



Exploring Load Balancer Options

NetApp Solutions

NetApp
May 17, 2024

Table of Contents

- Exploring load balancer options 1
- Installing F5 BIG-IP load balancers 1
- Installing MetalLB load balancers 6
- Installing SeeSaw load balancers 8

Exploring load balancer options

An application deployed in Anthos is exposed to the world by a service that is delivered by a load balancer deployed in the Anthos on-prem environment.

The following pages have additional information about load balancer options validated in the Anthos with NetApp solution:

- [Installing F5 BIG-IP load balancers](#)
- [Installing MetalLB load balancers](#)
- [Installing SeeSaw load balancers](#)

Installing F5 BIG-IP load balancers

F5 BIG-IP is an Application Delivery Controller (ADC) that offers a broad set of advanced, production-grade traffic management and security services like L4-L7 load balancing, SSL/TLS offload, DNS, firewall, and more. These services dramatically increase the availability, security, and performance of your applications.

F5 BIG-IP can be deployed and consumed in various ways, including on dedicated hardware, in the cloud, or as a virtual appliance on-premises. Refer to the documentation here to explore and deploy F5 BIG-IP.

F5 BIG-IP was the first of the bundled load balancer solutions available with Anthos On-Prem and was used in a number of the early Anthos Ready partner validations for the Anthos with NetApp solution.



F5 BIG-IP can be deployed in standalone mode or in cluster mode. For the purpose of this validation, F5 BIG-IP was deployed in standalone mode. However, for production purposes, NetApp recommends creating a cluster of BIG-IP instances to avoid a single point of failure.



An F5 BIG-IP system can be deployed on dedicated hardware, in the cloud, or as a virtual appliance on-premises with versions greater than 12.x for it to be integrated with F5 CIS. For the purpose of this document, the F5 BIG-IP system was validated as a virtual appliance, for example using the BIG-IP VE edition.

Validated releases

This solution makes use of the virtual appliance deployed in VMware vSphere. Networking for the F5 Big-IP virtual appliance can be configured in a two-armed or three-armed configuration based on your network environment. The deployment in this document is based on the two-armed configuration. Additional details on configuring the virtual appliance for use with Anthos can be found [here](#).

The Solutions Engineering Team at NetApp have validated the releases in the following table in our lab to work with deployments of Anthos On-Prem:

Make	Type	Version
F5	BIG-IP VE	15.0.1-0.0.11
F5	BIG-IP VE	16.1.0-0.0.19

Installation

To install F5 BIG-IP, complete the following steps:

1. Download the virtual application Open Virtual Appliance (OVA) file from F5 [here](#).



To download the appliance, a user must register with F5. They provide a 30-day demo license for the Big-IP Virtual Edition Load Balancer. NetApp recommends a permanent 10Gbps license for the production deployment of an appliance.

2. Right-click the Infrastructure Resource Pool and select Deploy OVF Template. A wizard launches that allows you to select the OVA file that you just downloaded in Step 1. Click Next.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

URL

Local file

BIGIP-15.0.1-0.....ALL-vmware.ova

CANCEL

BACK

NEXT

3. Click Next to continue through each step and accept the default values for each screen presented until you reach the storage selection screen. Select the VM_Datastore that you would like to deploy the virtual machine to, and then click Next.
4. The next screen presented by the wizard allows you to customize the virtual networks for use in the environment. Select VM_Network for the External field and select Management_Network for the Management field. Internal and HA are used for advanced configurations for the F5 Big-IP appliance and are not configured. These parameters can be left alone, or they can be configured to connect to non-infrastructure, distributed port groups. Click Next.

5. Review the summary screen for the appliance, and, if all the information is correct, click Finish to start the deployment.
6. After the virtual appliance is deployed, right-click it and power it up. It should receive a DHCP address on the management network. The appliance is Linux-based, and it has VMware Tools deployed, so you can view the DHCP address it receives in the vSphere client.
7. Open a web browser and connect to the appliance at the IP address from the previous step. The default login is admin/admin, and, after the first login, the appliance immediately prompts you to change the admin password. It then returns you to a screen where you must log in with the new credentials.



