verification of incoming certificates.

Field	Description
MACsec Status Report Timer	Specify the MACsec operational status periodic report timer in minutes.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Advanced

The fields in the **Advanced** tab are described in the following table. All of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
VRF Template	Specifies the VRF template for creating VRFs. These are pre-built, best practice templates for VRFs that are provided with NDFC. You do not have to specify a template but one is automatically selected.
Network Template	Specifies the network template for creating networks. These are pre-built, best practice templates for networks that are provided with NDFC. You do not have to specify a template but one is automatically selected.
Layer 2 Host Interface MTU	Specifies the MTU for the layer 2 host interface. This value should be an even number.
Unshut Host Interfaces by Default	Check this check box to unshut the host interfaces by default.
Power Supply Mode	Choose the appropriate power supply mode.
CoPP Profile	Choose the appropriate Control Plane Policing (CoPP) profile policy for the fabric. By default, the strict option is populated.

Field	Description
Brownfield Network Name Format	Enter the format to be used to build the overlay network name during a brownfield import or migration. The network name should not contain any white spaces or special characters except underscore and hyphen. The network name must not be changed once the brownfield migration has been initiated. See the *Creating Networks for the Standalone Fabric* section for the naming convention of the network name. The syntax is [<string> | VLAN_ID] and the default value is **Auto_Net_VLANVLAN_ID**. When you create networks, the name is generated according to the syntax you specify. The following list describes the variables in the syntax: · **VLAN_ID**: Specifies the VLAN ID associated with the network. VLAN ID is specific to switches, hence Nexus Dashboard Fabric Controller picks the VLAN ID from one of the switches, where the network is found, randomly and use it in the name. We recommend not to use this unless the VLAN ID is consistent across the fabric. · **<string>**: This variable is optional and you can enter any number of alphanumeric characters that meet the network name guidelines. An example overlay network name: Site_VLAN1234 Ignore this field for greenfield deployments.
Enable CDP for Bootstrapped Switch	Enables CDP on management (mgmt0) interface for bootstrapped switch. By default, for bootstrapped switches, CDP is disabled on the mgmt0 interface.
Enable Tenant DHCP	Check the check box to enable feature dhcp and associated configurations globally on all switches in the fabric. This is a pre-requisite for support of DHCP for overlay networks that are part of the tenant VRFs. Ensure that **Enable Tenant DHCP** is enabled before enabling DHCP-related parameters in the overlay profiles.
Enable NX-API	Specifies enabling of NX-API on HTTPS. · There are subtle behavior changes to the EPL feature when the **NX-API** option is enabled or disabled. See the section EPL Behavior with NX-API Configuration for more information. · Disabling the **NX-API** option will change the behavior for the Layer 4 to Layer 7 service configuration feature. See the Guidelines and Limitations for Layer 4 to Layer 7 Services for more information.

Field	Description
NX-API HTTPS Port Number	Field becomes active if the **Enable NX-API** option is enabled. Enter the NX-API HTTPS port number. Default value is 443.
Enable HTTP NX-API	Specifies enabling of NX-API on HTTP. Enable this check box and the **Enable NX-API** check box to use HTTP. This check box is checked by default. If you uncheck this check box, the applications that use NX-API and supported by Cisco Nexus Dashboard Fabric Controller, such as Endpoint Locator (EPL), Layer 4-Layer 7 services (L4-L7 services), VXLAN OAM, and so on, start using HTTPS instead of HTTP. If you check the **Enable NX-API** check box and the **Enable NX-API on HTTP** check box, applications use HTTP.
NX-API HTTP Port Number	Field becomes active if the **Enable HTTP NX-API** option is enabled.Enter the NX-API HTTPS port number. Default value is 80.
Enable Strict Config Compliance	Enable the Strict Config Compliance feature by selecting this check box. It enables bi-directional compliance checks to flag additional configs in the running config that are not in the intent/expected config. By default, this feature is disabled.
Enable AAA IP Authorization	Enables AAA IP authorization, when IP Authorization is enabled in the remote authentication server. This is required to support Nexus Dashboard Fabric Controller in scenarios where customers have strict control of which IP addresses can have access to the switches.
Enable NDFC as Trap Host	Select this check box to enable Nexus Dashboard Fabric Controller as an SNMP trap destination. Typically, for a native HA Nexus Dashboard Fabric Controller deployment, the eth1 VIP IP address will be configured as SNMP trap destination on the switches. By default, this check box is enabled.
Enable Aggregation/Access Auto Pairing	For back-to-back vPCs, enable this option to automatically pair aggregation and access devices based on topology.
Create Route-map fabric-rmap-redist-subnet	Enable this option to create a route map fabric-rmap-redist-subnet. This route-map matches tag 12345.
Greenfield Cleanup Option	Enable this field to clean the switch configuration without a reloads when PreserveConfig=no. Valid options are Enable or Disable.
Enable Real Time Interface Statistics Collection	Valid for NX-OS fabrics only. Check the box to enable the collection of real time interface statistics.
Interface Statistics Load Interval	Enter the interval for the interface statistics load, in seconds (min: 5, max: 300).

