



Port Channels

« VLANs & Trunking | Rapid Spanning Tree (RSTP) Traffic Engineering »

Last updated: January 18, 2017

Port Channels

Note:

Configure the task "VLANs & Trunking" prior to beginning this lab.

Objective

- Configure and verify link aggregation and load balancing on Nexus 5Ks and 7Ks with Port Channels.

Task

- Per the previous task the network should be configured as follows:
 - The first server's link to the first 5K is in VLAN 10 with address 10.0.0.1X/24, where X is the server number.
 - The first server's link to the second 5K is in VLAN 20 with address 20.0.0.1X/24, where X is the server number.
 - The second server's link to the first 5K is in VLAN 20 with address 20.0.0.1X/24, where X is the server number.
 - The second server's link to the second 5K is in VLAN 10 with address 10.0.0.1X/24, where X is the server number.
 - 5Ks and 7Ks have SVIs for VLANs 10 & 20 with IP addressing 10.0.0.X/24 and 20.0.0.X/24 respectively, where X is the device number.
 - 5Ks have access ports to the servers and trunk ports to the 7Ks, and the 7Ks have trunks between each other.
- Disable the diagonal links from the 5Ks to the 7Ks.
- Configure a port channel between the 5Ks as follows:
 - All links between the 5Ks should be aggregated together in Port-Channel 5.
 - Do not use dynamic negotiation for the port channel.
 - The port channel should be a trunk, but only allow VLAN 10 traffic.
- Configure a port channel between the 7Ks as follows:
 - All links between the 7Ks should be aggregated together in Port-Channel 7.
 - Use dynamic negotiation per 802.3ad for the port channel.
 - Control packets for negotiating the link should be transmitted every 1 second.
 - If the port channel has less than 2 member ports it should be disabled.
 - The port channel should be a trunk, but only allow VLAN 20 traffic.
- Once complete you should have IP reachability between the servers in both VLANs 10 & 20.
- Verify that VLAN 10 traffic is being load balanced on the port channel between the 5Ks, and that VLAN 20 traffic is being load balanced on the port channel between the 7Ks.

Configuration [Click to collapse](#)

```
NSK1:
vlan 10,20
!
interface Ethernet1/1
 switchport
 switchport mode access
 switchport access vlan 10
 no shutdown
!
interface Ethernet1/2
 switchport
 switchport mode access
 switchport access vlan 20
 no shutdown
!
interface Ethernet1/3 - 4
 switchport
 switchport mode trunk
 switchport trunk allowed vlan 10
 channel-group 5
 no shutdown
!
interface port-channel5
 switchport mode trunk
 switchport trunk allowed vlan 10
 speed 10000
!
interface Ethernet1/5
 switchport
 switchport mode trunk
 switchport trunk allowed vlan 10,20
 no shutdown
!
interface Ethernet1/6
 shutdown
```

```
NSK2:
vlan 10,20
!
interface Ethernet1/1
 switchport
 switchport mode access
 switchport access vlan 20
 no shutdown
!
interface Ethernet1/2
 switchport
 switchport mode access
 switchport access vlan 10
 no shutdown
!
interface Ethernet1/3 - 4
```

```

switchport
switchport mode trunk
switchport trunk allowed vlan 10
channel-group 5
no shutdown
!
interface port-channel5
switchport mode trunk
switchport trunk allowed vlan 10
speed 10000
!
interface Ethernet1/5
shutdown
!
interface Ethernet1/6
switchport
switchport mode trunk
switchport trunk allowed vlan 10,20
no shutdown

```

N7K1:

```

feature lacp
!
vlan 10,20
!
interface Ethernet1/1 - 2
lacp rate fast
switchport
switchport mode trunk
switchport trunk allowed vlan 20
channel-group 7 mode active
no shutdown
!
interface port-channel7
switchport
switchport mode trunk
switchport trunk allowed vlan 20
lacp min-links 2
!
interface Ethernet1/3 - 4
switchport
switchport mode trunk
switchport trunk allowed vlan 10,20
no shutdown

```

N7K2:

```

feature lacp
!
vlan 10,20
!
interface Ethernet1/1 - 2
lacp rate fast
switchport
switchport mode trunk
switchport trunk allowed vlan 20
channel-group 7 mode active
no shutdown
!
interface port-channel7
switchport
switchport mode trunk
switchport trunk allowed vlan 20
lacp min-links 2
!
interface Ethernet1/3 - 4
switchport
switchport mode trunk
switchport trunk allowed vlan 10,20
no shutdown

```

Verification

The Flags in the below output indicate that N5K1's channel to N5K2 has formed with no negotiation.

```

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-Channel	Type	Protocol	Member Ports
5	Po5(SU)	Eth	NONE	Eth1/3(P) Eth1/4(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-Channel	Type	Protocol	Member Ports
5	Po5(SU)	Eth	NONE	Eth1/3(P) Eth1/4(P)

From STP's point of view, the port channel is one logical link, as the lower level member interface details are hidden from the STP process. The result of this is that active/active forwarding can be achieved across all member ports.

