



Deployment Guide | Technical Doc

VYOS HIGH AVAILABILITY BORDER ROUTER

Index:

Overview	3
Lab Setup	3
Border Router (1) Configuration	4
Interfaces	4
VRRP	4
BFD	4
prefix-list and route-map	5
eBGP	6
iBGP	6
Border Router (2) Configuration	7
Interfaces	7
VRRP	7
BFD	7
prefix-list and route-map	8
eBGP	8
iBGP	8
Firewall (Optional)	9
Validation	9
Failure Scenarios	11
Internal link failure	12
Router failure	13



Overview

In this guide, we'll walk through the process of setting up high availability (HA) in a VyOS border router environment. High availability isn't just about using multiple providers—it's about maintaining network connectivity even if a router or link fails. To achieve this, our design includes the following components: VRRP for LAN-side redundancy and WAN failover using two VyOS routers connected to two different ISPs. For routing redundancy, we'll leverage both internal and external BGP, as well as BFD for rapid failure detection.

⚠️ This guide was validated against VyOS 1.5 LTS (Circinus) as of July 2026, and has been tested and verified.

Lab Setup

We have two VyOS routers in this setup: Router 1 connects to ISP 1, and Router 2 connects to ISP 2. On the LAN side, both routers use VRRP to provide a virtual default gateway via a switch, ensuring that internal devices can reach the gateway regardless of which router is currently active.




Figure-1: Lab setup topology

On the WAN side, both routers establish BGP sessions with their respective ISPs, which advertise a default route to the border routers. Additionally, the two VyOS routers form an internal BGP session to support failover routing. Router 1 is configured as the preferred exit point, even in the event of downstream link failure. If Router 1 becomes unavailable, traffic will automatically fail over to ISP 2 via Router 2.

Border Router (1) Configuration

Interfaces

LAN interface:

```
set interfaces ethernet eth0 description to_LAN
set interfaces ethernet eth0 address 192.168.10.2/24
```

WAN Interfaces:

```
set interfaces ethernet eth1 description to_BR-2
set interfaces ethernet eth1 address 172.16.0.1/30
set interfaces ethernet eth4 description to_ISP-1
set interfaces ethernet eth4 address 10.100.1.2/30
```

VRRP

The address **192.168.10.1** is configured under VRRP group **ha1**. This IP will act as the default gateway for LAN hosts, allowing them to access external networks, including the internet.

The hello source address is used by Router 1 to send hello packets to Router 2, allowing it to confirm that Router 1 is still active. The peer address refers to Router 2's IP address. A priority value of 200 designates Router 1 as the master. The **vrid** simply defines the VRRP group ID.

```
set high-availability vrrp group ha1 address 192.168.10.1/24
set high-availability vrrp group ha1 hello-source-address 192.168.10.2
set high-availability vrrp group ha1 peer-address 192.168.10.3
set high-availability vrrp group ha1 interface eth0
set high-availability vrrp group ha1 priority 200
set high-availability vrrp group ha1 preempt-delay 30
set high-availability vrrp group ha1 vrid 1
set high-availability vrrp sync-group sync1 member ha1
commit ; save
```

With this configuration, Router 1 will automatically take back the master role when it comes back online (failback). If you want to disable this behavior, you can use the **no-preempt** option, which prevents automatic failback once a router recovers.

Finally, a VRRP sync group is created to allow both routers to track each other's connectivity and maintain consistency in failover behavior.

BFD

Before we configure BGP, we'll need to configure BFD, a prefix-list and a route-map.

