



Fabric Extenders (FEX)

« Back to Back vPC | FEX Active/Active Host vPC »

Last updated: January 27, 2017

Fabric Extenders (FEX)

Note:

For information connecting to the KVM see [Managing UCS C-Series Server Using CIMC](#)

Objective

- Configure connectivity between the Nexus 5K and UCS C series server through the Fabric Extender (FEX).

Task

- The UCS C server is pre-configured to use an LACP based NIC Team, and with the IP address 10.0.0.20X/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 first 5K to pair with the first Fabric Extender as follows:
 - Enable the Fabric Extender feature.
 - Configure the 5K's links connecting to the 2K as FEX ports.
 - Aggregate the links from the 5K to the FEX as Port-Channel101.
 - Use module number 101 for this FEX.
- Configure the 5K to communicate with the UCS C server as follows:
 - Configure a VLAN 10 SVI on your 5K with IP address 10.0.0.5X/24, where X is the device number.
 - Configure the 5K's link to the UCS C server as a member of Port-Channel1 to accommodate this setup.
 - Port-Channel1 should be an access port in VLAN 10.
- When complete you should have IP reachability from your 5K to the UCS C server.

Configuration [Click to collapse](#)

```
N5K1:
feature interface-vlan
feature fex
feature lacp
!
vlan 10
!
interface Vlan10
no shutdown
ip address 10.0.0.51/24
!
interface port-channel1
switchport access vlan 10
speed 10000
!
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 1 mode active
```

Verification

Fabric Extenders (FEXes) are access switches that behave as remote line cards of a parent switch. After the FEX is paired with the parent switch such as a Nexus 5K or 7K, all configuration occurs on the upstream parent. From the parent switch's perspective, the FEX is simply another module/line card, and is configured as such.

The first step in configuring FEX is to discover what ports the extenders are connected on and their hardware details. This happens automatically as soon as `feature fex` is enabled, as seen in the output below:

```
N5K1# terminal monitor
N5K1# config t
Enter configuration commands, one per line. End with CNTL/Z.
N5K1(config)# feature fex
N5K1(config)#
2017 Jan 18 21:51:38 N5K1 %FEX-5-FEX_ENABLED: FEX Enabled
2017 Jan 18 21:51:40 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 1 of Fex 0 that is connected with Ethernet1/23 changed its status from Created to Discovered
2017 Jan 18 21:51:40 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 2 of Fex 0 that is connected with Ethernet1/24 changed its status from Created to Discovered
2017 Jan 18 21:51:41 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 2 of Fex 0 that is connected with Ethernet1/22 changed its status from Created to Discovered
2017 Jan 18 21:51:41 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 1 of Fex 0 that is connected with Ethernet1/21 changed its status from Created to Discovered

N5K1# show fex
```

FEX Number	FEX Description	FEX State	FEX Model	FEX Serial
---	-----	Discovered	N2K-C2232PP-10GE	SSI15030C1R
---	-----	Discovered	N2K-C2232PP-10GE	SSI16330GT8

Once FEXes are discovered, they need to be assigned a unique module number with the `flex associate` command.

```
N5K1# config t
Enter configuration commands, one per line. End with CNTL/Z.
N5K1(config)# interface e1/21 - 22
N5K1(config-if-range)# switchport mode flex-fabric
N5K1(config-if-range)# flex associate 101
2017 Jan 18 21:53:30 N5K1 %ETHPORT-5-IF_DOWN_CFG_CHANGE: Interface Ethernet1/21 is down(Config change)
2017 Jan 18 21:53:30 N5K1 %ETHPORT-5-IF_DOWN_CFG_CHANGE: Interface Ethernet1/22 is down(Config change)
<snip>
2017 Jan 18 21:53:30 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 1 of Fex 0 that is connected with Ethernet1/21 changed its status from Discovered to Created
2017 Jan 18 21:53:30 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 2 of Fex 0 that is connected with Ethernet1/22 changed its status from Discovered to Created
<snip>
2017 Jan 18 21:53:33 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 1 of Fex 0 that is connected with Ethernet1/21 changed its status from Created to Discovered
2017 Jan 18 21:53:33 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 2 of Fex 0 that is connected with Ethernet1/22 changed its status from Created to Discovered
2017 Jan 18 21:53:38 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 1 of Fex 101 that is connected with Ethernet1/21 changed its status from Discovered to Configured
2017 Jan 18 21:53:38 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 2 of Fex 101 that is connected with Ethernet1/22 changed its status from Discovered to Configured
<snip>
2017 Jan 18 21:53:41 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 0 of Fex 101 that is connected with Ethernet1/21 changed its status from Configured to Fabric Up
2017 Jan 18 21:53:41 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 0 of Fex 101 that is connected with Ethernet1/22 changed its status from Configured to Fabric Up
2017 Jan 18 21:53:41 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 1 of Fex 101 that is connected with Ethernet1/21 changed its status from Fabric Up to Connecting
2017 Jan 18 21:53:41 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 1 of Fex 101 that is connected with Ethernet1/21 changed its status from Connecting to Active
2017 Jan 18 21:53:41 N5K1 %ETHPORT-5-IF_UP: Interface Ethernet1/21 is up in mode Flex Fabric
2017 Jan 18 21:53:42 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 2 of Fex 101 that is connected with Ethernet1/22 changed its status from Fabric Up to Connecting
2017 Jan 18 21:53:42 N5K1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 2 of Fex 101 that is connected with Ethernet1/22 changed its status from Connecting to Active
2017 Jan 18 21:53:42 N5K1 %ETHPORT-5-IF_UP: Interface Ethernet1/22 is up in mode Flex Fabric
2017 Jan 18 21:54:46 N5K1 %PFMA-2-FEX_STATUS: Fex 101 is online
2017 Jan 18 21:54:46 N5K1 %NOHMS-2-NOHMS_ENV_FEX_ONLINE: FEX-101 On-line
2017 Jan 18 21:54:47 N5K1 %PFMA-2-FEX_STATUS: Fex 101 is online
```