```
N5K1# show spanning-tree vlan 10

VLAN0010
Spanning tree enabled protocol rstp
Root ID    Priority    32778
  Address   0026.980c.2142
  Cost      3
  Port      4100 (port-channel5)
  Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec

Bridge ID  Priority    32778 (priority 32768 sys-id-ext 10)
  Address   00de.fb12.1a7c
  Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec

Interface      Role Sts Cost      Prio.Nbr Type
-----
Po5             Root FWD 1        128.4100 P2p
Eth1/1          Desg FWD 2        128.129  P2p
Eth1/5          Desg FWD 2        128.133  P2p
```

From a layer 2 forwarding point of view, the member port details are likewise hidden. The result is that if one of the member ports goes down, or another member is added, the CAM table does not need to be flushed and re-learned.

```
N5K1# show mac address-table vlan 10
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   0000.0000.0010     dynamic   70       F F          Eth1/1
* 10   0000.0000.0021     dynamic   90       F F          Po5
* 10   00de.fb12.1a01     dynamic  190       F F          Po5

N5K2# show mac address-table vlan 10
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   0000.0000.0010     dynamic  420       F F          Po5
* 10   0000.0000.0021     dynamic  110       F F          Eth1/2
* 10   00de.fb12.1a7c     dynamic  450       F F          Po5
```

To check the load balancing result, Server 2 is set to listen for TCP connections. Server 1 then sends 8 TCP flows to Server 2, each with unique source ports.

```
Server 2:
C:\Users\Administrator\Desktop\iperf3>iperf3.exe -s
-----
Server listening on 5201
-----
```

```
Server 1:
C:\Users\Administrator\Desktop\iperf3>iperf3.exe -c 10.0.0.12 -t 0 -P 8
Connecting to host 10.0.0.12, port 5201
[ 4] local 10.0.0.11 port 59772 connected to 10.0.0.12 port 5201
[ 6] local 10.0.0.11 port 59773 connected to 10.0.0.12 port 5201
[ 8] local 10.0.0.11 port 59774 connected to 10.0.0.12 port 5201
[10] local 10.0.0.11 port 59775 connected to 10.0.0.12 port 5201
[12] local 10.0.0.11 port 59776 connected to 10.0.0.12 port 5201
[14] local 10.0.0.11 port 59777 connected to 10.0.0.12 port 5201
[16] local 10.0.0.11 port 59778 connected to 10.0.0.12 port 5201
[18] local 10.0.0.11 port 59779 connected to 10.0.0.12 port 5201
[ ID] Interval           Transfer             Bandwidth
[ 4]  0.00-1.05   sec   87.0 MBytes       697 Mbits/sec
[ 6]  0.00-1.05   sec  102 MBytes       821 Mbits/sec
[ 8]  0.00-1.05   sec   69.4 MBytes       556 Mbits/sec
[10]  0.00-1.05   sec   65.6 MBytes       526 Mbits/sec
[12]  0.00-1.05   sec   49.0 MBytes       393 Mbits/sec
[14]  0.00-1.05   sec   64.2 MBytes       515 Mbits/sec
[16]  0.00-1.05   sec   71.6 MBytes       574 Mbits/sec
[18]  0.00-1.05   sec   62.0 MBytes       497 Mbits/sec
[SUM] 0.00-1.05   sec   571 MBytes       4.58 Gbits/sec
```

The end result is that traffic is load balanced on a near 1:1 basis, as seen in both the counters of show interface , and the traffic percentages of show port-channel traffic .

```
N5K1# show int e1/3 - 4 | in "output rate"
30 seconds output rate 1473540712 bits/sec, 121143 packets/sec
  input rate 6.91 Mbps, 12.45 Kpps; output rate 2.07 Gbps, 170.27 Kpps
30 seconds output rate 1535883440 bits/sec, 126277 packets/sec
  input rate 6.94 Mbps, 12.51 Kpps; output rate 2.05 Gbps, 168.51 Kpps

N5K1# show port-channel traffic
ChanId      Port Rx-Ucst Tx-Ucst Rx-Mcst Tx-Mcst Rx-Bcst Tx-Bcst
-----
5           Eth1/3 51.07% 48.42% 49.13% 56.64% 8.08% 100.00%
5           Eth1/4 48.92% 51.57% 50.86% 43.35% 91.91% 0.0%
```

For VLAN 20 traffic, the 5Ks must first layer 2 switch the traffic to the 7Ks, as the East/West trunk link between them only allowed VLAN 10.