By default, BGP convergence can take several minutes due to its built-in timers. To accelerate failover detection, BFD (Bidirectional Forwarding Detection) is used to continuously monitor connectivity by



exchanging control packets at regular intervals.

```
set protocols bfd profile ISP interval multiplier 3
set protocols bfd profile ISP interval transmit 200
set protocols bfd profile ISP interval receive 200
```

In this setup, BFD is configured to send and expect a packet every 200 milliseconds. The multiplier value of 3 means that if three consecutive packets are missed, the link will be considered down, allowing for much faster failover than standard BGP timers.

prefix-list and route-map

The first prefix list is configured to permit only the default route received via BGP from our ISPs.

```
set policy prefix-list defaultroute description "Default Route"
set policy prefix-list defaultroute rule 10 action permit
set policy prefix-list defaultroute rule 10 prefix 0.0.0.0/0
```

The next prefix list will be used to advertise our internal network to the ISPs, allowing them to route return traffic properly. This route-map controls the advertisement of our publicly accessible internal prefix to the providers. In the lab scenario, this corresponds to the prefix assigned to the workstation. In a real-world deployment, it would usually be a public IP block used for NATing the internal network.

```
set policy prefix-list lan description "internal network"
set policy prefix-list lan rule 10 action permit
set policy prefix-list lan rule 10 prefix 192.168.10.0/24
```

The **import_isp1** route-map is used with the **defaultroute** prefix list to accept only the default route from the provider.

```
set policy route-map import_isp1 description "accept default route"
set policy route-map import_isp1 rule 10 match ip address prefix-list defaultroute
set policy route-map import_isp1 rule 10 action permit
```

The **export_isp1** route-map is used together with the **lan** prefix list to advertise only our internal prefix to the provider.

```
set policy route-map export_isp1 description "route-map to ISP-1"
set policy route-map export_isp1 rule 10 description "advertise internal network"
set policy route-map export_isp1 rule 10 match ip address prefix-list lan
set policy route-map export_isp1 rule 10 action permit
set policy route-map export_isp1 rule 100 description "default deny all"
set policy route-map export_isp1 rule 100 action deny
```

eBGP

eBGP will be used to establish communication with the provider. We begin by configuring our autonomous system number (ASN), and then redistribute connected routes. This setup supports failover routing between Router 1 and Router 2, and also enables the advertisement of our internal network to the ISP.

```
set protocols bgp system-as 65000
set protocols bgp address-family ipv4-unicast redistribute connected route-map export_isp1
```

Now we configure the ISP as a BGP neighbor. We apply the import and export policies using route-maps, which define the prefixes we accept from and advertise to our peer.

```
set protocols bgp neighbor 10.100.1.1 description ISP-1
set protocols bgp neighbor 10.100.1.1 address-family ipv4-unicast route-map export export_isp1
set protocols bgp neighbor 10.100.1.1 address-family ipv4-unicast route-map import import_isp1
set protocols bgp neighbor 10.100.1.1 remote-as 65001
set protocols bgp neighbor 10.100.1.1 bfd profile ISP
```

You can optionally configure a password for the BGP session.

```
set protocols bgp neighbor 10.100.1.1 password <text> # optional
```

iBGP

We begin by configuring iBGP between the two routers to enable failover routing. Router 1 is designated as the preferred path regardless of the downstream ISP link status. Traffic will only switch to Router 2 if Router 1 becomes unreachable.

As part of this setup, we'll create a prefix list and a route-map to control the export of the default route. The route-map will reference the **defaultroute** prefix list created earlier.

In this lab, we're simulating a scenario where ISP-1 offers better connectivity—either higher bandwidth, lower latency, or reduced cost. To prioritize this link, we modify the BGP local preference to 200 for routes received from ISP-1 and advertise them to Router 2.

This ensures that traffic continues to prefer the ISP-1 path, even if VRRP fails over to Router 2. By default, BGP assigns a local preference of 100. Increasing it to 200 signals Router 2 to prefer the route through Router 1 over using its own ISP connection.

```
set policy route-map export_br2 description "route-map to border router 2"
set policy route-map export_br2 rule 10 action permit
set policy route-map export_br2 rule 10 match ip address prefix-list defaultroute
set policy route-map export_br2 rule 10 set local-preference 200
set policy route-map export_br2 rule 100 action permit
```

Now we configure the iBGP neighbor. We apply the export policy using route-map **export_br2**.

```
set protocols bgp neighbor 172.16.0.2 description BR-2
set protocols bgp neighbor 172.16.0.2 address-family ipv4-unicast nexthop-self
set protocols bgp neighbor 172.16.0.2 address-family ipv4-unicast route-map export export_br2
set protocols bgp neighbor 172.16.0.2 remote-as 65000
set protocols bgp neighbor 172.16.0.2 bfd profile ISP
```

Border Router (2) Configuration

Router 2 is configured in a similar way to Router 1.