The following fields are available in the **Advanced** tab in NDFC release 12.2.2, but are moved to the Freeform tab in NDFC release 12.2.3.

Field	Description
Aggregation Freeform Config	Additional CLIs for all Aggregation devices as captured from show running configuration.
Access Freeform Config	Additional CLIs for all Access devices as captured from show running configuration.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Freeform

This tab is introduced in NDFC release 12.2.3.

The fields in this tab are shown below. For more information, see Enabling Freeform Configurations on Fabric Switches.

Field	Description
Aggregation Pre-Interfaces Freeform Config	Enter additional CLIs, added before interface configurations, for all Aggregation devices as captured from Show Running Configuration.
Access Pre-Interfaces Freeform Config	Enter additional CLIs, added before interface configurations, for all Access devices as captured from Show Running Configuration.
Aggregation Post-Interfaces Freeform Config	Enter additional CLIs, added after interface configurations, for all Aggregation devices as captured from Show Running Configuration.
Access Post-Interfaces Freeform Config	Enter additional CLIs, added after interface configurations, for all Access devices as captured from Show Running Configuration.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Resources

The fields in the **Resources** tab are described in the following table. Most of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Network VLAN Range	VLAN range for the per switch overlay network (min:2, max:4094).

Field	Description
Aggregation-Core/Aggregation-Edge Connectivity	Specify the option for the VRF Lite Aggregation-Core and Aggregation-Edge Router Inter-Fabric connection. Options are: • **Auto**: Automatically generates the VRF Lite configuration on the Aggregation and Core switches. This option is applicable only if you are using the Cisco Nexus 7000 or 9000 Series switches for the Core layer. • **Manual**: If you are using the Cisco Catalyst 9000 series switches or Cisco ASR 9000 Series Aggregation Services Routers for the Core layer, select **Manual** in this field. You must manually create a policy using the necessary policy provided to you through NDFC. For more information, see VRF Lite.
VRF-Lite Subinterface dot1q Range	Specifies the per Aggregation dot1q Range for VRF Lite connectivity (min:2, max:4093).
Auto Generate VRF Lite Configuration on Aggregation and Core/Edge	Option that controls the automatic generation of the VRF Lite sub-interface and peering configurations on the Aggregation & Core/Edge devices. When this option is enabled, the automatically created VRF Lite links will have the 'Auto Generate Flag' enabled.
VRF Lite IP Version	Select the IP version for VRF Lite.Options are: • **IPv4_only** • **IPv6_only** • **IPv4_and_IPv6**
IPv4 VRF Subnet IP Range and **IPv4 VRF Subnet Mask Length**	The IPv4 address range to assign peer-to-peer Aggregation-Core connections, and peering between vPC Aggregation switches. Update the fields as needed. The values shown in your screen are automatically generated. If you want to update the IP address ranges or the VRF/Network VLAN ranges, ensure that it does not overlap with other ranges. You should only 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 Layer 2 and Layer 3 ranges, you should do the following: 1. Update the Layer 2 range and click **Save**. 2. Click the **Edit Fabric** option again, update the Layer 3 range and click **Save**.

Field	Description
IPv6 VRF Subnet IP Range and **IPv6 VRF Subnet Mask Length**	The IPv6 address range to assign peer-to-peer Aggregation-Core connections, and peering between vPC Aggregation switches. Update the fields as needed. The values shown in your screen are automatically generated. If you want to update the IP address ranges or the VRF/Network VLAN ranges, ensure that it does not overlap with other ranges. You should only 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 Layer 2 and Layer 3 ranges, you should do the following: 1. Update the Layer 2 range and click **Save**. 2. Click the **Edit Fabric** option again, update the Layer 3 range and click **Save**.
VRF Lite VLAN Range	VLAN range for Per VRF SVI Peering between Aggregation pairs (min:2, max:4094).
Use Specific vPC/Port-Channel ID Ranges	Specifies the custom range for a vPC ID for leaf-ToR switch pairing. The minimum allowed value is 1 and the maximum allowed value is 4099.
vPC/Port-Channel ID Ranges	Specifies the custom vPC ID range for auto-allocating a vPC ID for leaf-ToR switch pairing. The minimum allowed value is 1 and the maximum allowed value is 4099.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Manageability