If successful, the FEX should move to Online state. If the state is listed as "Image Download" it means that an automatic code upgrade/downgrade is occurring from the parent switch to the FEX.

```
N5K1# show flex
```

FEX Number	FEX Description	FEX State	FEX Model	FEX Serial
101	FEX0101	Online	N2K-C2232PP-10GE	SSI15030C1R
---	-----	Discovered	N2K-C2232PP-10GE	SSI16330GT8

Once the module is fully online the interfaces of the FEX appear as directly connected ports of the parent switch.

```
N5K1# show interface e101/1/1 - 32 status
```

Port	Name	Status	Vlan	Duplex	Speed	Type
Eth101/1/1	--	connected	1	full	10G	10Gbase-SR
Eth101/1/2	--	sfpAbsent	1	full	10G	--
Eth101/1/3	--	sfpAbsent	1	full	10G	--
Eth101/1/4	--	sfpAbsent	1	full	10G	--
Eth101/1/5	--	sfpAbsent	1	full	10G	--
Eth101/1/6	--	sfpAbsent	1	full	10G	--
Eth101/1/7	--	sfpAbsent	1	full	10G	--
Eth101/1/8	--	sfpAbsent	1	full	10G	--
Eth101/1/9	--	sfpAbsent	1	full	10G	--
Eth101/1/10	--	sfpAbsent	1	full	10G	--
Eth101/1/11	--	sfpAbsent	1	full	10G	--
Eth101/1/12	--	sfpAbsent	1	full	10G	--
Eth101/1/13	--	sfpAbsent	1	full	10G	--
Eth101/1/14	--	sfpAbsent	1	full	10G	--
Eth101/1/15	--	sfpAbsent	1	full	10G	--
Eth101/1/16	--	sfpAbsent	1	full	10G	--
Eth101/1/17	--	sfpAbsent	1	full	10G	--
Eth101/1/18	--	sfpAbsent	1	full	10G	--
Eth101/1/19	--	sfpAbsent	1	full	10G	--
Eth101/1/20	--	sfpAbsent	1	full	10G	--
Eth101/1/21	--	sfpAbsent	1	full	10G	--
Eth101/1/22	--	sfpAbsent	1	full	10G	--
Eth101/1/23	--	sfpAbsent	1	full	10G	--
Eth101/1/24	--	sfpAbsent	1	full	10G	--
Eth101/1/25	--	sfpAbsent	1	full	10G	--
Eth101/1/26	--	sfpAbsent	1	full	10G	--
Eth101/1/27	--	sfpAbsent	1	full	10G	--
Eth101/1/28	--	sfpAbsent	1	full	10G	--
Eth101/1/29	--	sfpAbsent	1	full	10G	--
Eth101/1/30	--	sfpAbsent	1	full	10G	--
Eth101/1/31	--	sfpAbsent	1	full	10G	--

