FEX Active/Active Host vPC

« Fabric Extenders (FEX) | FEX Active/Active Fabric vPC »

Last updated: January 27, 2017

FEX Active/Active Host vPC

Note:

For information connecting to the KVM see Managing UCS C-Series Server Using CIMC

Objective

- Configure connectivity between the Nexus 5Ks and UCS C series server via a vPC from the FEXes to the server.

Task

- The UCS C server is pre-configured to use an LACP based NIC Team, and with the IP address 10.0.0.20**X**/24, where ***X*** is the server number (i.e. C200-1 is 10.0.0.201).
  - The GUI of the UCS C server can be accessed through the CIMC KVM and Windows RDP via the links on the jumpbox desktop link to "Z:\UCS C Servers". Both the CIMC and Windows credentials are cisco/cisco.
- Configure your 5Ks to communicate with the server as follows:
  - Configure the 5Ks in vPC domain 5.
  - Use the mgmt0 port as the vPC Peer Keepalive link.
  - Configure the links between the 5Ks as Port-Channel 5, and use this as the vPC Peer Link.
  - Configure the links from the first 5K to the first 2K as Port-Channel 101 and FEX number 101.
  - Configure the links from the second 5K to the second 2K as Port-Channel 102 and FEX number 102.
  - Configure the 5Ks' links to the server via the FEXes as Port-Channel 10 and vPC 10.
  - Port-Channel 10 should be an access port in VLAN 10.
  - Configure a VLAN 10 SVI on both your 5Ks with IP addresses 10.0.0.5**X**/24, where ***X*** is the device number.
- When complete you should have IP reachability from both your 5Ks to the UCS C server, and verify that the MAC address of the server's Team interface is being learned via the vPC member ports and not across the vPC Peer Link.
- Verify that IP reachability is maintained if the links from one of the 5Ks to its FEX go down or if the link from the FEX to the server goes down.

Configuration Click to *collapse*

N5K1:

```
feature interface-vlan
feature lacp
feature vpc
feature fex
!
vlan 10
!
vpc domain 5
  peer-keepalive destination 192.168.0.52
!
interface Vlan10
  no shutdown
  ip address 10.0.0.51/24
!
interface port-channel5
  switchport mode trunk
  spanning-tree port type network
  speed 10000
  vpc peer-link
!
interface port-channel10
  switchport access vlan 10
  speed 10000
  vpc 10
!
interface port-channel101
  switchport mode fex-fabric
  fex associate 101
!
interface Ethernet1/21
  switchport mode fex-fabric
  fex associate 101
  channel-group 101
!
interface Ethernet1/22
  switchport mode fex-fabric
  fex associate 101
  channel-group 101
!
interface Ethernet101/1/1
  switchport access vlan 10
  channel-group 10 mode active
```

N5K2:

```
feature interface-vlan
feature lacp
feature vpc
feature fex
!
vlan 10
!
vpc domain 5
  peer-keepalive destination 192.168.0.51
!
interface Vlan10
  no shutdown
```

```
  ip address 10.0.0.51/24
!
interface port-channel5
  switchport mode trunk
  spanning-tree port type network
  speed 10000
  vpc peer-link
!
interface port-channel10
  switchport access vlan 10
  speed 10000
  vpc 10
!
interface port-channel102
  switchport mode fex-fabric
  fex associate 102
!
interface Ethernet1/23
  switchport mode fex-fabric
  fex associate 102
  channel-group 102
!
interface Ethernet1/24
  switchport mode fex-fabric
  fex associate 102
  channel-group 102
!
interface Ethernet102/1/1
  switchport access vlan 10
  channel-group 10 mode active
```

Verification

Fabric Extender (FEX) and vPC topologies come in three forms. The first is a Host vPC, which has two FEXes connected straight-through to two parent switches, the FEXes dual homed to the server, and the vPC configured from the FEXes southbound to the server.

The second is a Fabric vPC, which has a single FEX dual homed to two parent switches, the FEX single homed to the server, and the vPC configured from the FEXes northbound to the parent switches.

The third is an Enhanced vPC (EvPC), which has two FEXes dual homed to two parent switches, the FEXes dual homed to the server, and the vPC configured from the FEXes both northbound to the parent switches and southbound to the server.

This example uses the Host vPC to achieve active/active forwarding to the server, and redundancy for any single link or node failure (with the exception of the server node itself). In most cases this is the preferred design of vPC and FEXes, as it simplifies change control between the two parent switches, allows for simpler code upgrades or ISSU, and still offers redundancy.