Interfaces

LAN interface:

```
set interfaces ethernet eth0 description to_LAN
set interfaces ethernet eth0 address 192.168.10.3/24
```

WAN Interfaces:

```
set interfaces ethernet eth1 description to_BR-1
set interfaces ethernet eth1 address 172.16.0.2/30
set interfaces ethernet eth4 description to_ISP-2
set interfaces ethernet eth4 address 10.100.2.2/30
```

VRRP

Same configuration as Router 1, with addresses flipped accordingly. The main change here is the priority value. Since Router 2 is intended to operate as the secondary (not the master), we set its priority to 100.

```
set high-availability vrrp group ha1 address 192.168.10.1/24
set high-availability vrrp group ha1 hello-source-address 192.168.10.3
set high-availability vrrp group ha1 peer-address 192.168.10.2
set high-availability vrrp group ha1 interface eth0
set high-availability vrrp group ha1 priority 100
set high-availability vrrp group ha1 vrid 1
set high-availability vrrp sync-group sync1 member ha1
commit ; save
```

BFD

```
set protocols bfd profile ISP interval multiplier 3
set protocols bfd profile ISP interval transmit 200
set protocols bfd profile ISP interval receive 200
```



prefix-list and route-map

The prefix-list **defaultroute** is used to receive only a default-route from our ISP.

```
set policy prefix-list defaultroute description "Default Route"
set policy prefix-list defaultroute rule 10 action permit
set policy prefix-list defaultroute rule 10 prefix 0.0.0.0/0
```

We'll also create a prefix list to match our internal network.

```
set policy prefix-list lan description "internal network"
set policy prefix-list lan rule 10 action permit
set policy prefix-list lan rule 10 prefix 192.168.10.0/24
```

Just like on Router 1, we'll advertise our internal network to the ISP and deny all other outbound route advertisements.

```
set policy route-map export_isp2 description "route-map to ISP-2"
set policy route-map export_isp2 rule 10 description "advertise internal network"
set policy route-map export_isp2 rule 10 match ip address prefix-list lan
set policy route-map export_isp2 rule 10 action permit
set policy route-map export_isp2 rule 100 description "default deny all"
set policy route-map export_isp2 rule 100 action deny
```

Only the default route will be accepted from the ISP.

```
set policy route-map import_isp2 description "accept default route"
set policy route-map import_isp2 rule 10 match ip address prefix-list defaultroute
set policy route-map import_isp2 rule 10 action permit
```

eBGP

```
set protocols bgp system-as 65000
set protocols bgp address-family ipv4-unicast redistribute connected route-map export_isp2
set protocols bgp neighbor 10.100.2.1 address-family ipv4-unicast route-map export export_isp2
set protocols bgp neighbor 10.100.2.1 description ISP-2
set protocols bgp neighbor 10.100.2.1 address-family ipv4-unicast route-map import import_isp2
set protocols bgp neighbor 10.100.2.1 remote-as 65002
set protocols bgp neighbor 10.100.2.1 bfd profile ISP
set protocols bgp neighbor 10.100.2.1 password <text> # optional
```

iBGP

No route-map will be applied to the iBGP session, allowing Router 1 to maintain the default route from the ISP as its preferred path. In the absence of import or export statements, all routes are permitted.

```
set protocols bgp neighbor 172.16.0.1 description BR-1
set protocols bgp neighbor 172.16.0.1 address-family ipv4-unicast nexthop-self
set protocols bgp neighbor 172.16.0.1 remote-as 65000
set protocols bgp neighbor 172.16.0.1 bfd profile ISP
```

Firewall (Optional)

We can optionally configure a firewall filter on both Router 1 and Router 2. The default action is set to drop all traffic.

