



## Active Active NIC Teaming with vPC

« STP Edge Ports | vPC and HSRP »

Last updated: January 27, 2017

### Active Active NIC Teaming with vPC

#### Objective

- Establish active/active connectivity between the Nexus 5Ks and their attached dual-homed servers.

#### Task

- Configure a vPC Domain between your 5Ks as follows:
  - Use vPC Domain 5.
  - Use the mgmt0 ports for the vPC Peer Keepalive Link.
  - Use Port-Channel 5 between the 5Ks for the the vPC Peer Link.
  - The vPC Peer Link should use LACP negotiation, be an 802.1q trunk link, and be an STP Network Port.
- Configure the vPC Member Ports on your 5Ks as follows:
  - The links to your first Server should be in Port-Channel 1, and be vPC 1.
  - The links to your second Server should be in Port-Channel 2, and be vPC 2.
  - Both vPCs should use LACP negotiation, be access ports in VLAN 10, and be STP Edge Ports.
- Configure your Servers as follows:
  - Configure the links on the Servers going to the 5Ks as a NIC Team using LACP.
  - Use IP addresses 10.0.0.1X/24 on the Team interface, where X is the Server number.
- Configure interface VLAN 10 on your 5Ks with IP addresses 10.0.0.X/24, where X is the device number.
- Once complete, you should have IP reachability between the 5Ks and the Servers.
- Verify that traffic between the Servers does not forward over the vPC Peer Link unless one of the vPC Member Ports is down.

#### Configuration [Click to collapse](#)



#### Note:

For Server configuration see the Intel support document [Setting up NIC teaming for Windows Server 2012/2012 R2/2016](#)

```
N5K1:
feature lacp
feature vpc
feature interface-vlan
!
vlan 10
!
interface Vlan10
ip address 10.0.0.51/24
no shutdown
!
vpc domain 5
peer-keepalive destination 192.168.0.52
!
interface Ethernet1/3 - 4
switchport
switchport mode trunk
spanning-tree port type network
channel-group 5 mode active
no shutdown
!
interface port-channel5
switchport
switchport mode trunk
spanning-tree port type network
vpc peer-link
no shutdown
!
interface Ethernet1/1
switchport
switchport mode access
switchport access vlan 10
spanning-tree port type edge
channel-group 1 mode active
no shutdown
!
interface port-channel1
switchport
switchport mode access
spanning-tree port type edge
vpc 1
no shutdown
!
interface Ethernet1/2
switchport
switchport mode access
switchport access vlan 10
spanning-tree port type edge
channel-group 2 mode active
no shutdown
!
interface port-channel2
switchport
switchport mode access
spanning-tree port type edge
vpc 2
no shutdown
```

```

no shutdown
N5K2:
feature lacp
feature vpc
feature interface-vlan
!
vlan 10
!
interface Vlan10
ip address 10.0.0.52/24
no shutdown
!
vpc domain 5
peer-keepalive destination 192.168.0.51
!
interface Ethernet1/3 - 4
switchport
switchport mode trunk
spanning-tree port type network
channel-group 5 mode active
no shutdown
!
interface port-channel5
switchport
switchport mode trunk
spanning-tree port type network
vpc peer-link
no shutdown
!
interface Ethernet1/1
switchport
switchport mode access
switchport access vlan 10
spanning-tree port type edge
channel-group 1 mode active
no shutdown
!
interface port-channel1
switchport
switchport mode access
spanning-tree port type edge
vpc 1
no shutdown
!
interface Ethernet1/2
switchport
switchport mode access
switchport access vlan 10
spanning-tree port type edge
channel-group 2 mode active
no shutdown
!
interface port-channel2
switchport
switchport mode access
spanning-tree port type edge
vpc 2
no shutdown

```

## Verification

In this design, the end servers are dual attached to separate access switches, N5K1 and N5K2. Additionally, N5K1 and N5K2 are configured for Virtual Port Channel (vPC), which is a type of Multi-Chassis EtherChannel (MEC). vPC means that the downstream devices, Server 1 and Server 2 in this case, see the upstream switches (the vPC Peers) as a single switch. In other words, while the physical topology is a triangle, the logical topology is a point-to-point port channel.

vPC configuration is made up of three main components, the vPC Peer Keepalive Link, the vPC Peer Link, and the vPC Member Ports. The vPC Keepalive Link is any layer 3 interface, including the mgmt0 port, that is used to send UDP pings between the vPC peers. If the UDP ping is successful over the keepalive link, the peers are considered to be reachable. The second portion, the vPC Peer Link, is used to synchronize the control plane between the vPC Peers. The Peer Link is used for operations such as MAC address table synchronization, ARP table synchronization, IGMP Snooping synchronization, and so on. The Peer Link is a port channel made up of at least two 10Gbps links, and is a layer 2 trunk link that runs as STP port type network. The final portions, the vPC member ports, are the port channel interfaces that go connect to the end hosts or downstream devices.