The fields in the **Manageability** tab are described in the following table. Most of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
DNS Server IPs	Specifies the comma separated list of IP addresses (v4/v6) of the DNS servers.
DNS Server VRFs	Specifies one VRF for all DNS servers or a comma separated list of VRFs, one per DNS server.
NTP Server IPs	Specifies comma separated list of IP addresses (v4/v6) of the NTP server.
NTP Server VRFs	Specifies one VRF for all NTP servers or a comma separated list of VRFs, one per NTP server.
Syslog Server IPs	Specifies the comma separated list of IP addresses (v4/v6) IP address of the syslog servers, if used.
Syslog Server Severity	Specifies the comma separated list of syslog severity values, one per syslog server. The minimum value is 0 and the maximum value is 7. To specify a higher severity, enter a higher number.

Field	Description
Syslog Server VRFs	Specifies one VRF for all syslog servers or a comma separated list of VRFs, one per syslog server.
AAA Freeform Config	Specifies the AAA freeform configurations. If AAA configurations are specified in the fabric settings, **switch_freeform** PTI with source as **UNDERLAY_AAA** and description as **AAA Configurations** will be created.
Banner	Specifies the message of the day banner.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Bootstrap

The fields in the **Bootstrap** tab are described in the following table. Most of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Enable Bootstrap	Select this check box to enable the bootstrap feature. Bootstrap allows easy day-0 import and bring-up of new devices into an existing fabric. Bootstrap leverages the NX-OS POAP functionality. To add more switches and for POAP capability, chose check box for **Enable Bootstrap** and **Enable Local DHCP Server**. After you enable bootstrap, you can enable the DHCP server for automatic IP address assignment using one of the following methods: • External DHCP Server: Enter information about the external DHCP server in the **Switch Mgmt Default Gateway** and **Switch Mgmt IP Subnet Prefix** fields. • Local DHCP Server: Enable the **Local DHCP Server** check box and enter details for the remaining mandatory fields.
Enable Local DHCP Server	Select this check box to initiate enabling of automatic IP address assignment through the local DHCP server. When you select this check box, the **DHCP Scope Start Address** and **DHCP Scope End Address** fields become editable. If you do not select this check box, Nexus Dashboard Fabric Controller uses the remote or external DHCP server for automatic IP address assignment.

Field	Description
DHCP Version	Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you select DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled.  Cisco Nexus 9000 and 3000 Series Switches support IPv6 POAP only when switches are either Layer-2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 are not supported.
DHCP Scope Start Address and **DHCP Scope End Address**	Specifies the first and last IP addresses of the IP address range to be used for the switch out of band POAP.
Switch Mgmt Default Gateway	Specifies the default gateway for the management VRF on the switch.
Switch Mgmt IP Subnet Prefix	Specifies the prefix for the Mgmt0 interface on the switch. The prefix should be between 8 and 30. *DHCP scope and management default gateway IP address specification*: If you specify the management default gateway IP address 10.0.1.1 and subnet mask 24, ensure that the DHCP scope is within the specified subnet, between 10.0.1.2 and 10.0.1.254.
DHCPv4 Multi Subnet Scope	Specifies the field to enter one subnet scope per line. This field is editable after you check the **Enable Local DHCP Server** check box.The format of the scope should be defined as: **DHCP Scope Start Address, DHCP Scope End Address, Switch Management Default Gateway, Switch Management Subnet Prefix** For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24
Enable AAA Config	Select this check box to include AAA configurations from the Manageability tab as part of the device start-up config post bootstrap.
Bootstrap Freeform Config	(Optional) Enter additional commands as needed. For example, if you require some additional configurations to be pushed to the device and be available post device bootstrap, they can be captured in this field, to save the desired intent. After the devices boot up, they will contain the configuration defined in the **Bootstrap Freeform Config** field. Copy-paste the running-config to a **freeform config** field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. For more information, see Enabling Freeform Configurations on Fabric Switches.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Configuration Backup

