128 TECHNOLOGY

HOW-TO: SHARE ONE PHYSICAL LINK FOR A HA SYNC AND INTER-NODE FABRIC

Abstract

In situations where physical ports are scarce or limited in number for whatever reason, management of the router can be achieved via one of the forwarding plane interfaces (either LAN or WAN). However, the need for HA sync between the two nodes still requires a dedicated physical connection. This article provides a way to share one set of physical interfaces on an HA pair of nodes for both the dog-leg and HA Sync connections.

18 April 2019

INTENDED AUDIENCE

This How-To Guide is intended for network administrators and architects. It presumes that the reader has a working 128T router and is familiar with basic administration of the 128T Networking Platform.

PREREQUISITES

We assume the 128T router is using inter-node HA, as shown in the diagram here. There are redundant interfaces on the LAN side, and non-redundant interfaces on the WAN side. Inbound (SSH, Web) and outbound (conductor) management traffic will use WAN forwarding interface(s) – please see the following article for further details on enabling this functionality: How-To: Share Interface for Management and Data Planes. Also, there is an HA Sync link provided by a straight cable. The starting point finally assumes that there is a separate "dog-leg" fabric link. The purpose of the exercise below is to move the dog-leg over to the same physical link that is also being used by the HA Sync traffic.

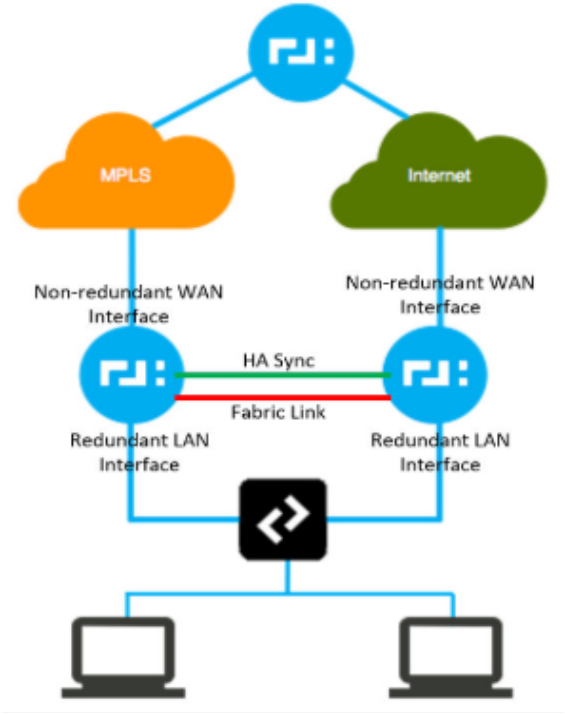

PROCEDURE

The outline of the procedure is quite straightforward:

1. On each node create a tagged interface from a parent interface that is currently being used by HA Sync, using an "ifcfg" script. Put both nodes' tagged interfaces on the same VLAN (we used 101 for purposes of this article).
2. In the 128T router configuration, for each node, change the existing Fabric Link device-interface, or create a new one, with type "bridged" and specify the target to be the tagged interface created in step 1 above.
3. If changing from an existing separate dog-leg configuration, no additional steps are needed, just commit the changes. If building a new config, create and configure network-interfaces inside the bridged device interfaces, just like you would if it were a separate fabric link, and then commit your changes.

STEP 1

Here is an example of an HA Sync interface ifcfg script on each of the two nodes in directory /etc/sysconfig/network-scripts/:

- Node 1 (ifcfg-enp1s0):

```
BOOTPROTO=static
DEFROUTE=no
DEVICE=enp1s0
GATEWAY=30.254.255.2
HWADDR=0C:C4:7A:7F:80:A4
IPADDR=30.254.255.1
IPV4_FAILURE_FATAL=no
IPV6INIT=no
MTU=1500
NAME=enp1s0
NM_CONTROLLED=yes
ONBOOT=yes
PREFIX=30
TYPE=Ethernet
ZONE=trusted
```

- Node 2 (ifcfg-enp1s0):

```
BOOTPROTO=static
DEFROUTE=no
DEVICE=enp1s0
GATEWAY=30.254.255.1
HWADDR=00:ec:ac:cd:09:35
IPADDR=30.254.255.2
IPV4_FAILURE_FATAL=no
IPV6INIT=no
MTU=1500
NAME=enp1s0
ONBOOT=yes
PREFIX=30
TYPE=Ethernet
NM_CONTROLLED=yes
ZONE=trusted
```

And here is the script for the tagged interfaces to be used on both nodes, assuming the name of the parent interface is the same, "enp1s0" (ifcfg-enp1s0.101):

```
BOOTPROTO=none
DEFROUTE=no
DEVICE=enp1s0.101
IPV4_FAILURE_FATAL=no
IPV6INIT=no
MTU=1500
NAME=enp1s0
NM_CONTROLLED=no
ONBOOT=yes
PHYSDEV=enp1s0
TYPE=Ethernet
VLAN=yes
VLAN_ID=101
ZONE=trusted
```