Like in the previous example, the first step in implementing this design is to pair the FEXes to their respective parent switches, which is verified below. Note that each parent switch only pairs with one of the downstream FEXes each, not both.

```
N5K1# show fex
  FEX         FEX           FEX             FEX             Fex
Number    Description      State           Model           Serial
------------------------------------------------------------------------
101         FEX0101                Online    N2K-C2232PP-10GE   SSI15030C1R
---         --------            Discovered    N2K-C2232PP-10GE   SSI16330GT8

N5K2# show fex
  FEX         FEX           FEX             FEX             Fex
Number    Description      State           Model           Serial
------------------------------------------------------------------------
102         FEX0102                Online    N2K-C2232PP-10GE   SSI16330GT8
---         --------            Discovered    N2K-C2232PP-10GE   SSI15030C1R
```

Link redundancy from the parent switches to the FEXes is achieved with a straight-through static port channel.

```
N5K1# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
5     Po5(SU)     Eth      LACP      Eth1/3(P)    Eth1/4(P)
10    Po10(SU)    Eth      LACP      Eth101/1/1(P)
101   Po101(SU)   Eth      NONE      Eth1/21(P)   Eth1/22(P)

N5K2# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
5     Po5(SU)     Eth      LACP      Eth1/3(P)    Eth1/4(P)
10    Po10(SU)    Eth      LACP      Eth102/1/1(P)
102   Po102(SU)   Eth      NONE      Eth1/23(P)   Eth1/24(P)
```

The links through the FEXes to the server are configured as the vPC member ports. From a vPC point of view there is no difference between this design and one without the FEXes, since the FEX ports appear as directly connected modules of the parent switches.

```
N5K1# show vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 5
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 1
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Operational Layer3 Peer-router    : Disabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
--------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po5    up     1,10

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
10     Po10        up     success     success                    10

N5K2# show vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 5
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 1
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Operational Layer3 Peer-router    : Disabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
--------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po5    up     1,10

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
10     Po10        up     success     success                    10
```

Similar to other vPC examples, the MAC address of the device attached to the vPC member ports should be learned via the vPC member port channel, and not via the vPC Peer Link.

```
N5K1# ping 10.0.0.201
PING 10.0.0.201 (10.0.0.201): 56 data bytes
64 bytes from 10.0.0.201: icmp_seq=0 ttl=127 time=0.709 ms
64 bytes from 10.0.0.201: icmp_seq=1 ttl=127 time=0.547 ms
64 bytes from 10.0.0.201: icmp_seq=2 ttl=127 time=0.572 ms
64 bytes from 10.0.0.201: icmp_seq=3 ttl=127 time=0.592 ms
64 bytes from 10.0.0.201: icmp_seq=4 ttl=127 time=0.553 ms

--- 10.0.0.201 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.547/0.594/0.709 ms

N5K1# show ip arp

Flags: * - Adjacencies learnt on non-active FHRP router
       + - Adjacencies synced via CFSoE
       # - Adjacencies Throttled for Glean
       D - Static Adjacencies attached to down interface

IP ARP Table for context default
Total number of entries: 1
Address         Age       MAC Address     Interface
10.0.0.201      00:00:35  d48c.b5bd.460c  Vlan10

N5K1# show mac address-table address d48c.b5bd.460c
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       d48c.b5bd.460c    dynamic   0          F    F  Po10

N5K2# show mac address-table address d48c.b5bd.460c
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
```

```
* 10       d48c.b5bd.460c    dynamic   0          F    F  Po10
```

In the case of a link or node failure, traffic should reconverge to the vPC Peer Link.

```
N5K1# config t
Enter configuration commands, one per line.  End with CNTL/Z.
N5K1(config)# int e1/21 - 22
N5K1(config-if-range)# shut
2017 Jan 19 00:28:54 N5K1 %ETH_PORT_CHANNEL-5-PORT_DOWN: port-channel101: Ethernet1/22 is down
2017 Jan 19 00:28:54 N5K1 %ETH_PORT_CHANNEL-5-PORT_DOWN: port-channel101: Ethernet1/21 is down
2017 Jan 19 00:28:54 N5K1 %ETH_PORT_CHANNEL-5-PORT_DOWN: port-channel101: port-channel101 is down
<snip>
2017 Jan 19 00:28:54 N5K1 %PFMA-2-FEX_STATUS: Fex 101 is offline


N5K1(config-if-range)# ping 10.0.0.201
PING 10.0.0.201 (10.0.0.201): 56 data bytes
64 bytes from 10.0.0.201: icmp_seq=0 ttl=127 time=1.71 ms
64 bytes from 10.0.0.201: icmp_seq=1 ttl=127 time=0.595 ms
64 bytes from 10.0.0.201: icmp_seq=2 ttl=127 time=0.558 ms
64 bytes from 10.0.0.201: icmp_seq=3 ttl=127 time=0.512 ms
64 bytes from 10.0.0.201: icmp_seq=4 ttl=127 time=0.523 ms
--- 10.0.0.201 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.512/0.779/1.71 ms

N5K1(config-if-range)# show mac address-table address d48c.b5bd.460c
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
+ 10       d48c.b5bd.460c    dynamic   0          F    F  Po5
```

« Fabric Extenders (FEX) | FEX Active/Active Fabric vPC »

^ back to top