Rule 10 permits return traffic back into our network, while Rule 20 allows all outbound traffic originating from inside the network.

```
set firewall ipv4 forward filter default-action drop
set firewall ipv4 forward filter rule 10 description "Allow return traffic through the router"
set firewall ipv4 forward filter rule 10 state related
set firewall ipv4 forward filter rule 10 state established
set firewall ipv4 forward filter rule 10 action accept
set firewall ipv4 forward filter rule 20 description "Allow all traffic from LAN interface"
set firewall ipv4 forward filter rule 20 inbound-interface name eth0
set firewall ipv4 forward filter rule 20 action accept
```

Validation

1. VRRP. We verify that VRRP is working.

```
vyos@BR-1# run show vrrp
Name      Interface  VRID  State      Priority  Last Transition
-----
ha1       eth0       1     MASTER     200      17m10s
[edit]
vyos@BR-1#
```

```
vyos@BR-2# run show vrrp
Name      Interface  VRID  State      Priority  Last Transition
-----
ha1       eth0       1     BACKUP     100      5m15s
[edit]
vyos@BR-2#
```

The output confirms that VRRP is running on group ha1 using interface eth0. It also shows the configured VRID. BR-1 is currently in the MASTER state, while BR-2 is in BACKUP. The priority values are set to 200 and 100, respectively.



1. Ping test to default gateway from LAN host.

```
host-1> ping 192.168.10.1

84 bytes from 192.168.10.1 icmp_seq=1 ttl=64 time=1.019 ms
84 bytes from 192.168.10.1 icmp_seq=2 ttl=64 time=1.023 ms
84 bytes from 192.168.10.1 icmp_seq=3 ttl=64 time=0.888 ms
84 bytes from 192.168.10.1 icmp_seq=4 ttl=64 time=0.815 ms
84 bytes from 192.168.10.1 icmp_seq=5 ttl=64 time=0.764 ms

host-1>
```

1. Verify BGP.

```
vyos@BR-1# run show bgp summary

IPv4 Unicast Summary (VRF default):
BGP router identifier 192.168.10.2, local AS number 65000 vrf-id 0
BGP table version 19
RIB entries 7, using 672 bytes of memory
Peers 2, using 40 KiB of memory
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd	PfxSnt	Desc
10.100.1.1	4	65001	225	224	19	0	0	01:38:53	1	1	ISP-1
172.16.0.2	4	65000	204	203	19	0	0	01:38:54	3	4	BR-2

```
Total number of neighbors 2
[edit]
vyos@BR-1#
```

```
vyos@BR-2# run show bgp summary

IPv4 Unicast Summary (VRF default):
BGP router identifier 192.168.10.3, local AS number 65000 vrf-id 0
BGP table version 12
RIB entries 7, using 672 bytes of memory
Peers 2, using 40 KiB of memory
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd	PfxSnt	Desc
10.100.2.1	4	65002	225	221	12	0	0	03:32:10	1	1	ISP-2
172.16.0.1	4	65000	306	254	12	0	0	01:39:25	4	3	BR-1

```
Total number of neighbors 2
[edit]
vyos@BR-2#
```

1. Ping test to "internet" (10.0.0.1)

```
host-1> ping 10.0.0.1

84 bytes from 10.0.0.1 icmp_seq=1 ttl=62 time=2.190 ms
84 bytes from 10.0.0.1 icmp_seq=2 ttl=62 time=2.986 ms
84 bytes from 10.0.0.1 icmp_seq=3 ttl=62 time=2.723 ms
84 bytes from 10.0.0.1 icmp_seq=4 ttl=62 time=3.230 ms
84 bytes from 10.0.0.1 icmp_seq=5 ttl=62 time=2.612 ms

host-1>
```



1. Ping test to default gateway from LAN host.

```
host-1> ping 192.168.10.1
84 bytes from 192.168.10.1 icmp_seq=1 ttl=64 time=1.019 ms
84 bytes from 192.168.10.1 icmp_seq=2 ttl=64 time=1.023 ms
84 bytes from 192.168.10.1 icmp_seq=3 ttl=64 time=0.888 ms
84 bytes from 192.168.10.1 icmp_seq=4 ttl=64 time=0.815 ms
84 bytes from 192.168.10.1 icmp_seq=5 ttl=64 time=0.764 ms
host-1>
```

Failure Scenarios

In this section, we'll simulate two failure scenarios. First, a failure on the link connecting to the LAN. Second, a failure of the primary router. Under normal conditions, the traffic path follows the one shown in Figure-2.