The first step in vPC verification is to ensure that the vPC Peer Keepalive is up and that the vPC Peer Link is up, as shown below.

```

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 : 2
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
<snip>

```

<https://t.me/learningnets>

Details of the Keepalive can be seen as follows:

```

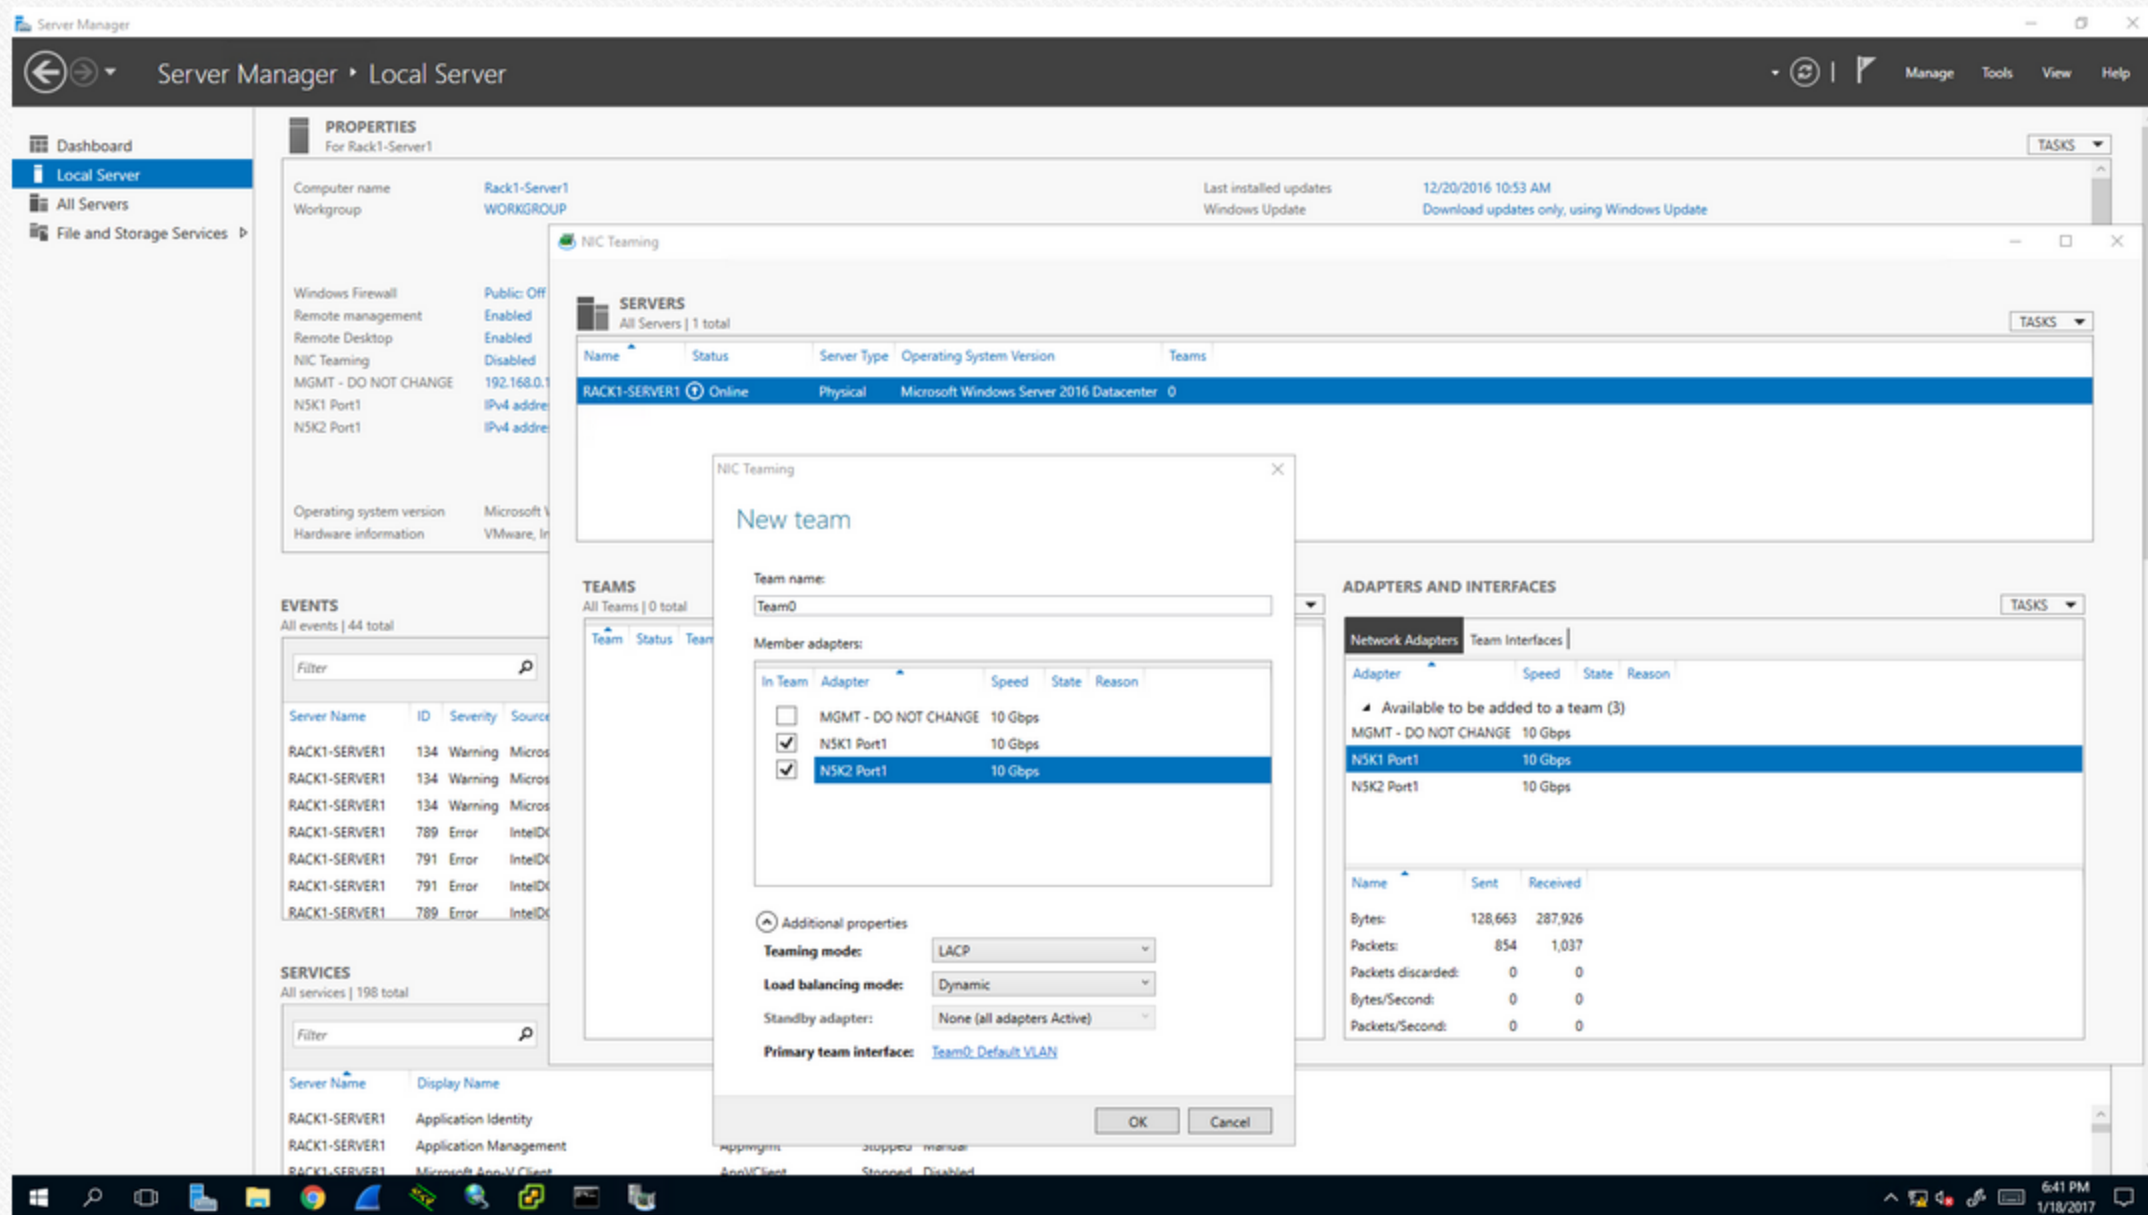
NSK1# show vpc peer-keepalive

vPC keep-alive status      : peer is alive
--Peer is alive for       : (1321) seconds, (922) msec
--Send status             : Success