The fields in the **Configuration Backup** tab are described in the following table. Most of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Hourly Fabric Backup	Select the check box to enable an hourly backup of fabric configurations and the intent. The hourly backups are triggered during the first 10 minutes of the hour.
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**. The scheduled backups are triggered exactly at the time you specify with a delay of up to two minutes. The scheduled backups are triggered regardless of the configuration deployment status. The number of fabric backups that will be retained on NDFC is decided by the **Admin > System Settings > Server Settings > LAN Fabric > Maximum Backups per Fabric**. The number of archived files that can be retained is set in the **# Number of archived files per device to be retained:** field in the **Server Properties** window. Note: To trigger an immediate backup, do the following: 1. Choose **Overview > Topology**. 2. Click within the specific fabric box. The fabric topology screen comes up. 3. Right-click on a switch within the fabric, then select **Preview Config**. 4. On the **Preview Config** page for this fabric, click **Re-Sync All**. You can also initiate the fabric backup on the fabric topology page. Click **Backup Now** in the **Actions** pane.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Flow Monitor

The fields in the **Flow Monitor** tab are described in the following table. Most of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can

update the fields if needed.

Field	Description
Enable Netflow	Check this check box to enable Netflow on Aggregation devices for this fabric. By default, Netflow is disabled. On Enable, NetFlow configuration will be applied to all Aggregation devices that support Netflow. When Netflow is enabled on the fabric, you can choose not to have netflow on a particular switch by having a dummy no_netflow PTI. If netflow is not enabled at the fabric level, an error message is generated when you enable netflow at the interface, network, or VRF level. For information about Netflow support for Cisco NDFC, see section "Netflow Support" in Understanding LAN Fabrics.

In the **Netflow Exporter** area, click **Actions > Add** to add one or more Netflow exporters. This exporter is the receiver of the netflow data. The fields on this screen are:

Field	Description
Exporter Name	Specifies the name of the exporter.
IP	Specifies the IP address of the exporter.
VRF	Specifies the VRF over which the exporter is routed.
Source Interface	Enter the source interface name.
UDP Port	Specifies the UDP port over which the netflow data is exported.

Click **Save** to configure the exporter. Click **Cancel** to discard. You can also choose an existing exporter and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

In the **Netflow Record** area, click **Actions > Add** to add one or more Netflow records. The fields on this screen are:

Field	Description
Record Name	Specifies the name of the record.
Record Template	Specifies the template for the record. Enter one of the record templates names.

The following two record templates are available for use. You can create custom netflow record templates. Custom record templates saved in the template library are available for use here.

- **netflow_ipv4_record** – to use the IPv4 record template.
- **netflow_l2_record** – to use the Layer 2 record template.
    - **Is Layer2 Record** – Check this check box if the record is for Layer2 netflow.

Click **Save** to configure the report. Click **Cancel** to discard. You can also choose an existing record and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

In the **Netflow Monitor** area, click **Actions > Add** to add one or more Netflow monitors. The fields on this screen are:

Field	Description
Monitor Name	Specifies the name of the monitor.
Record Name	Specifies the name of the record for the monitor.
Exporter1 Name	Specifies the name of the exporter for the netflow monitor.
Exporter2 Name (optional)	Specifies the name of the secondary exporter for the netflow monitor.

The record name and exporters referred to in each netflow monitor must be defined in **Netflow Record** and **Netflow Exporter**.

In the **Netflow Sampler** area, click **Actions > Add** to add one or more Netflow samplers. These are optional fields and are applicable only when there are N7K aggregation switches in the fabric. The fields on this screen are:

Field	Description
Sampler Name	Specifies the name of the sampler.
Number of Samples	Specifies the number of samples.
Number of Packets in Each Sampling	Specifies the number of packets in each sampling.

Click **Save** to configure the monitor. Click **Cancel** to discard. You can also choose an existing monitor and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

About Aggregation-Access Pairing in an Enhanced Classic LAN Fabric

With the NDFC 12.1.3 release, NDFC added a one-click vPC feature for automatically detecting and pairing aggregation and access switches for optimal traffic engineering. By default, the auto aggregation-access pairing option is enabled, which means that after you perform a **Recalculate and Deploy** operation, NDFC automatically detects the connectivity between the aggregation and the access switches and generates the appropriate configurations based on the detected supported topologies. The configurations include vPC domains that NDFC automatically pushes to the paired aggregation and access switches. The links between these aggregation-access pairs are bundled into a common vPC logical construct.