Detailed output about the FEX can be seen as follows:

```

NSK1# show fex detail
FEX: 101 Description: FEX0101 state: Online
FEX version: 7.3(0)N1(1) [Switch version: 7.3(0)N1(1)]
FEX Interim version: 7.3(0)N1(1)
Switch Interim version: 7.3(0)N1(1)
Extender Serial: SSI15030C1R
Extender Model: N2K-C2232PP-10GE, Part No: 73-12533-04
Card Id: 82, Mac Addr: c8:9c:1d:3a:19:82, Num Macs: 64
Module Sw Gen: 12594 [Switch Sw Gen: 21]
Post level: complete
Pinning-mode: static Max-links: 1
Fabric port for control traffic: Eth1/21
FCoE Admin: false
FCoE Oper: true
FCoE FEX AA Configured: false
Fabric interface state:
Eth1/21 - Interface Up. State: Active
Eth1/22 - Interface Up. State: Active
Fex Port      State Fabric Port
Eth101/1/1    Up      Eth1/21
Eth101/1/2    Down    None
Eth101/1/3    Down    None
Eth101/1/4    Down    None
Eth101/1/5    Down    None
Eth101/1/6    Down    None
Eth101/1/7    Down    None
Eth101/1/8    Down    None
Eth101/1/9    Down    None
Eth101/1/10   Down    None
Eth101/1/11   Down    None
Eth101/1/12   Down    None
Eth101/1/13   Down    None
Eth101/1/14   Down    None
Eth101/1/15   Down    None
Eth101/1/16   Down    None
Eth101/1/17   Down    None
Eth101/1/18   Down    None
Eth101/1/19   Down    None
Eth101/1/20   Down    None
Eth101/1/21   Down    None
Eth101/1/22   Down    None
Eth101/1/23   Down    None
Eth101/1/24   Down    None
Eth101/1/25   Down    None
Eth101/1/26   Down    None
Eth101/1/27   Down    None
Eth101/1/28   Down    None
Eth101/1/29   Down    None
Eth101/1/30   Down    None
Eth101/1/31   Down    None
Eth101/1/32   Down    None

Logs:
01/18/2017 21:54:43.423077: Module register received
01/18/2017 21:54:43.426524: Registration response sent
01/18/2017 21:54:43.460300: create module inserted event.
01/18/2017 21:54:43.461407: Module Online Sequence
01/18/2017 21:54:46.685622: Module Online

```

In the above output we can see the model number (2232PP), the serial number (SSI15030C1R), if FCoE is enabled, and which interfaces are being used to forward traffic from the downstream server interfaces upstream to the parent switch. By default, the "max-links" is set to 1, which means only one uplink is active. To increase backplane bandwidth and reduce oversubscription to the downstream servers, either the pinning configuration can be changed, or the uplinks can be aggregated together into a port channel.

```

NSK1# conf t
Enter configuration commands, one per line. End with CNTL/Z.
NSK1(config)# interface e1/21 - 22
NSK1(config-if-range)# channel-group 101
2017 Jan 18 22:13:05 NSK1 %ETH_PORT_CHANNEL-5-CREATED: port-channel101 created
2017 Jan 18 22:13:05 NSK1 %ETHPORT-5-IF_DOWN_CHANNEL_MEMBERSHIP_UPDATE_IN_PROGRESS: Interface Ethernet1/22 is down (Channel membership update in progress)
2017 Jan 18 22:13:05 NSK1 %ETHPORT-5-IF_DOWN_CHANNEL_MEMBERSHIP_UPDATE_IN_PROGRESS: Interface Ethernet1/21 is down (Channel membership update in progress)
<snip>
2017 Jan 18 22:13:06 NSK1 %NOHMS-2-NOHMS_ENV_FEX_OFFLINE: FEX-101 Off-line (Serial Number SSI15030C1R)
2017 Jan 18 22:13:06 NSK1 %PFMA-2-FEX_STATUS: Fex 101 is offline
<snip>
2017 Jan 18 22:13:12 NSK1 %ETH_PORT_CHANNEL-5-PORT_UP: port-channel101: Ethernet1/21 is up
2017 Jan 18 22:13:12 NSK1 %ETH_PORT_CHANNEL-5-FOP_CHANGED: port-channel101: first operational port changed from none to Ethernet1/21
2017 Jan 18 22:13:12 NSK1 %ETHPORT-5-IF_UP: Interface port-channel101 is up in mode Fex Fabric
2017 Jan 18 22:13:12 NSK1 %ETHPORT-5-IF_UP: Interface Ethernet1/21 is up in mode Fex Fabric
2017 Jan 18 22:13:13 NSK1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 2 of Fex 101 that is connected with Ethernet1/22 changed its status from Fabric Up to Connecting
2017 Jan 18 22:13:13 NSK1 %FEX-5-FEX_PORT_STATUS_NOTI: Uplink-ID 2 of Fex 101 that is connected with Ethernet1/22 changed its status from Connecting to Active
2017 Jan 18 22:13:13 NSK1 %ETH_PORT_CHANNEL-5-PORT_UP: port-channel101: Ethernet1/22 is up
2017 Jan 18 22:13:13 NSK1 %ETHPORT-5-IF_UP: Interface Ethernet1/22 is up in mode Fex Fabric