--Last send at           : 2017.01.27 19:45:32 1 ms
--Sent on interface      : mgmt0
--Receive status         : Success
--Last receive at       : 2017.01.27 19:45:32 3 ms
--Received on interface  : mgmt0
--Last update from peer  : (0) seconds, (36) msec

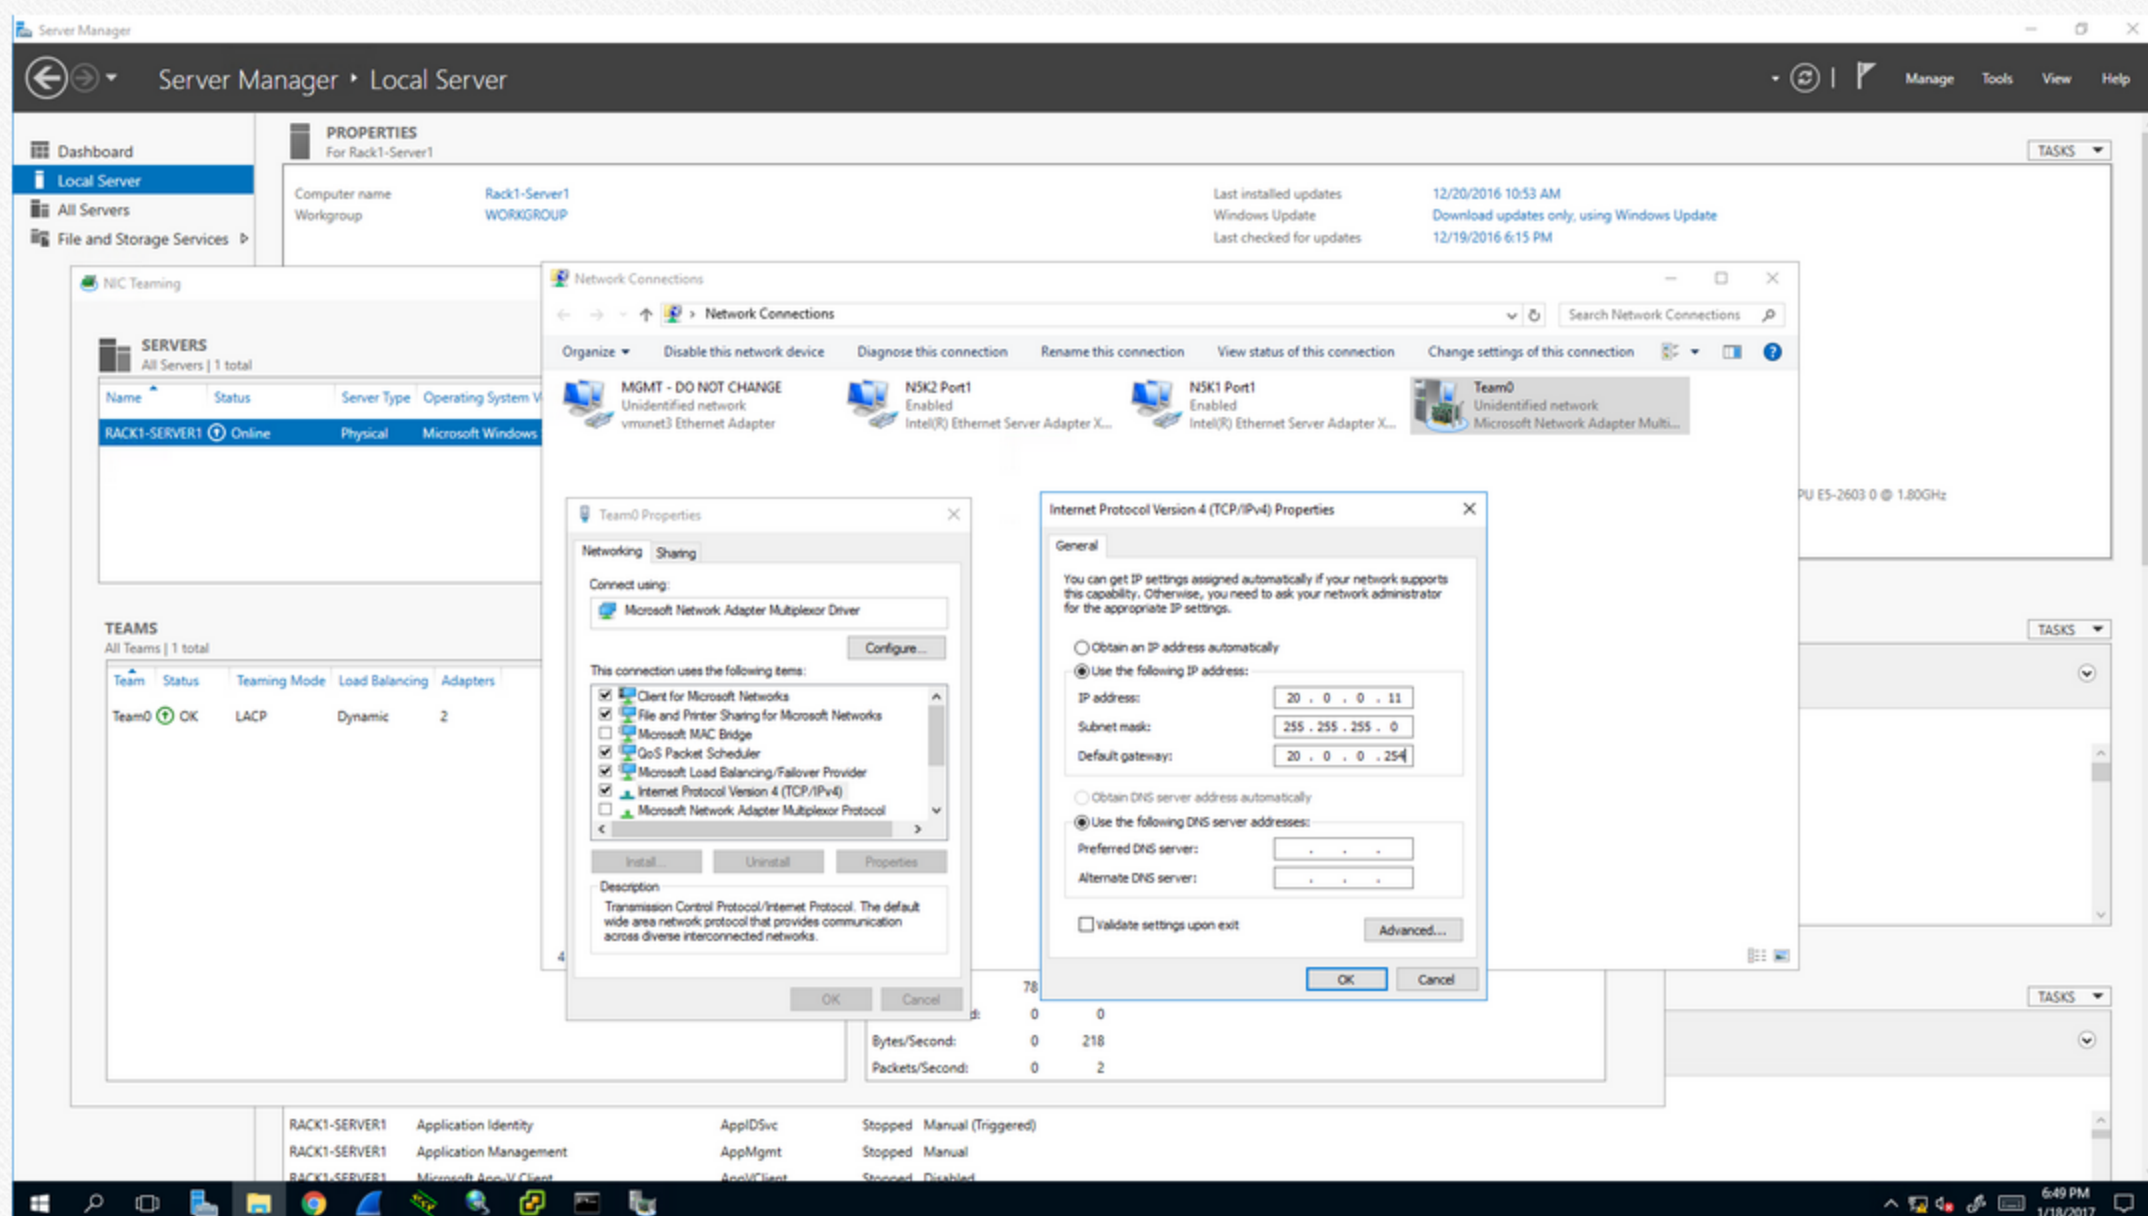
vPC Keep-alive parameters
--Destination            : 192.168.0.52
--Keepalive interval     : 1000 msec
--Keepalive timeout      : 5 seconds
--Keepalive hold timeout : 3 seconds
--Keepalive vrf          : management
--Keepalive udp port     : 3200
--Keepalive tos          : 192