```
N5K1# show mac address-table vlan 20
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
-----+-----+-----+-----+-----+-----
* 20    0000.0000.0011  dynamic  10     F    F    Eth1/5
* 20    0000.0000.0020  dynamic  10     F    F    Eth1/2
* 20    0026.980c.2142  dynamic  280    F    F    Eth1/5
* 20    00de.fb12.1a01  dynamic  300    F    F    Eth1/5

N5K2# show mac address-table vlan 20
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
-----+-----+-----+-----+-----+-----
* 20    0000.0000.0011  dynamic  10     F    F    Eth1/1
* 20    0000.0000.0020  dynamic  10     F    F    Eth1/6
* 20    0026.980c.2142  dynamic  310    F    F    Eth1/6
* 20    00de.fb12.1a7c  dynamic  310    F    F    Eth1/6
```

The end result of this topology is that VLAN 20 traffic is forced to transit the port channel between N7K1 and N7K2.

```
N7K1# show mac address-table vlan 20
Note: MAC table entries displayed are getting read from software.
Use the 'hardware-age' keyword to get information related to 'Age'

Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since last seen,+ - primary entry using vPC Peer-Link, E - EVPN entry
(T) - True, (F) - False , ~~~ - use 'hardware-age' keyword to retrieve age info
VLAN/BD  MAC Address    Type    age    Secure NTFY Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----
G 20     68bd.abd7.6042  static   -      F    F    sup-eth1(R)
* 20     0000.0000.0011  dynamic  ~~~    F    F    Po7
* 20     0000.0000.0020  dynamic  ~~~    F    F    Eth1/3
* 20     0026.980c.2142  dynamic  ~~~    F    F    Po7
* 20     00de.fb12.1a01  dynamic  ~~~    F    F    Po7
* 20     00de.fb12.1a7c  dynamic  ~~~    F    F    Eth1/3

N7K2# show mac address-table vlan 20
Note: MAC table entries displayed are getting read from software.
Use the 'hardware-age' keyword to get information related to 'Age'

Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since last seen,+ - primary entry using vPC Peer-Link, E - EVPN entry
(T) - True, (F) - False , ~~~ - use 'hardware-age' keyword to retrieve age info
VLAN/BD  MAC Address    Type    age    Secure NTFY Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----
G 20     0026.980c.2142  static   -      F    F    sup-eth1(R)
* 20     0000.0000.0011  dynamic  ~~~    F    F    Eth1/4
* 20     0000.0000.0020  dynamic  ~~~    F    F    Po7
* 20     00de.fb12.1a01  dynamic  ~~~    F    F    Eth1/4
* 20     00de.fb12.1a7c  dynamic  ~~~    F    F    Po7
```

Similar to the above VLAN 10 example, Server 1 & 2 send TCP flows to each other to test load balancing.

```
Server 2:
C:\Users\Administrator\Desktop\iperf3>iperf3.exe -s
-----
Server listening on 5201
-----

Server 1:
C:\Users\Administrator\Desktop\iperf3>iperf3.exe -c 20.0.0.12 -t 0 -P 8
Connecting to host 20.0.0.12, port 5201
[ 4] local 20.0.0.11 port 59791 connected to 20.0.0.12 port 5201
[ 6] local 20.0.0.11 port 59792 connected to 20.0.0.12 port 5201
[ 8] local 20.0.0.11 port 59793 connected to 20.0.0.12 port 5201
[10] local 20.0.0.11 port 59794 connected to 20.0.0.12 port 5201
[12] local 20.0.0.11 port 59795 connected to 20.0.0.12 port 5201
[14] local 20.0.0.11 port 59796 connected to 20.0.0.12 port 5201
[16] local 20.0.0.11 port 59797 connected to 20.0.0.12 port 5201
[18] local 20.0.0.11 port 59798 connected to 20.0.0.12 port 5201
[ ID] Interval           Transfer         Bandwidth
[ 4]  0.00-1.00   sec   60.5 MBytes   506 Mbits/sec
[ 6]  0.00-1.00   sec   46.9 MBytes   392 Mbits/sec
[ 8]  0.00-1.00   sec   45.8 MBytes   382 Mbits/sec
[10]  0.00-1.00   sec   50.6 MBytes   423 Mbits/sec
[12]  0.00-1.00   sec   56.5 MBytes   472 Mbits/sec
[14]  0.00-1.00   sec   51.5 MBytes   430 Mbits/sec
[16]  0.00-1.00   sec   47.1 MBytes   394 Mbits/sec
[18]  0.00-1.00   sec   59.0 MBytes   493 Mbits/sec
[SUM] 0.00-1.00   sec   418 MBytes   3.49 Gbits/sec
```

Like the 5Ks, the traffic across the 7Ks port channel is load balanced nearly 1:1 between members of the port channel.