8. The first screen prompts the user to complete the Setup Utility. Begin the utility by clicking Next.
9. The next screen prompts for activation of the license for the appliance. Click Activate to begin. When prompted on the next page, paste either the 30-day evaluation license key you received when you registered for the download or the permanent license you acquired when you purchased the appliance. Click Next.



For the device to perform activation, the network defined on the management interface must be able to reach the internet.

10. On the next screen, the End User License Agreement (EULA) is presented. If the terms in the license are acceptable, click Accept.
11. The next screen counts the elapsed time as it verifies the configuration changes that have been made so far. Click Continue to resume with the initial configuration.
12. The Configuration Change window closes, and the Setup Utility displays the Resource Provisioning menu. This window lists the features that are currently licensed and the current resource allocations for the virtual appliance and each running service.
13. Clicking the Platform menu option on the left enables additional modification of the platform. Modifications include setting the management IP address configured with DHCP, setting the host name and the time zone the appliance is installed in, and securing the appliance from SSH accessibility.
14. Next click the Network menu, which enables you to configure standard networking features. Click Next to begin the Standard Network Configuration wizard.
15. The first page of the wizard configures redundancy; leave the defaults and click Next. The next page enables you to configure an internal interface on the load balancer. Interface 1.1 maps to the VMNIC labeled Internal in the OVF deployment wizard.



The spaces in this page for Self IP Address, Netmask, and Floating IP address can be filled with a non-routable IP for use as a placeholder. They can also be filled with an internal network that has been configured as a distributed port group for virtual guests if you are deploying the three-armed configuration. They must be completed to continue with the wizard.

16. The next page enables you to configure an external network that is used to map services to the pods deployed in Kubernetes. Select a static IP from the VM_Network range, the appropriate subnet mask, and a floating IP from that same range. Interface 1.2 maps to the VMNIC labeled External in the OVF deployment wizard.
17. On the next page, you can configure an internal-HA network if you are deploying multiple virtual appliances in the environment. To proceed, you must fill the Self-IP Address and the Netmask fields, and you must select interface 1.3 as the VLAN Interface, which maps to the HA network defined by the OVF template wizard.
18. The next page enables you to configure the NTP servers. Then click Next to continue to the DNS setup. The DNS servers and domain search list should already be populated by the DHCP server. Click Next to

accept the defaults and continue.

19. For the remainder of the wizard, click Next to continue through the advanced peering setup, the configuration of which is beyond the scope of this document. Then click Finish to exit the wizard.
20. Create individual partitions for the Anthos admin cluster and each user cluster deployed in the environment. Click System in the menu on the left, navigate to Users, and click Partition List.
21. The displayed screen only shows the current common partition. Click Create on the right to create the first additional partition, and name it `GKE-Admin`. Then click Repeat, and name the partition `User-Cluster-1`. Click the Repeat button again to name the next partition `User-Cluster-2`. Finally click Finished to complete the wizard. The Partition list screen returns with all the partitions now listed.

Integration with Anthos

There is a section in each configuration file, respectively for the admin cluster, and each user cluster that you choose to deploy to configure the load balancer so that it is managed by Anthos On Prem.

The following script is a sample from the configuration of the partition for the GKE-Admin cluster. The values that need to be uncommented and modified are placed in bold text below:

```
# (Required) Load balancer configuration
loadBalancer:
# (Required) The VIPs to use for load balancing
vips:
# Used to connect to the Kubernetes API
controlPlaneVIP: "10.61.181.230"
# # (Optional) Used for admin cluster addons (needed for multi cluster
features). Must
# # be the same across clusters
# # addonsVIP: ""
# (Required) Which load balancer to use "F5BigIP" "Seesaw" or
"ManualLB". Uncomment
# the corresponding field below to provide the detailed spec
kind: F5BigIP
# # (Required when using "ManualLB" kind) Specify pre-defined nodeports
# manualLB:
# # NodePort for ingress service's http (only needed for user cluster)
# ingressHTTPONodePort: 0
# # NodePort for ingress service's https (only needed for user
cluster)
# ingressHTTPSNodePort: 0
# # NodePort for control plane service
# controlPlaneNodePort: 30968
# # NodePort for addon service (only needed for admin cluster)
# addonsNodePort: 31405
```

```

# # (Required when using "F5BigIP" kind) Specify the already-existing
partition and
# # credentials
f5BigIP:
  address: "172.21.224.21"
  credentials:
    username: "admin"
    password: "admin-password"
  partition: "GKE-Admin"
# # # (Optional) Specify a pool name if using SNAT
# # snatPoolName: ""
# (Required when using "Seesaw" kind) Specify the Seesaw configs
# seesaw:
# (Required) The absolute or relative path to the yaml file to use for
IP allocation
# for LB VMs. Must contain one or two IPs.
# ipBlockFilePath: ""
# (Required) The Virtual Router Identifier of VRRP for the Seesaw
group. Must
# be between 1-255 and unique in a VLAN.
# vrid: 0
# (Required) The IP announced by the master of Seesaw group
# masterIP: ""
# (Required) The number CPUs per machine
# cpus: 4
# (Required) Memory size in MB per machine
# memoryMB: 8192
# (Optional) Network that the LB interface of Seesaw runs in (default:
cluster
# network)
# vCenter:
# vSphere network name
# networkName: VM_Network
# (Optional) Run two LB VMs to achieve high availability (default:
false)
# enableHA: false

