By Gregory Coward, Solution Architect, F5 Networks (https://www.f5.com/)

This tutorial shows how to set up the F5 BIG-IP (https://www.f5.com/products/big-ip-services) Application Delivery Controller (ADC) before you integrate with GKE On-Prem (/gke-on-prem/docs/how-to/installation/install) using the manual load-balancing mode on GKE On-Prem (/gke-on-prem/docs/how-to/installation/manual-lb). If you're interested in installing F5 BIG-IP ADC on GKE On-Prem to automatically provision L4 load-balancing services, see Installing F5 BIG-IP ADC for GKE On-Prem (/solutions/partners/installing-f5-big-ip-adc-for-gke-on-prem).

F5 is a leading provider of ADC services. The F5 BIG-IP platform provides various services to help you enhance the security, availability, and performance of your apps. These services include, L7 load balancing, network firewalling, web application firewalling (WAF) (https://wikipedia.org/wiki/Web_application_firewall), DNS services, and more. For GKE On-Prem, BIG-IP provides external access and L3/4 load-balancing services.

When deployed in integrated mode, Anthos (/anthos/) uses a version of F5 container ingress services (https://clouddocs.f5.com/containers/v2/kubernetes/) (CIS) to automatically provision L4 load-balancing services on the BIG-IP platform. CIS continues to monitor and update BIG-IP when the GKE On-Prem cluster is modified. However, CIS comes with limitations.

At the time of publication, you cannot add L7 services such as F5 Advanced WAF (https://www.f5.com/products/security/advanced-waf) or Access Policy Manager (F5 APM (https://www.f5.com/products/security/access-policy-manager)) to the virtual IP address endpoints when the environment is deployed using the integrated mode. This limitation is due to the nature of CIS. Any modifications made to the BIG-IP partitons are overwritten by the CIS controller when it's updated.

By deploying the GKE On-Prem environment using the manual load-balancer mode, on the other hand, you create the required virtual servers and related BIG-IP resources prior to deploying GKE On-Prem. This type of deployment lets you customize and secure the BIG-IP hosted environment endpoints. The trade-off is that as the environment changes, for example when cluster node instances are added or removed, you need to manually update the BIG-IP.

- Learn about the BIG-IP architecture.
- Configure the BIG-IP for GKE On-Prem external endpoints.
- Create virtual servers.

This tutorial uses the following billable components of Google Cloud:

- Anthos (/anthos/pricing)

To generate a cost estimate based on your projected usage, use the pricing calculator (/products/calculator). New Google Cloud users might be eligible for a free trial (/free-trial).

1. Obtain an F5 BIG-IP Application Delivery Controller and license. The F5 BIG-IP ADC is available in various hardware platforms and virtual editions. Regardless of the platform you use, the solution is supported, and the following configuration process is applicable.

   There are three types of licenses for F5 BIG-IP.

License type	Production	Evaluation	Demonstration
Expiration date	No	Yes, 45 days	Yes, 30 days
Throughput limitations	Up to 40 Gbps on GKE On-Prem	25 Mbps - 10 Gbps	1 Mbps
Use case	Production	Proof of concept, demonstration	Proof of concept, demonstration

   - If you have a production license, you can use that license.

- For the purpose of this tutorial, you can request a free trial license key (https://www.f5.com/trials).
- If your throughput requirements are greater than 1 Mbps provided by the trial license, you can request an evaluation license from F5:
  - Request a BIG-IP evaluation license (https://www.google.com/url?sa=D&q=https%3A%2F%2Ff5.com%2Fproducts%2Ftrials%2Fbig-ip-evaluation-request%2Fpopup%2Ftrue%3Fsbtitle%3DBIG-IP%2BEvaluation%2BTrial%2BRequest%26utm_source%3Df5com%26utm_medium%3Dweb%26utm_campaign%3Dhtb-trial-page).
  - Send email to F5 to request a BIG-IP evaluation license (mailto:googleteam@f5.com).
- If the BIG-IP system contains an evaluation or demonstration license, the BIG-IP system stops processing traffic when the license expires.

2. Activate a license key for BIG-IP (https://www.f5.com/products/get-f5).
3. Make sure your environment meets the following minimum system requirements:
   - 8 vCPUs that aren't shared between other hosts on that system
   - 16 GB memory that isn't shared between other hosts on that system

There are two common scenarios to deploy BIG-IP ADC to GKE On-Prem clusters. Because the BIG-IP acts as a proxy for external access to the clusters, it's common to deploy a BIG-IP with three or more interfaces, as illustrated in the following diagram.