```

N7K2# show interface e1/2 | in "output rate"
 30 seconds output rate 2403984720 bits/sec, 197615 packets/sec
input rate 8.53 Mbps, 15.39 Kpps; output rate 2.40 Gbps, 197.62 Kpps
300 seconds output rate 1243230768 bits/sec, 102146 packets/sec
input rate 4.43 Mbps, 7.95 Kpps; output rate 1.24 Gbps, 102.15 Kpps
30 seconds output rate 2416879520 bits/sec, 198697 packets/sec
input rate 8.52 Mbps, 15.39 Kpps; output rate 2.42 Gbps, 198.70 Kpps
300 seconds output rate 1267232944 bits/sec, 104114 packets/sec
input rate 4.49 Mbps, 8.07 Kpps; output rate 1.27 Gbps, 104.11 Kpps

N7K2# show port-channel traffic
NOTE: Clear the port-channel member counters to get accurate statistics

ChanId      Port Rx-Ucst Tx-Ucst Rx-Mcst Tx-Mcst Rx-Bcst Tx-Bcst
-----
7           Eth1/1 49.52% 49.43% 52.27% 94.79% 66.66% 55.55%
7           Eth1/2 50.47% 50.56% 47.72%  5.20% 33.33% 44.44%

```

The default timers for an LACP negotiated port channel are 30 second hellos and a 90 second dead interval, as seen on the 5Ks below. Note that the channel shows as invalid since N5K1 isn't using LACP for negotiation, hence the 0 packets sent and received, but the default timer values still show up.

```

N5K1# show lacp interface e1/3
Interface Ethernet1/3 is invalid
Channel group is 0 port channel is
PDU's sent: 0
PDU's rcvd: 0
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, 0-0-0-0-0-0, 0, 0, 0), (0, 0-0-0-0-0-0, 0, 0, 0)] ]
Operational as aggregated link since Thu Jan 1 00:00:00 1970

Local Port: Eth1/3  MAC Address= 0-de-fb-12-1a-7c
System Identifier=0x8000, Port Identifier=0x8000,0x0
Operational key=0
LACP_Activity=passive
LACP_Timeout=Long Timeout (30s)
Synchronization=NOT_IN_SYNC
Collecting=false
Distributing=false
Partner information refresh timeout=Long Timeout (90s)
Actor Admin State=(Ac-0:To-0:Ag-0:Sy-0:Co-0:Di-0:De-0:Ex-0)
Actor Oper State=(Ac-0:To-0:Ag-0:Sy-0:Co-0:Di-0:De-0:Ex-0)
Neighbor: 0x0
MAC Address= 0-0-0-0-0-0
System Identifier=0x0, Port Identifier=0x0,0x0
Operational key=0
LACP_Activity=unknown
LACP_Timeout=Long Timeout (30s)
Synchronization=NOT_IN_SYNC
Collecting=false
Distributing=false
Partner Admin State=(Ac-0:To-0:Ag-0:Sy-0:Co-0:Di-0:De-0:Ex-0)
Partner Oper State=(Ac-0:To-0:Ag-0:Sy-0:Co-0:Di-0:De-0:Ex-0)
Aggregate or Individual(True=1)= 2

```

The "fast" LACP timers send hellos every 1 second and have a dead timer of 3 seconds, as seen on the 7K below. Also note that LACP PDUs are being both sent and received.

