

How to use only one network interface

Abstract

The LoadBalancer is built with two network interfaces in mind. This has been made in order to provide a better network architecture: administrations services are on an internal (usually private) network interface, and all user-defined services along with Virtual IPs are on a dedicated public interface.

The drawback of this approach is that you have to define two different network mask and some route definition. When this is not possible due to business rules, it's possible to configure the load balancer to act as if it was with only one interface.

Before jumping to the configuration, here are some things to keep in mind to understand the changes:

- cluster replica is done through SSH connection, using the IP assigned on the internal interface;
- public IPs aren't a strict requirement, if not useful for publishing services;
- Virtual IPs are bound to the public interface, and are managed through ARP messages.

How to configure

The first thing to do, is to make sure that both the interfaces of the cluster nodes are connected to the same "physical" (i.e. of the Hypervisor) network. Since we are trying to melt the two interfaces, the physical communication must be shared.

The configuration for the nodes could be the following (e.g. network 192.168.1.0/24 with gateway 192.168.1.254/24):

Configuration	Master node	Backup node
Private IP	192.168.1.12/24	192.168.1.34/24
Private GW	default 192.168.1.254/24	default 192.168.1.254/24

Configuration	Master node	Backup node
Public IP	127.1.0.1/8	127.1.0.2/8
Public GW	not set	not set
Virtual IP	192.168.1.56	

And that's all. Note that local IPs are assigned to public interfaces and that internal interfaces are "published" just like the virtual IPs.

×**Security** By exposing the internal interface you are responsible for firewalling the application. All traffic between the cluster nodes need to be allowed, but access to ports 22 and 1443 must be limited to administrator IPs or network