```

After the port channel is added from the 5K down to the FEX, all downstream server traffic is pinned to the port channel instead of the physical links, and traffic is load balanced accordingly.

```

NSK1# show fex detail
FEX: 101 Description: FEX0101 state: Online
FEX version: 7.3(0)N1(1) [Switch version: 7.3(0)N1(1)]

```

```

FEX Version: 7.3(0)N1(1)
FEX Interim version: 7.3(0)N1(1)
Switch Interim version: 7.3(0)N1(1)
Extender Serial: SSI15030C1R
Extender Model: N2K-C2232PP-10GE, Part No: 73-12533-04
Card Id: 82, Mac Addr: c8:9c:1d:3a:19:82, Num Macs: 64
Module Sw Gen: 12594 [Switch Sw Gen: 21]
Post level: complete
Pinning-mode: static Max-links: 1
Fabric port for control traffic: Eth1/21
FCoE Admin: false
FCoE Oper: true
FCoE FEX AA Configured: false
Fabric interface state:
  Po101 - Interface Up. State: Active
  Eth1/21 - Interface Up. State: Active
  Eth1/22 - Interface Up. State: Active
Fex Port      State Fabric Port
  Eth101/1/1   Up      Po101
  Eth101/1/2   Down    None
<snip>

```

Since the UCS C server is preconfigured for an LACP based port channel, the 5K needs to run LACP even though there is only one physical link to the server in this example.

```

N5K1# show lacp interface e101/1/1
Interface Ethernet101/1/1 is up
Channel group is 1 port channel is Po1
PDU sent: 60
PDU rcvd: 7
Markers sent: 0
Markers rcvd: 0
Marker response sent: 0
Marker response rcvd: 0
Unknown packets rcvd: 0
Illegal packets rcvd: 0
Lag Id: [ [(0, d4-8c-b5-bd-46-c, 0, 0, 1), (8000, 0-de-fb-12-1a-7c, 8000, 8000, 2201)] ]
Operational as aggregated link since Wed Jan 18 22:46:20 2017

Local Port: Eth101/1/1 MAC Address= 0-de-fb-12-1a-7c
System Identifier=0x8000,0-de-fb-12-1a-7c
Port Identifier=0x8000,0x2201
Operational key=32768
LACP_Activity=active
LACP_Timeout=Long Timeout (30s)
Synchronization=IN_SYNC
Collecting=true
Distributing=true
Actor Admin State=(Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0)
Actor Oper State=(Ac-1:To-0:Ag-1:Sy-1:Co-1:Di-1:De-0:Ex-0)
Neighbor: 0/1
MAC Address= d4-8c-b5-bd-46-c
System Identifier=0x0,d4-8c-b5-bd-46-c
Port Identifier=0x0,0x1
Operational key=0
LACP_Activity=active
LACP_Timeout=short Timeout (1s)
Synchronization=IN_SYNC
Collecting=true
Distributing=true
Partner Admin State=(Ac-0:To-1:Ag-0:Sy-0:Co-0:Di-0:De-0:Ex-0)
Partner Oper State=(Ac-1:To-1:Ag-1:Sy-1:Co-1:Di-1:De-0:Ex-0)

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
-----
1      Po1(SU)   Eth      LACP    Eth101/1/1(P)
101   Po101(SU) Eth      NONE    Eth1/21(P)  Eth1/22(P)

```

Note that there are some behavioral differences between FEX host ports and other physical links. For example, the FEX ports always run as STP Edge Ports with BPDU Guard enabled. This can be seen in the below show output, although this configuration does not show up at the port level in the running config.

```

N5K1# show run int po1 membership

!Command: show running-config interface port-channel1 membership
!Time: Wed Jan 18 22:51:49 2017

version 7.3(0)N1(1)

interface port-channel1
 switchport access vlan 10
 speed 10000

interface Ethernet101/1/1

 switchport access vlan 10
 channel-group 1 mode active

N5K1# show spanning-tree interface po1 detail

Port 4096 (port-channel1) of VLAN0010 is designated forwarding

```

```
Port path cost 2, Port priority 128, Port Identifier 128.4096
Designated root has priority 32778, address 00de.fb12.1a7c
Designated bridge has priority 32778, address 00de.fb12.1a7c
Designated port id is 128.4096, designated path cost 0
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 1
The port type is edge
Link type is point-to-point by default
Bpdu guard is enabled
Bpdu filter is not enabled by default
BPDU: sent 11, received 0
```



[« Back to Back vPC | FEX Active/Active Host vPC »](#)

[^ back to top](#)