```

N7K1# show lacp interface e1/1
Interface Ethernet1/1 is up
Channel group is 7 port channel is Po7
PDU's sent: 994
PDU's rcvd: 1001
Markers sent: 0
Markers rcvd: 0
Marker response sent: 0
Marker response rcvd: 0
Unknown packets rcvd: 0
Illegal packets rcvd: 0
Lag Id: [ [(8000, 0-26-98-c-21-42, 6, 8000, 101), (8000, 68-bd-ab-d7-60-42, 6, 8000, 101)] ]
Operational as aggregated link since Wed Jan 18 04:37:09 2017

Local Port: Eth1/1  MAC Address= 68-bd-ab-d7-60-42
System Identifier=0x8000, Port Identifier=0x8000,0x101
Operational key=6
LACP_Activity=active
LACP_Timeout=Short Timeout (1s)
Synchronization=IN_SYNC
Collecting=true
Distributing=true
Partner information refresh timeout=Short Timeout (3s)
Actor Admin State=63
Actor Oper State=63
Neighbor: 0x101
MAC Address= 0-26-98-c-21-42
System Identifier=0x8000, Port Identifier=0x8000,0x101
Operational key=6
LACP_Activity=active
LACP_Timeout=short Timeout (1s)
Synchronization=IN_SYNC
Collecting=true
Distributing=true
Partner Admin State=63
Partner Oper State=63
Aggregate or Individual(True=1)= 1

```

To verify the "min-links" configuration on the 7Ks, shut down one of the member ports. The end result should be that the other member port is suspended, as seen below:

```
N7K1# config t
Enter configuration commands, one per line. End with CNTL/Z.
N7K1(config)# int e1/1
N7K1(config-if)# shutdown
2017 Jan 18 04:31:36 N7K1 %ETH_PORT_CHANNEL-5-PORT_DOWN: port-channel7: Ethernet1/1 is down
2017 Jan 18 04:31:36 N7K1 %ETHPORT-5-IF_BANDWIDTH_CHANGE: Interface port-channel7,bandwidth changed to 10000000 Kbit
2017 Jan 18 04:31:36 N7K1 %ETHPORT-5-IF_DOWN_CFG_CHANGE: Interface Ethernet1/1 is down(Config change)
2017 Jan 18 04:31:36 N7K1 %ETHPORT-5-IF_DOWN_ADMIN_DOWN: Interface Ethernet1/1 is down (Administratively down)
2017 Jan 18 04:31:36 N7K1 %ETHPORT-5-IF_DOWN_INITIALIZING: Interface port-channel7 is down (Initializing)
2017 Jan 18 04:31:36 N7K1 %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty by cisco on 192.168.0.1@pts/0
2017 Jan 18 04:31:36 N7K1 %ETH_PORT_CHANNEL-5-FOP_CHANGED: port-channel7: first operational port changed from Ethernet1/2 to none
2017 Jan 18 04:31:36 N7K1 %ETH_PORT_CHANNEL-5-PORT_DOWN: port-channel7: Ethernet1/2 is down
2017 Jan 18 04:31:36 N7K1 %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-channel7 is down (No operational members)
2017 Jan 18 04:31:36 N7K1 %ETHPORT-5-IF_BANDWIDTH_CHANGE: Interface port-channel7,bandwidth changed to 100000 Kbit
2017 Jan 18 04:31:36 N7K1 %ETHPORT-5-IF_DOWN_CFG_CHANGE: Interface Ethernet1/2 is down(Config change)
2017 Jan 18 04:31:37 N7K1 %ETHPORT-5-SPEED: Interface Ethernet1/2, operational speed changed to 10 Gbps
2017 Jan 18 04:31:37 N7K1 %ETHPORT-5-IF_DUPLEX: Interface Ethernet1/2, operational duplex mode changed to Full
2017 Jan 18 04:31:37 N7K1 %ETHPORT-5-IF_RX_FLOW_CONTROL: Interface Ethernet1/2, operational Receive Flow Control state changed to off
2017 Jan 18 04:31:37 N7K1 %ETHPORT-5-IF_TX_FLOW_CONTROL: Interface Ethernet1/2, operational Transmit Flow Control state changed to off
2017 Jan 18 04:31:37 N7K1 %ETHPORT-5-SPEED: Interface port-channel7, operational speed changed to 10 Gbps
2017 Jan 18 04:31:37 N7K1 %ETHPORT-5-IF_DUPLEX: Interface port-channel7, operational duplex mode changed to Full
2017 Jan 18 04:31:37 N7K1 %ETHPORT-5-IF_RX_FLOW_CONTROL: Interface port-channel7, operational Receive Flow Control state changed to off
2017 Jan 18 04:31:37 N7K1 %ETHPORT-5-IF_TX_FLOW_CONTROL: Interface port-channel7, operational Transmit Flow Control state changed to off
2017 Jan 18 04:31:40 N7K1 %ETH_PORT_CHANNEL-5-PORT_SUSPENDED: Ethernet1/2: Ethernet1/2 is suspended by protocol, min-link requirement not met
```



[« VLANs & Trunking | Rapid Spanning Tree \(RSTP\) Traffic Engineering »](#)

[^ back to top](#)