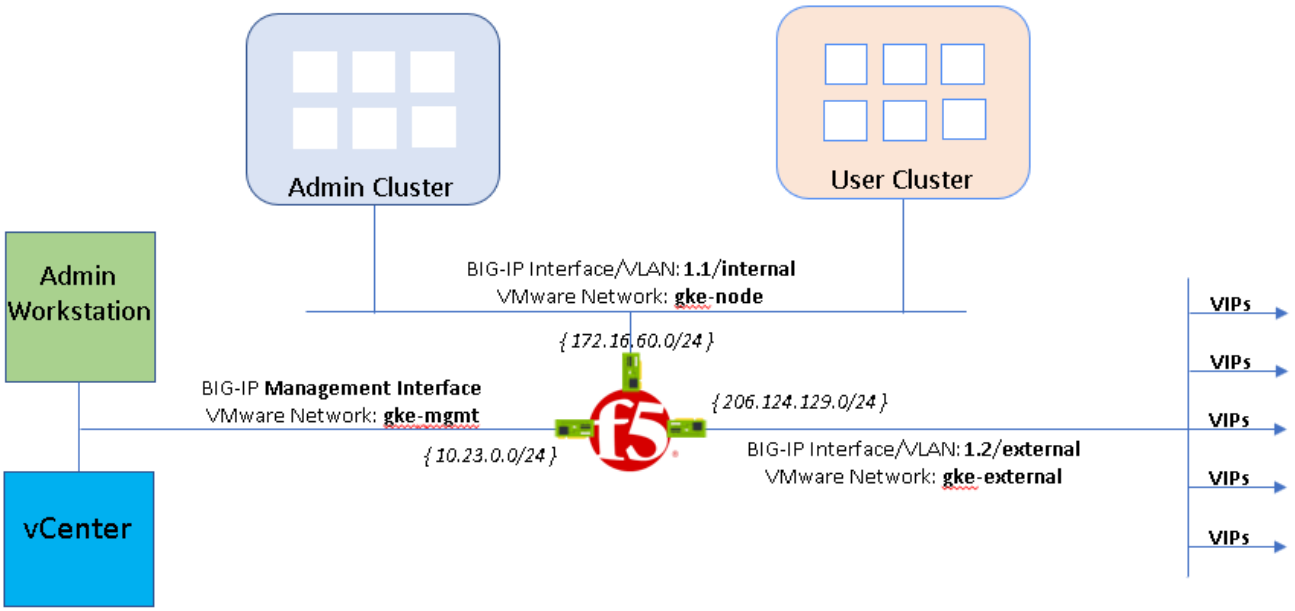

In the preceding diagram, separate interfaces serve internal *private*- and external *public*-facing traffic independently. This architecture provides better visibility for monitoring and troubleshooting, and increased throughput.