```

Next, the end server is configured for NIC Teaming. In Windows 2016 Server, this is configured under the Server Manager as seen below:



IP addressing is configured on the logical Team adapter.



After the server signals the switch with LACP, the channel can form and the vPC comes up, as shown below.

```

NSK1#
2017 Jan 27 19:44:13 NSK1 %ETHPORT-5-SPEED: Interface port-channel1, operational speed changed to 10 Gbps
2017 Jan 27 19:44:13 NSK1 %ETHPORT-5-IF_DUPLEX: Interface port-channel1, operational duplex mode changed to Full
2017 Jan 27 19:44:13 NSK1 %ETHPORT-5-IF_RX_FLOW_CONTROL: Interface port-channel1, operational Receive Flow Control state changed to off
2017 Jan 27 19:44:13 NSK1 %ETHPORT-5-IF_TX_FLOW_CONTROL: Interface port-channel1, operational Transmit Flow Control state changed to off
2017 Jan 27 19:44:17 NSK1 %ETH_PORT_CHANNEL-5-PORT_UP: port-channel1: Ethernet1/1 is up
2017 Jan 27 19:44:17 NSK1 %ETH_PORT_CHANNEL-5-FOP_CHANGED: port-channel1: first operational port changed from none to Ethernet1/1
2017 Jan 27 19:44:17 NSK1 %ETHPORT-5-IF_UP: Interface Ethernet1/1 is up in mode access
2017 Jan 27 19:44:17 NSK1 %ETHPORT-5-IF_UP: Interface port-channel1 is up in mode access

```

```

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      Eth1/1(P)
2      Po2(SU)    Eth       LACP      Eth1/2(P)
5      Po5(SU)    Eth       LACP      Eth1/3(P)  Eth1/4(P)

N5K1# show vpc 1

vPC status
-----
id      Port      Status Consistency Reason      Active vlans
-----
1      Po1      up      success    success
                                         10