```

Installing MetalLB load balancers

This page lists the installation and configuration instructions for the MetalLB managed load balancer.

Installing The MetalLB Load Balancer

The MetalLB load balancer is fully integrated with Anthos Clusters on VMware and has automated deployment performed as part of the Admin and User cluster setups starting with the 1.11 release. There are blocks of text

in the respective `cluster.yaml` configuration files that you must modify to provide load balancer info. It is self-hosted on your Anthos cluster instead of requiring the deployment of external resources like the other supported load balancer solutions. It also allows you to create an ip-pool that automatically assigns addresses with the creation of Kubernetes services of type load balancer in clusters that do not run on a cloud provider.

Integration with Anthos

When enabling the MetalLB load balancer for Anthos admin, you must modify a few lines in the `loadBalancer:` section that exists in the `admin-cluster.yaml` file. The only values that you must modify are to set the `controlPlaneVIP:` address and then set the `kind:` as MetalLB. See the following code snippet for an example:

```
# (Required) Load balancer configuration
loadBalancer:
  # (Required) The VIPs to use for load balancing
  vips:
    # Used to connect to the Kubernetes API
    controlPlaneVIP: "10.61.181.230"
    # # (Optional) Used for admin cluster addons (needed for multi cluster
features). Must
    # # be the same across clusters
    # addonsVIP: ""
  # (Required) Which load balancer to use "F5BigIP" "Seesaw" "ManualLB" or
"MetalLB".
  # Uncomment the corresponding field below to provide the detailed spec
  kind: MetalLB
```

When enabling the MetalLB load balancer for Anthos user clusters, there are two areas in each `user-cluster.yaml` file that you must update. First, in a manner similar to the `admin-cluster.yaml` file, you must modify the `controlPlaneVIP:`, `ingressVIP:`, and `kind:` values in the `loadBalancer:` section. See the following code snippet for an example:

```
loadBalancer:
  # (Required) The VIPs to use for load balancing
  vips:
    # Used to connect to the Kubernetes API
    controlPlaneVIP: "10.61.181.240"
    # Shared by all services for ingress traffic
    ingressVIP: "10.61.181.244"
  # (Required) Which load balancer to use "F5BigIP" "Seesaw" "ManualLB" or
"MetalLB".
  # Uncomment the corresponding field below to provide the detailed spec
  kind: MetalLB
```



The `ingressVIP` IP address must exist within the pool of IP addresses assigned to the MetalLB load balancer later in the configuration.

You then need to navigate to the `metalLB:` subsection and modify the `addressPools:` section by naming the pool in the `- name: variable`. You must also create a pool of ip-addresses that MetalLB can assign to services of type `LoadBalancer` by providing a range to the `addresses: variable`.

```
# # (Required when using "MetalLB" kind in user clusters) Specify the
MetalLB config
  metalLB:
    # # (Required) A list of non-overlapping IP pools used by load balancer
typed services.
    # # Must include ingressVIP of the cluster.
    addressPools:
    # # (Required) Name of the address pool
    - name: "default"
    # # (Required) The addresses that are part of this pool. Each address
must be either
    # # in the CIDR form (1.2.3.0/24) or range form (1.2.3.1-1.2.3.5).
    addresses:
    - "10.61.181.244-10.61.181.249"
```



The address pool can be provided as a range like in the example, limiting it to a number of addresses in a particular subnet, or it can be provided as a CIDR notation if the entire subnet is made available.