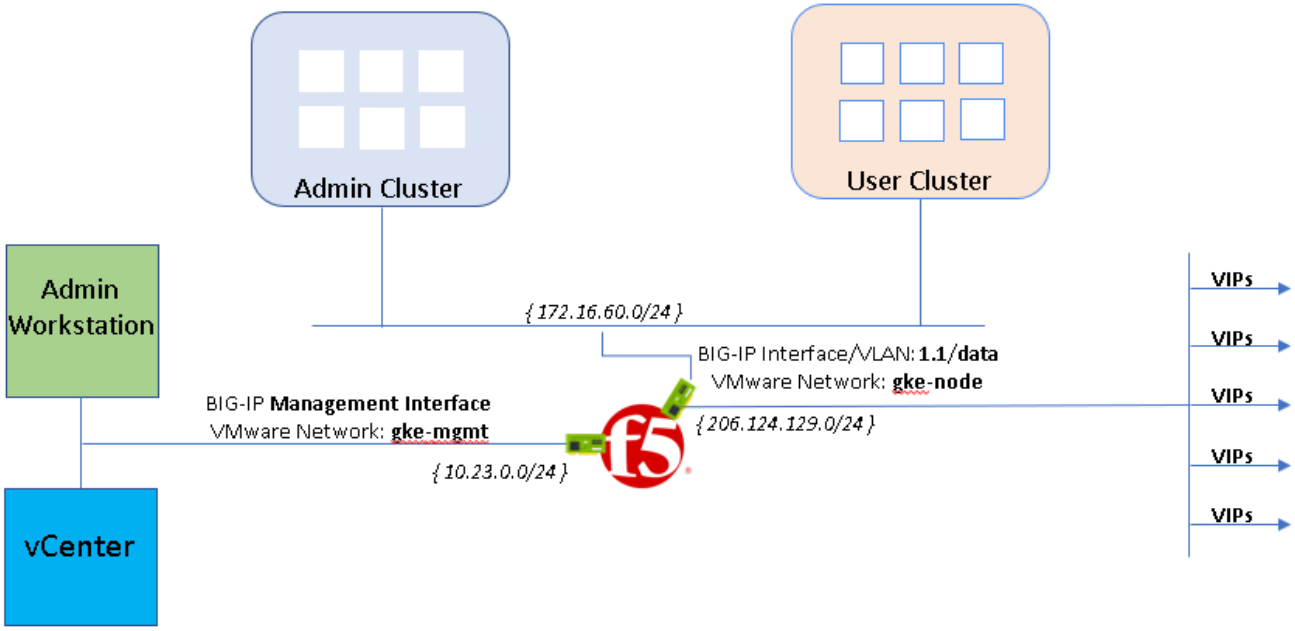

While this configuration isn't considered a best practice, if you're integrating into an existing environment with a pre-defined network architecture, you might need this type of configuration.

1. Follow the instructions to set up BIG-IP virtual edition deployed on VMware ESXi 6.5 (https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-ve-setup-vmware-esxi-13-1-0/3.html)

   .

   The OVF template requires configuring four interfaces. The fourth interface is designated for HA heartbeat traffic between BIG-IP pairs. For this three-arm deployment, assign the internal network **gke-node**.

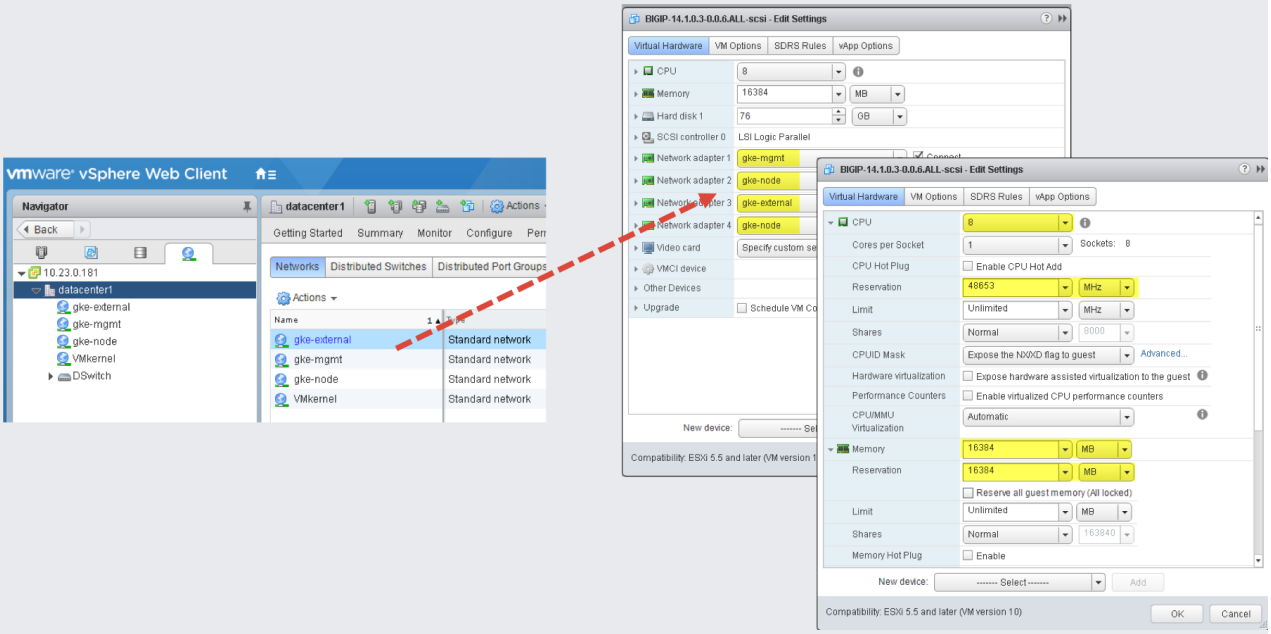

2. After the VM boots, use the F5 BIG-IP's Setup utility (https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/big-ip-system-initial-configuration-14-0-0/01.html#guid-5cfb2400-eaa3-4ba6-858d-8dd8db68b2cb) for initial configuration. The setup utility walks you through the following configuration tasks:

   a. From a network accessible workstation on which you configured the **gke-mgmt** interface, go to the following URL **https://*management_IP_address***, where ***management_IP_address*** is the address you configured for your device.

   b. When prompted, enter the default username as **admin** and the password as **admin**.

   c. Click **Log in**.