```

Once traffic flows to the servers, the 5Ks synchronize the MAC address table.

```

N5K1# ping 10.0.0.11
PING 10.0.0.11 (10.0.0.11): 56 data bytes
36 bytes from 10.0.0.51: Destination Host Unreachable
Request 0 timed out
64 bytes from 10.0.0.11: icmp_seq=1 ttl=127 time=0.911 ms
64 bytes from 10.0.0.11: icmp_seq=2 ttl=127 time=0.729 ms
64 bytes from 10.0.0.11: icmp_seq=3 ttl=127 time=0.585 ms
64 bytes from 10.0.0.11: icmp_seq=4 ttl=127 time=0.631 ms

--- 10.0.0.11 ping statistics ---
5 packets transmitted, 4 packets received, 20.00% packet loss
round-trip min/avg/max = 0.585/0.713/0.911 ms

N5K1# ping 10.0.0.12
PING 10.0.0.12 (10.0.0.12): 56 data bytes
36 bytes from 10.0.0.51: Destination Host Unreachable
Request 0 timed out
64 bytes from 10.0.0.12: icmp_seq=1 ttl=127 time=1.012 ms
64 bytes from 10.0.0.12: icmp_seq=2 ttl=127 time=0.679 ms
64 bytes from 10.0.0.12: icmp_seq=3 ttl=127 time=0.654 ms
64 bytes from 10.0.0.12: icmp_seq=4 ttl=127 time=0.654 ms

--- 10.0.0.12 ping statistics ---
5 packets transmitted, 4 packets received, 20.00% packet loss
round-trip min/avg/max = 0.654/0.749/1.012 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: 2
Address      Age      MAC Address  Interface
10.0.0.11    00:00:04  001b.218d.3d99  Vlan10
10.0.0.12    00:00:05  001b.2188.8075  Vlan10

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      001b.2188.8075   dynamic   0         F   F   Po2
* 10      001b.218d.3d99   dynamic   10        F   F   Po1
* 10      00de.fb12.1a01   static    0         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      001b.2188.8075   dynamic   0         F   F   Po2
* 10      001b.218d.3d99   dynamic   0         F   F   Po1
* 10      00de.fb12.1a7c   static    0         F   F   Po5

```

To test load balancing, multiple TCP flows are sent between the servers.

```

Server 2:
C:\Users\Administrator\Desktop\iperf3>ipconfig | find "10.0"
IPv4 Address. . . . . : 10.0.0.12
Default Gateway . . . . . : 10.0.0.254

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 -P 8 -t 0

```

<https://t.me/learningnets>

```

Connecting to host 10.0.0.12, port 5201
[ 4] local 10.0.0.11 port 49787 connected to 10.0.0.12 port 5201
[ 6] local 10.0.0.11 port 49788 connected to 10.0.0.12 port 5201
[ 8] local 10.0.0.11 port 49789 connected to 10.0.0.12 port 5201
[10] local 10.0.0.11 port 49790 connected to 10.0.0.12 port 5201
[12] local 10.0.0.11 port 49791 connected to 10.0.0.12 port 5201
[14] local 10.0.0.11 port 49792 connected to 10.0.0.12 port 5201
[16] local 10.0.0.11 port 49793 connected to 10.0.0.12 port 5201
[18] local 10.0.0.11 port 49794 connected to 10.0.0.12 port 5201
[ ID] Interval      Transfer      Bandwidth
[ 4]  0.00-1.01    sec  44.9 MBytes  373 Mbits/sec
[ 6]  0.00-1.01    sec  78.4 MBytes  651 Mbits/sec
[ 8]  0.00-1.01    sec   130 MBytes 1.08 Gbits/sec
[10]  0.00-1.01    sec  75.5 MBytes  627 Mbits/sec
[12]  0.00-1.01    sec  29.0 MBytes  241 Mbits/sec
[14]  0.00-1.01    sec  41.0 MBytes  341 Mbits/sec
[16]  0.00-1.01    sec  32.4 MBytes  269 Mbits/sec
[18]  0.00-1.01    sec  33.8 MBytes  280 Mbits/sec
[SUM] 0.00-1.01    sec  464 MBytes  3.86 Gbits/sec
-----
[ 4]  7.00-8.00    sec  46.1 MBytes  387 Mbits/sec
[ 6]  7.00-8.00    sec  56.5 MBytes  474 Mbits/sec
[ 8]  7.00-8.00    sec   172 MBytes 1.45 Gbits/sec
[10]  7.00-8.00    sec  57.8 MBytes  484 Mbits/sec
[12]  7.00-8.00    sec  45.1 MBytes  378 Mbits/sec
[14]  7.00-8.00    sec  45.9 MBytes  385 Mbits/sec
[16]  7.00-8.00    sec  37.9 MBytes  318 Mbits/sec
[18]  7.00-8.00    sec  58.9 MBytes  494 Mbits/sec
[SUM] 7.00-8.00    sec  521 MBytes  4.37 Gbits/sec
<snip>

