November 27, 2018

AUTOMATIC BINDING OF VIRTUAL INTERNET PROTOCOL ADDRESS TO THE APPROPRIATE INTERFACE

Ravi Kumar Vadapalli

Srinivas Jakkam Shivaji

Kousik Nandy

AUTOMATIC BINDING OF VIRTUAL INTERNET PROTOCOL ADDRESS TO THE APPROPRIATE INTERFACE

AUTHORS:
Ravi Kumar Vadapalli
Srinivas Jakkam Shivaji
Kousik Nandy

ABSTRACT

Techniques are described herein for providing an innovative configuration model that requires a user to configure Virtual Internet Protocol (VIP) addresses in just one place, thus avoiding the possibilities of configuration error and also making the network infrastructure easy to deploy. For example, to configure four VIP addresses in a three-node cluster, state-of-the-art configuration models would require the user to touch/configure twelve values spanning across three appliances, which is cumbersome and error prone. By contrast, the configuration model described herein has just one touch point in the whole cluster.

DETAILED DESCRIPTION

A cluster is a distributed system made of multiple appliances to provide High Availability (HA). The cluster is perceived as a single logical entity to its user, and individual members of the cluster are not exposed to the users. Therefore, even if individual members of the cluster may have Internet Protocol (IP) addresses of their own, the cluster itself is assigned an IP address for users to communicate to it. Since this IP address does not belong to any physical member but instead the logical entity, it is called a Virtual IP (VIP) address. As a part of HA, the cluster's VIP address must also be highly available. This IP address is used by the external devices (e.g., identity services entity, routers, switches, etc.) to reach the cluster. If the cluster IP address is not highly available, the backend services, although are highly available, are not highly reachable. To achieve the HA for IP addresses, a network infrastructure may use a Virtual IP (VIP) address (e.g., managed by keepalived). Since the communication to the VIP must be responded to by some member of the cluster, keepalived may ensure that the VIP address is always configured on one of the available appliances. The configuration of the VIP address gains complexity because there can be multiple interfaces in the appliances. Each appliance

Published by Technical Disclosure Commons, 2018

supports multiple interfaces to be able to connect to different networks through which different external devices are reachable. For example, one interface might be connected to a (colloquially termed) management network that is used to reach out to the identity services entity, policy servers and other management stations. Second and third interfaces might be respectively connected to a (colloquially termed) customer enterprise network and a cloud network. So, with multiple Network Interface Cards (NICs), the network infrastructure must also support multiple VIP addresses in such a way that the cluster is highly reachable on all its connected networks. For example, the cluster may be highly reachable on (say) VIP-1 on the management network and also on (say) VIP-2 on the enterprise network at the same time. Adding to this complexity, there may be future use cases demanding multiple VIP addresses on each NIC. For example, an identity services entity might want to reach the network infrastructure on (say) VIP-1A while a management station might want to reach the network infrastructure on (say) VIP-1B, both via the same management network connected to the same NIC on the network infrastructure appliance. This happens if the identity services entity and management station cannot route packets to VIP-1B and VIP-1A respectively (because of their routing policies or because of the intermediate networks connected to them).

With the possibility of multiple VIP addresses per NIC and per appliance, exposing the same to the user in an easy to configure model is a challenge. Consider the fact that the user should have the option to configure multiple VIPs on each NIC in each appliance. And at the same time, the user configuration should be consistent across the appliances in the network infrastructure cluster. For instance, if the user configures a VIP address on the management NIC of the first appliance and the same VIP address on the enterprise NIC of the second appliance, it results in incorrect networking. The problem is to design a configuration model that offers flexibility of configuring the VIP addresses per business needs while at the same time keeping it easy to configure with lesser chances of making an error.

Provided herein is a configuration model to configure the VIP addresses, by not tying the VIP address configuration to any single NIC nor to any single appliance. The VIP addresses are configured globally at the cluster level in a single cluster level configuration placeholder. This configuration takes just a list of VIP addresses and does not expect any

interface details. The network infrastructure takes the list of VIPs, internally identifies the best NIC that can host each VIP address in each appliance, and internally configures the same in the Keepalived service of each appliance. Effectively, the network infrastructure converts the list of VIP addresses into an interface-to-VIP mapping that Keepalived expects. This mapping is expected from the user in state-of-the-art configuration models.

The network infrastructure uses reverse route lookup to identify the best NIC that can host each VIP. It may assume that the routing in the appliances is symmetric. Asymmetric routing is an advanced form of routing that is not commonly used in appliances, making this assumption a realistic one. The NIC identification may be based on the logic that the best interface to reach a VIP as per the route table of an appliance is the only interface that can host that VIP. This logic may be used to simplify the configuration of the VIP addresses. In some simplified scenarios (but not always), the same logic boils down to the fact that the best interface to configure a VIP address is the interface that is attached to a subnet comprising the VIP address.

This novel configuration model reduces the number of touch points to configure VIP addresses in a network infrastructure cluster. For example, to configure four VIP addresses in a three node network infrastructure cluster, other state-of-the-art configuration models would have twelve touch points spanning across three appliances, whereas the instant configuration model has just one touch point.