1. When Kubernetes services of type `LoadBalancer` are created, MetalLB automatically assigns an externalIP to the services and advertises the IP address by responding to ARP requests.

Installing SeeSaw load balancers

This page lists the installation and configuration instructions for the SeeSaw managed load balancer.

Seesaw is the default managed network load balancer installed in an Anthos Clusters on VMware environment from versions 1.6 to 1.10.

Installing The SeeSaw load balancer

The SeeSaw load balancer is fully integrated with Anthos Clusters on VMware and has automated deployment performed as part of the Admin and User cluster setups. There are blocks of text in the `cluster.yaml` configuration files that must be modified to provide load balancer info, and then there is an additional step prior to cluster deployment to deploy the load balancer using the built in `gkectl` tool.



SeeSaw load balancers can be deployed in HA or non-HA mode. For the purpose of this validation, the SeeSaw load balancer was deployed in non-HA mode, which is the default setting. For production purposes, NetApp recommends deploying SeeSaw in an HA configuration for fault tolerance and reliability.

Integration with Anthos

There is a section in each configuration file, respectively for the admin cluster, and in each user cluster that you choose to deploy to configure the load balancer so that it is managed by Anthos On-Prem.

The following text is a sample from the configuration of the partition for the GKE-Admin cluster. The values that need to be uncommented and modified are placed in bold text below:

```
loadBalancer:
# (Required) The VIPs to use for load balancing
vips:
# Used to connect to the Kubernetes API
controlPlaneVIP: "10.61.181.230"
# # (Optional) Used for admin cluster addons (needed for multi cluster
features). Must
# # be the same across clusters
# # addonsVIP: ""
# (Required) Which load balancer to use "F5BigIP" "Seesaw" or
"ManualLB". Uncomment
# the corresponding field below to provide the detailed spec
kind: Seesaw
# # (Required when using "ManualLB" kind) Specify pre-defined nodeports
# manualLB:
# # NodePort for ingress service's http (only needed for user cluster)
# ingressHTTPTNodePort: 0
# # NodePort for ingress service's https (only needed for user
cluster)
# ingressHTTPSNodePort: 0
# # NodePort for control plane service
# controlPlaneNodePort: 30968
# # NodePort for addon service (only needed for admin cluster)
# addonsNodePort: 31405
# # (Required when using "F5BigIP" kind) Specify the already-existing
partition and
# # credentials
# f5BigIP:
# address:
# credentials:
# username:
# password:
# partition:
# # # (Optional) Specify a pool name if using SNAT
# # snatPoolName: ""
# (Required when using "Seesaw" kind) Specify the Seesaw configs
seesaw:
# (Required) The absolute or relative path to the yaml file to use for
IP allocation
```

```

# for LB VMs. Must contain one or two IPs.
ipBlockFilePath: "admin-seesaw-block.yaml"
# (Required) The Virtual Router Identifier of VRRP for the Seesaw
group. Must
# be between 1-255 and unique in a VLAN.
vrid: 100
# (Required) The IP announced by the master of Seesaw group
masterIP: "10.61.181.236"
# (Required) The number CPUs per machine
cpus: 1
# (Required) Memory size in MB per machine
memoryMB: 2048
# (Optional) Network that the LB interface of Seesaw runs in (default:
cluster
network)
# network)
vCenter:
# vSphere network name
networkName: VM_Network
# (Optional) Run two LB VMs to achieve high availability (default:
false)
enableHA: false

```

The SeeSaw load balancer also has a separate static `seesaw-block.yaml` file that you must provide for each cluster deployment. This file must be located in the same directory relative to the `cluster.yaml` deployment file, or the full path must be specified in the section above.

A sample of the `admin-seesaw-block.yaml` file looks like the following script:

```

blocks:
- netmask: "255.255.255.0"
  gateway: "10.63.172.1"
  ips:
- ip: "10.63.172.152"
  hostname: "admin-seesaw-vm"

```



This file provides the gateway and netmask for the network that the load balancer provides to the underlying cluster, as well as the management IP and hostname for the virtual machine that is deployed to run the load balancer.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.