```

Under normal circumstances, traffic between the servers should only forward over the member interfaces, and not over the vPC Peer Link. Note that in the below output, interfaces E1/3 - 4 (the vPC Peer Link) have little to no utilization.

```

NSK1# show interface e1/1 - 4 | include "Ethernet1/|output rate"
Ethernet1/1 is up
  30 seconds output rate 4838040 bits/sec, 9414 packets/sec
  input rate 997.33 Mbps, 82.25 Kpps; output rate 2.64 Mbps, 5.09 Kpps
Ethernet1/2 is up
  30 seconds output rate 1792037328 bits/sec, 147886 packets/sec
  input rate 2.66 Mbps, 5.13 Kpps; output rate 997.33 Mbps, 82.25 Kpps
Ethernet1/3 is up
  30 seconds output rate 1264 bits/sec, 1 packets/sec
  input rate 776 bps, 1 pps; output rate 800 bps, 1 pps
Ethernet1/4 is up
  30 seconds output rate 72 bits/sec, 0 packets/sec
  input rate 336 bps, 0 pps; output rate 152 bps, 0 pps

NSK2# show interface e1/1 - 4 | include "Ethernet1/|output rate"
Ethernet1/1 is up
  30 seconds output rate 4226080 bits/sec, 8223 packets/sec
  input rate 1.19 Gbps, 98.22 Kpps; output rate 2.53 Mbps, 4.88 Kpps
Ethernet1/2 is up
  30 seconds output rate 2177190128 bits/sec, 179693 packets/sec
  input rate 2.53 Mbps, 4.88 Kpps; output rate 1.19 Gbps, 98.22 Kpps
Ethernet1/3 is up
  30 seconds output rate 1008 bits/sec, 1 packets/sec
  input rate 760 bps, 1 pps; output rate 760 bps, 1 pps
Ethernet1/4 is up
  30 seconds output rate 360 bits/sec, 0 packets/sec
  input rate 152 bps, 0 pps; output rate 320 bps, 0 pps

```

Traffic only forwards over the vPC Peer Link in the case that a vPC Member Port fails and the end host becomes single attached (i.e. an Orphan Port).

```

NSK1# config t ; int e1/2 ; shut ; end
Enter configuration commands, one per line. End with CNTL/Z.
NSK1#
2017 Jan 27 20:12:35 NSK1 %ETH_PORT_CHANNEL-5-PORT_DOWN: port-channel2: Ethernet1/2 is down
2017 Jan 27 20:12:35 NSK1 %ETH_PORT_CHANNEL-5-PORT_DOWN: port-channel2: port-channel2 is down
2017 Jan 27 20:12:35 NSK1 %ETH_PORT_CHANNEL-5-FOP_CHANGED: port-channel2: first operational port changed from Ethernet1/2 to none
2017 Jan 27 20:12:35 NSK1 %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-channel2 is down (No operational members)
2017 Jan 27 20:12:35 NSK1 %ETHPORT-5-IF_DOWN_CFG_CHANGE: Interface Ethernet1/2 is down(Config change)
2017 Jan 27 20:12:35 NSK1 %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-channel2 is down (No operational members)
2017 Jan 27 20:12:35 NSK1 %ETHPORT-5-IF_DOWN_ADMIN_DOWN: Interface Ethernet1/2 is down (Administratively down)
2017 Jan 27 20:12:35 NSK1 %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty by brian on 192.168.0.1@pts/0

NSK1# clear counters
NSK1#
2017 Jan 27 20:12:41 NSK1 %LIBIFMGR-5-ALL_COUNTERS_CLEARED: All interface counters cleared by user

NSK1# show interface e1/1 - 4 | include "Ethernet1/|output rate"
Ethernet1/1 is up
  30 seconds output rate 783416 bits/sec, 1520 packets/sec
  input rate 1.72 Gbps, 142.01 Kpps; output rate 783.42 Kbps, 1.52 Kpps
Ethernet1/2 is down (Administratively down)
  30 seconds output rate 358644440 bits/sec, 29592 packets/sec
  input rate 787.62 Kbps, 1.53 Kpps; output rate 358.64 Mbps, 29.59 Kpps
Ethernet1/3 is up
  30 seconds output rate 563888320 bits/sec, 45931 packets/sec
  input rate 3.20 Kbps, 1 pps; output rate 563.89 Mbps, 45.93 Kpps
Ethernet1/4 is up
  30 seconds output rate 816158536 bits/sec, 66475 packets/sec
  input rate 832 bps, 0 pps; output rate 816.16 Mbps, 66.47 Kpps

```