Figure 1 below illustrates the state-of-the-art configuration model.

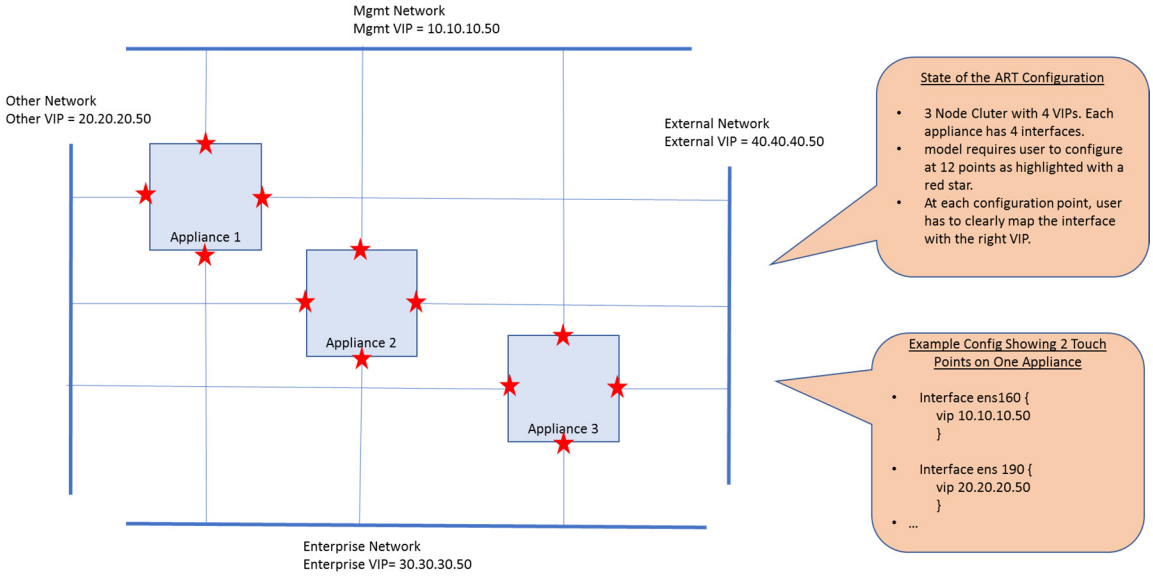

Figure 1

Figure 2 below illustrates the configuration model described herein.

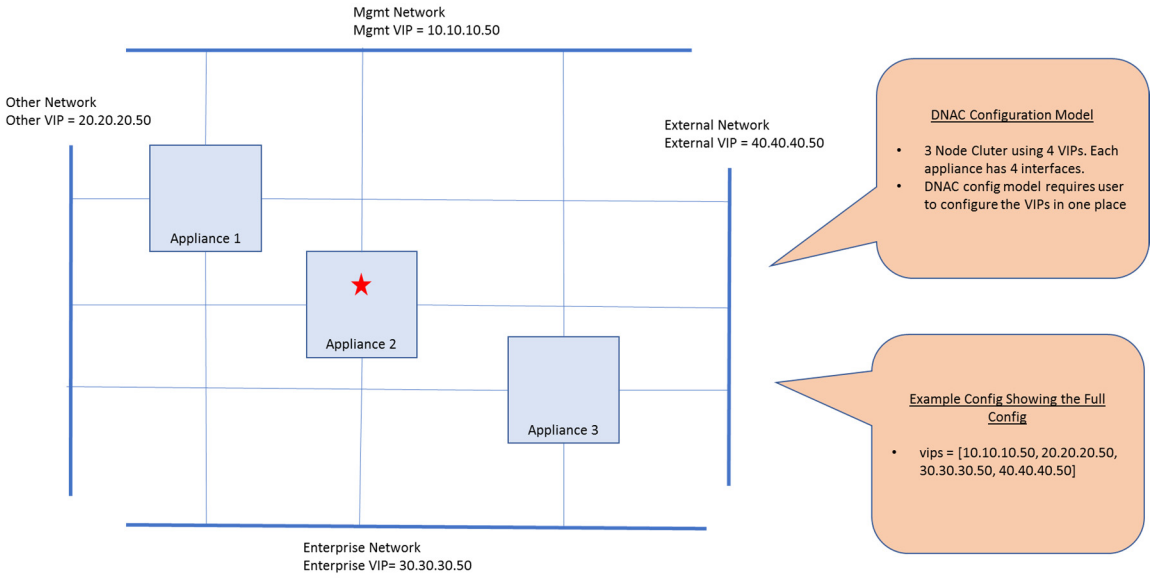

Figure 2

In summary, techniques are described herein for providing an innovative configuration model that requires a user to configure VIP addresses in just one place, thus avoiding the possibilities of configuration error and also making the network infrastructure easy to deploy. For example, to configure four VIP addresses in a three-node cluster, state-of-the-art configuration models would require the user to touch/configure twelve values

spanning across three appliances, which is cumbersome and error prone. By contrast, the configuration model described herein has just one touch point in the whole cluster.