Figure-2: Normal condition traffic path

The packet trace confirms the traffic path.

```
host-1> trace 10.0.0.1
trace to 10.0.0.1, 8 hops max, press Ctrl+C to stop
 1  192.168.10.2    0.896 ms  0.668 ms  0.584 ms
 2  10.100.1.1     1.362 ms  1.203 ms  1.264 ms
 3  *10.0.0.1     2.846 ms
host-1>
```

Internal link failure

When a failure occurs on the link connecting to the LAN, traffic follows the path shown in Figure 3.




Figure-3: Internal link failure traffic path

On router one we can see the state has changed to **FAULT**.

```
vyos@BR-1# run show vrrp
Name      Interface  VRID  State      Priority  Last Transition
-----
ha1       eth0       1     FAULT      200      7s
[edit]
vyos@BR-1#
```

On router two, we can see the state has changed to **MASTER**.

```
vyos@BR-2# run show vrrp
Name      Interface  VRID  State      Priority  Last Transition
-----
ha1       eth0       1     MASTER     100      17s
[edit]
vyos@BR-2#
```

The ping test is still successful, and the packet trace confirms the traffic path.

```
host-1> ping 10.0.0.1

84 bytes from 10.0.0.1 icmp_seq=1 ttl=61 time=3.134 ms
84 bytes from 10.0.0.1 icmp_seq=2 ttl=61 time=3.201 ms
84 bytes from 10.0.0.1 icmp_seq=3 ttl=61 time=3.746 ms
84 bytes from 10.0.0.1 icmp_seq=4 ttl=61 time=3.300 ms
84 bytes from 10.0.0.1 icmp_seq=5 ttl=61 time=3.342 ms

host-1>
```

```

host-1> trace 10.0.0.1
trace to 10.0.0.1, 8 hops max, press Ctrl+C to stop
 1  192.168.10.3    0.553 ms  0.408 ms  0.333 ms    # BR-2 interface
 2  172.16.0.1     1.672 ms  1.106 ms  0.988 ms    # BR-1 HA interface
 3  10.100.1.1     1.308 ms  1.003 ms  0.862 ms    # ISP-1 interface
 4  *10.0.0.1     1.347 ms
host-1>

```

Router failure

In the event of a failure on BR-1, traffic seamlessly transitions to the path depicted in Figure-4.




Figure-4: Primary router failure traffic path

On router two, we can see the state has changed to **MASTER**.

```

vyos@BR-2# ru show vrrp
Name      Interface  VRID  State      Priority  Last Transition
-----
ha1       eth0       1     MASTER     100      8s
[edit]
vyos@BR-2#

```

The ping test is still successful, and the packet trace confirms the traffic path.

```

host-1> ping 10.0.0.1

84 bytes from 10.0.0.1 icmp_seq=1 ttl=62 time=3.098 ms
84 bytes from 10.0.0.1 icmp_seq=2 ttl=62 time=3.031 ms
84 bytes from 10.0.0.1 icmp_seq=3 ttl=62 time=2.498 ms
84 bytes from 10.0.0.1 icmp_seq=4 ttl=62 time=2.954 ms
84 bytes from 10.0.0.1 icmp_seq=5 ttl=62 time=2.165 ms

host-1>

```

```
host-1> trace 10.0.0.1
trace to 10.0.0.1, 8 hops max, press Ctrl+C to stop
 1  192.168.10.3    0.982 ms  0.855 ms  0.649 ms    # BR-2 interface
 2  10.100.2.1     1.235 ms  1.338 ms  1.249 ms    # ISP-2 interface
 3  *10.0.0.1      2.451 ms
host-1>
```

