

Step-by-Step Guide: Setup Keepalived



Keepalived on Docker Swarm A Custom Raspberry Pi Setup

Keepalived is an open-source tool used to ensure high availability and redundancy for services. It does this by monitoring the health of your network and services and automatically switching to a backup server if the primary one fails. In this setup, we are using a custom Keepalived Docker image built specifically for Raspberry Pi (ARMv8) by Takabu, a friend of mine. This allows us to implement Keepalived in a Docker Swarm environment on Raspberry Pis.

In our Docker Swarm setup, Keepalived is used to provide a Virtual IP (VIP) that will always point to a healthy node, ensuring seamless failover if a node goes offline.

Priority List and How it Works

Keepalived assigns a priority to each node. The node with the **highest priority** becomes the primary holder of the VIP. If that node goes down, the node with the **next highest priority** takes over the VIP. This ensures that the VIP is always assigned to a functioning node.

Here's how the priority system works in our Docker Swarm:

- **Swarm1 (192.168.0.10):** Priority `100` (lowest priority)
- **Swarm2 (192.168.0.11):** Priority `101`

- **Swarm3 (192.168.0.12):** Priority `102` (highest priority)

The node with priority 102 (`Swarm3`) will be the primary node, and if it goes down, the VIP will switch to the node with priority 101 (`Swarm2`), and so on.

Setting Up the Priority List

Before deploying Keepalived, you need to set the priority for each node. Use the following commands to label each node with its priority:

```
docker node ls # List all nodes in the swarm
```

```
# Assign priority labels to each node:
```

```
docker node update Swarm1 --label-add KEEPALIVED_PRIORITY=100
```

```
docker node update Swarm2 --label-add KEEPALIVED_PRIORITY=101
```

```
docker node update Swarm3 --label-add KEEPALIVED_PRIORITY=102
```

With the priorities set, Keepalived will ensure that the node with the highest priority is always assigned the VIP. If that node fails, Keepalived automatically shifts the VIP to the next highest-priority node, maintaining service availability.

Docker Compose File for Keepalived

```
version: '3.8'
```

```
services:
```

```
  keepalived:
```

```
    image: takabu/public:docker-swarm-keepalived # Custom Keepalived image for Raspberry Pi (ARMv8)
```

```
    volumes:
```

```
      - /var/run/docker.sock:/var/run/docker.sock # Mount Docker socket to interact with Docker API and control
```

```
  the nodes
```

```
  networks:
```

```
    - host # Use the host network for direct access to networking features
```

deploy:

mode: global # Ensures Keepalived runs on all manager nodes

placement:

constraints: [node.role == manager] # Limit deployment to Swarm manager nodes only

environment:

KEEPALIVED_VIRTUAL_IPS: "192.168.0.200" # Virtual IP (VIP) for Keepalived to manage and switch between nodes

networks:

host:

external: true

name: host # Leverage the host network to manage VIP switching and network traffic directly

Explanation of Key Sections:

- **/var/run/docker.sock:/var/run/docker.sock**: This mounts the Docker socket into the Keepalived container, allowing Keepalived to communicate with Docker and control the failover between nodes based on the priority list.
- **KEEPALIVED_VIRTUAL_IPS: "192.168.0.200"**: This environment variable specifies the VIP (192.168.0.200) that Keepalived will manage. This IP will always point to the highest-priority, healthy node.

Revision #8

Created 8 September 2024 12:47:01 by aeoneros

Updated 9 September 2024 09:22:13 by aeoneros