

Cisco Nexus: vPC Cluster on Switches

Summary

The vPC 'cluster' is an approach of building multiple vPC domains (each with up to two peers), forming separate Multichassis EtherChannel (MCEC). This allows the access or distribution layer architecture to be scaled beyond two boxes, maintaining high availability, aggregating bandwidth and eliminating points of failure.

What is a vPC Cluster?

Limitation of two peers per domain

Each vPC domain supports a maximum of two switch instances (peers). You cannot have three or more devices in one vPC domain.

Multi-domain architecture

To have more than two switches, multiple vPC domains are created. Each domain combines a pair of devices into a separate MCEC. End devices (servers, top-of-rack) connect via EtherChannel to a single pair, and different domains can serve different network segments.

vPC Cluster architecture

Device pairing

- Two Nexus instances in one vPC domain.
- Unique domain number (e.g. 10, 20, 30)

Peer-link and Peer-keepalive

- Peer-link - physical Port-Channel, trunking VLANs and carrying MCEC traffic.
- Peer-keepalive - a separate channel (ICMP/BFD) that monitors the control state.

Top-of-Rack and traffic aggregation

- Servers use LACP, connecting to both switches of a given pair.
- Subsequent domain pairs can support other servers or VLANs

Spanning Tree and redundancy

- STP treats the entire vPC domain as a single link - no loop blocking
- In case of port or switch failure, traffic is taken over by the other peer.

Step by step: Configuring Multiple vPC Domains

Enable vPC on all devices

```
feature vpc
feature interface-vlan
```

Configuration of vPC Domain 10 (Peer A-B)

```
vpc domain 10
  peer-keepalive destination 192.0.2.2 source 192.0.2.1

interface ethernet1/1-2
  channel-group 100 mode active

interface port-channel100
  switchport mode trunk
  vpc peer-link
```

vPC Domain configuration 20 (Peer C-D)

```
vpc domain 20
  peer-keepalive destination 192.0.2.4 source 192.0.2.3

interface ethernet1/3-4
  channel-group 200 mode active

interface port-channel200
  switchport mode trunk
  vpc peer-link
```

3.4 Aggregation of server traffic in Domain 10

```
interface ethernet1/10-11
  channel-group 110 mode active

interface port-channel110
  switchport mode access
  switchport access vlan 100
```

```
vpc 110
```

Aggregation of server traffic in Domain 20

```
interface ethernet1/12-13
  channel-group 210 mode active

interface port-channel210
  switchport mode access
  switchport access vlan 200
  vpc 210
```

Verification of the status of all domains

```
show vpc brief
show vpc consistency-parameters
```

Best practice

- Each vPC domain must have a unique number.
- Port-channel parameters (VLAN, MTU, trunk/access) must be identical on both sides.
- Peer-link should consist of at least two physical links for redundancy.
- Regular monitoring with `show vpc` commands and log analysis.

Applications

- Data centres - redundancy and aggregation of server connections.
- Distribution layer - scalability over two boxes.
- Critical applications - minimising points of failure and high availability.

Sources

<https://www.firewall.cx/cisco/cisco-data-center/nexus-vpc-configuration-design-operation-troubleshooting.html>

https://www.reddit.com/r/networking/comments/l1yqv1/question_can_we_create_multiple_vpc_domains_on_n9k/

https://www.cisco.com/c/dam/en/us/products/collateral/switches/nexus-7000-series-switches/C07-572835-00_NX-OS_vPC_DG.pdf

https://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/sw/design/vpc_design/vpc_best_practices_design_guide.pdf

<https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/nx-os-software/217274-understand-virtual-port-channel-vpc-en.html>

<https://community.cisco.com/t5/switching/multichassis-etherchannel/td-p/1390836>