Issuing the *ifup enp1s0.101* command after the scripts are created will bring up the tagged interface on each node. Now the *ip a* command produces the following output for the relevant interfaces (example shown is for Node-2):

```
1: enp1s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    link/ether 00:ec:ac:cd:09:35 brd ff:ff:ff:ff:ff:ff
    inet 30.254.255.2/30 brd 30.254.255.3 scope global enp1s0
       valid_lft forever preferred_lft forever
    inet6 fe80::2ec:acff:fecd:935/64 scope link
       valid_lft forever preferred_lft forever
2: enp1s0.101@enp1s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP qle
n 1000
    link/ether 00:ec:ac:cd:09:35 brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2ec:acff:fecd:935/64 scope link
       valid_lft forever preferred_lft forever
Steps 2-3
```

STEPS 2-3

Now we will create the device-interface configuration on each 128T node with the new fields/values highlighted:

- Node-1

```
device-interface           2
    name                   2
    description            "HA Router Inter-Node Fabric Link Interface"
    type                   bridged
    target-interface       enp1s0.101
    bridge-name            enp101_br

    enabled                true

    network-interface      Inter-Node-Fabric
        name                     Inter-Node-Fabric
        global-id                6
        type                     fabric
        inter-router-security    internal
        source-nat               false
        mtu                      1500

        address                  169.254.255.0
            ip-address       169.254.255.0
            prefix-length    31
        exit
        icmp                     allow
        dhcp                     disabled
    exit
exit
```

- Node-2

```
device-interface           2
    name                   2
    description            "HA Router Inter-Node Fabric Link Interface"
    type                   bridged
    target-interface       enp1s0.101
    bridge-name            enp101_br

    enabled                true

    network-interface      Inter-Node-Fabric
        name                     Inter-Node-Fabric
        global-id                7
        type                     fabric
        inter-router-security    internal
        source-nat               false
        mtu                      1500

        address                  169.254.255.1
            ip-address       169.254.255.1
            prefix-length    31
        exit
        icmp                     allow
        dhcp                     disabled
    exit
exit
```

At this point we will commit the changes, and this concludes the migration from two links to a single link for both HA Sync and Inter-Node Fabric.

VERIFYING OPERATION

After the changes have been committed, dropping down to Linux shell and issuing ip a command again will produce a slightly different result, which now shows the new bridge and KNI interfaces we just created:

```
1: enp1s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1
000
    link/ether 00:ec:ac:cd:09:35 brd ff:ff:ff:ff:ff:ff
    inet 30.254.255.2/30 brd 30.254.255.3 scope global enp1s0
       valid_lft forever preferred_lft forever
    inet6 fe80::2ec:acff:fecd:935/64 scope link
       valid_lft forever preferred_lft forever
2: kni2: <BROADCAST,MULTICAST,PROMISC,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast master en
p101_br state UNKNOWN qlen 1000
    link/ether 1e:70:04:f6:be:ca brd ff:ff:ff:ff:ff:ff
3: enp101_br: <BROADCAST,MULTICAST,PROMISC,UP,LOWER_UP> mtu 1500 qdisc noqueue state U
P qlen 1000
    link/ether 1e:70:04:f6:be:ca brd ff:ff:ff:ff:ff:ff
    inet6 fe80::1c70:4ff:fef6:beca/64 scope link
       valid_lft forever preferred_lft forever
4: enp1s0.101@enp1s0: <BROADCAST,MULTICAST,PROMISC,UP,LOWER_UP> mtu 1500 qdisc noqueue
 master enp101_br state UP qlen 1000
    link/ether 1e:70:04:f6:be:ca brd ff:ff:ff:ff:ff:ff
    inet6 fe80::1c70:4ff:fef6:beca/64 scope link
       valid_lft forever preferred_lft forever
```

And here is PCLI output of the *show device-interface summary* and *show network-interface* commands (notice the device 2 is type "bridged").

```
admin@Node-2.GeneShRouter# show device-interface summary
Thu 2018-05-03 19:02:12 UTC

=========== ============== ==================== ==================== ================
==
 Name         Admin Status   Operational Status   Redundancy Status   MAC Address
=========== ============== ==================== ==================== ================
==
 Node-2:11    up             up                   standby             00:ec:ac:cd:09:3
6
 Node-2:100   up             up                   non-redundant       00:ec:ac:cd:09:3
7
 Node-2:2     up             up                   non-redundant       00:ec:ac:cd:09:3
5
 Node-2:24    up             up                   non-redundant       06:de:13:9d:93:0
f

admin@Node-2.GeneShRouter# show network-interface
Thu 2018-05-03 19:03:06 UTC

============= ======== ======== ============== ====== ============= ========== =====
============= =============== ========== ============== =============
 Router         Node     Device   Name            VLAN   Device Type   DHCP       Addr
ess            Gateway         Hostname   Admin Status   Oper Status
============= ======== ======== ============== ====== ============= ========== =====
============= =============== ========== ============== =============
 GeneShRouter   Node-2       11   PrivateLAN         0   ethernet      disabled   192.
168.1.1/24     --              --         up             up
 GeneShRouter   Node-2      100   WAN_Interface      0   ethernet      disabled   192.
168.15.92/24   192.168.15.10   --         up             up
 GeneShRouter   Node-2        2   Inter-Node-HA      0   bridged       disabled   169.
254.255.1/31   --              --         up             up
 GeneShRouter   Node-2       24   mgmt-outbound      0   host          disabled   169.
254.2.1/30     169.254.2.2     --         up             up
```