For more information on NDFC aggregation-access pairing, see the Enhanced Classic LAN in Cisco Nexus Dashboard Fabric Controller (NDFC) Release 12.1.3 white paper.

Workflow for Configuring Aggregation-Access Pairing

1. Create an Enhanced Classic LAN fabric. For more information, see Creating an Enhanced Classic LAN Fabric.
2. Discover the switches in the fabric. For more information, see the section "Adding Switches to a Fabric" in Add Switches for LAN Operational Mode.
3. Add the switches using a bootstrap. For more information, see the section "Adding Switches Using Bootstrap Mechanism" in Add Switches for LAN Operational Mode.
4. Define the roles for the aggregation and access switches. For more information, see the section "Assigning Switch Roles" in Add Switches for LAN Operational Mode.
5. Configure the vPC pairing. For more information, see the section "Creating a vPC Setup" in Add Switches for LAN Operational Mode.
6. Recalculate and deploy.

Create Aggregation-Access Pairings

1. Perform the following procedure to configure an aggregation and an access switch, where aggregation switches are connected to access switches through a port channel.
2. Add an aggregation and an access switch to an Enhanced Classic LAN fabric and set the role as either **Access** or **Aggregation** depending on the type of switch.

   With an Enhanced Classic LAN fabric, NDFC supports a minimum of two aggregation switches and the aggregation switches must be in a vPC pair.

3. On the **Fabric Overview > Switches** page, choose an aggregation switch.
4. Click **Actions > Access Pairing**.

   The **Access Pairing** page displays the aggregation switches on the top and a list of potential pairing access switches below the aggregation switches. The status of the aggregation switches display in the **Details** column.

5. Click **Save**.
6. On the **Fabric Overview** page, click **Actions > Recalculate and Deploy**.
7. After the configuration deployment is completed on the **Deploy Configuration** page, click **Close**.

Unpair Aggregation-Access Switches

1. Uncheck the **Enable <switch-name> Pairing as Access Pairing** check box to unpair the switches.

   You cannot unpair an aggregation-access pair if overlays are attached.

2. On the **Fabric Overview** page, click **Actions > Recalculate and Deploy** to complete the unpairing operation.

Configuring a Specific vPC/Port-Channel ID Range for Aggregation-Access Pairing

With this feature, you can:

- Configure a specific vPC/port-channel ID range for aggregation-access pairing by enabling the **Use Specific vPC/Port-Channel ID Range** field. NDFC then displays the **vPC/Port-Channel ID Range** field with the recommended vPC/port-channel ID range.
- Edit a vPC/port-channel ID for paired switches by clicking the **Action > Edit Pairing** option on the **Access Pairing** page.

Configure Fabric Settings for Specifying a vPC/Port-Channel ID Range for Aggregation-Access Pairing

1. On the **Fabric Overview** page, create an **Enhanced Classic LAN** fabric. For more information, see Creating an Enhanced Classic LAN Fabric.
2. Click on the **vPC** tab.
3. Check the **Use Specific vPC/Port-Channel ID Range** check box to use a specific vPC/port-channel ID range for aggregation-access pairing.

   The **vPC/Port-Channel ID Range** field displays the recommended values.

   The recommended values are from 1-499.

   

   You can increase the existing range or add more ranges if the values are exhausted.

4. Specify a range for the **vPC/Port-Channel ID Range** field if you do not want to use the recommended values.
5. Click **Save**.

   The new range applies to the new pairing.

Edit the Aggregation or the Access vPC/Port Channel IDs

1. On the **Fabric Overview > Switches** page, choose the aggregation switch you want to edit and click **Actions > Access Pairing**.

   The **Access Pairing** page appears with a horizontal bar of the paired aggregation switches.

2. Click **Edit Pairing** under the **Action** column.

   The access-aggregation paired switches page displays.

   The **Enable <switch-name> Pairing as Access Pairing** check box is checked due to auto aggregation-access pairing.

3. Click the arrow on the right-hand column of the page to view the fields.

4. Modify the access or the aggregation vPC/port-channel IDs if you want to change the values.
5. Click **Save**.

If you have overlays attached to the paired switches, you cannot change the vPC/port-channel IDs.

6. Navigate to the **Fabric Overview > Switches** page and click **Actions > Recalculate and Deploy**.

   The **Deploy Configuration** page displays with the list of aggregation switches.

   After successful deployment, the **Fabric Status** column displays as **In-Sync**.

Copyright