3. To install a license, in the **Base Registration Key** field, enter your key. The type of license dictates the BIG-IP's services and bandwidth limits.

4. To enhance performance when working with GKE clusters, set the **Management (MGMT)** plane provisioning to **Large**.

5. To provide L3/4 load balancing to the GKE On-Prem environment, set the **Local Traffic (LTM)** module is set to **Nominal**

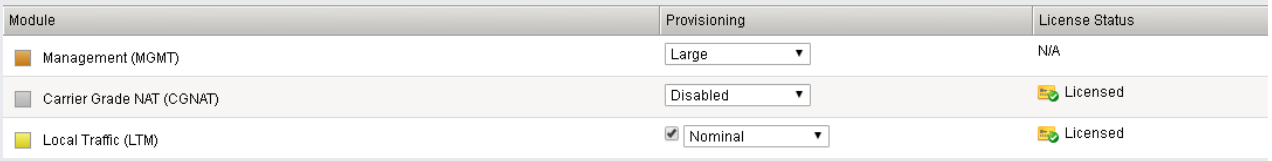

6. On the **Host and User Information** page, you provide the hostname, FQDN of BIG-IP system, and update the *admin* and *root* account passwords.

7. On the **Networking** page, you walk through configuring the BIG-IP's basic networking. The utility creates the internal, `gke-node` and external, `gke-external` interfaces, VLANs, and self-IP addresses.

The 4th interface deployed by VMware is left unconfigured.

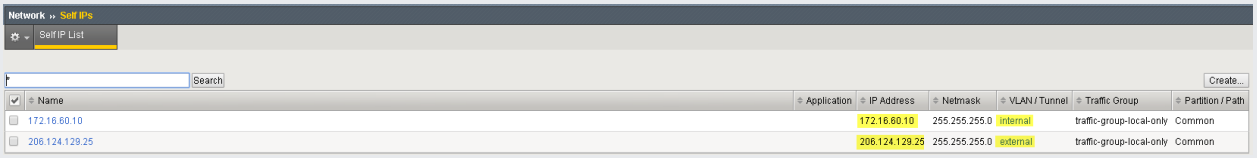

After the Setup utility completes, you have a functioning BIG-IP with a management plane interface attached to the `gke-mgmt` VMware network and two data plane interfaces attached to VMware networks,`gke-node` and `gke-external`.

Before you deploy GKE On-Prem (/gke-on-prem/docs/how-to/installation/install), more configuration of the BIG-IP is required.

- Create an administrative partition
  (https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-user-account-administration-12-0-0/3.html)
  for each admin and user cluster you intend to expose and access.

  Initially, you define two partitions: one for the admin cluster, and one for the first user cluster. Don't use cluster partitions for anything else. Each of the clusters must have a

partition that is for the sole use of that cluster.

The existing Administrator role provides enough permissions for use with GKE On-Prem. For more information, see User roles (https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-user-account-administration-11-6-0/3.html#taskid). You can also learn how to create additional users (https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-user-account-administration-13-1-0/4.html).

Before deploying GKE On-Prem, you must configure the BIG-IP with six virtual servers (/gke-on-prem/docs/how-to/installation/manual-lb), (VIPs), corresponding to the following GKE On-Prem endpoints:

- **Admin partition**
  - VIP for admin cluster control plane (port exposed: `443`)
  - VIP for admin cluster ingress controller (port exposed: `443`)
  - VIP for admin cluster ingress controller (port exposed: `80`)
  - VIP for user control plane (port exposed: `443`)
- **User partition**
  - VIP for user cluster ingress controller (port exposed: `443`)
  - VIP for user cluster ingress controller (port exposed: `80`)

Perform the following steps from both the admin and user partitions to create node objects on the BIG-IP for each host specified in the corresponding host configuration files.

GKE On-Prem clusters can run with one of two load-balancing modes: *integrated* or *manual*. For manual mode, cluster nodes (both admin and user clusters) must be assigned static IP

addresses (/gke-on-prem/docs/how-to/installation/static-ips). These addresses are in turn used to configure node objects on the BIG-IP system. You will create a node object for each GKE On-Prem cluster node. The nodes are added to backend pools that are then associated with virtual servers.

1. To log in to the BIG-IP management console, go to the IP address. The address is provided during the installation.
2. Click the **Administrative partition** that you previously created.
3. Go to **Local Traffic** > **Nodes** > **Node List**.
4. Click **Create**.
5. Enter a name and IP address for each cluster host and click **Finished**.
6. Repeat these steps each admin cluster member.
7. Repeat these steps for each user cluster member, but click the **User partition** instead of the Administrative partition.

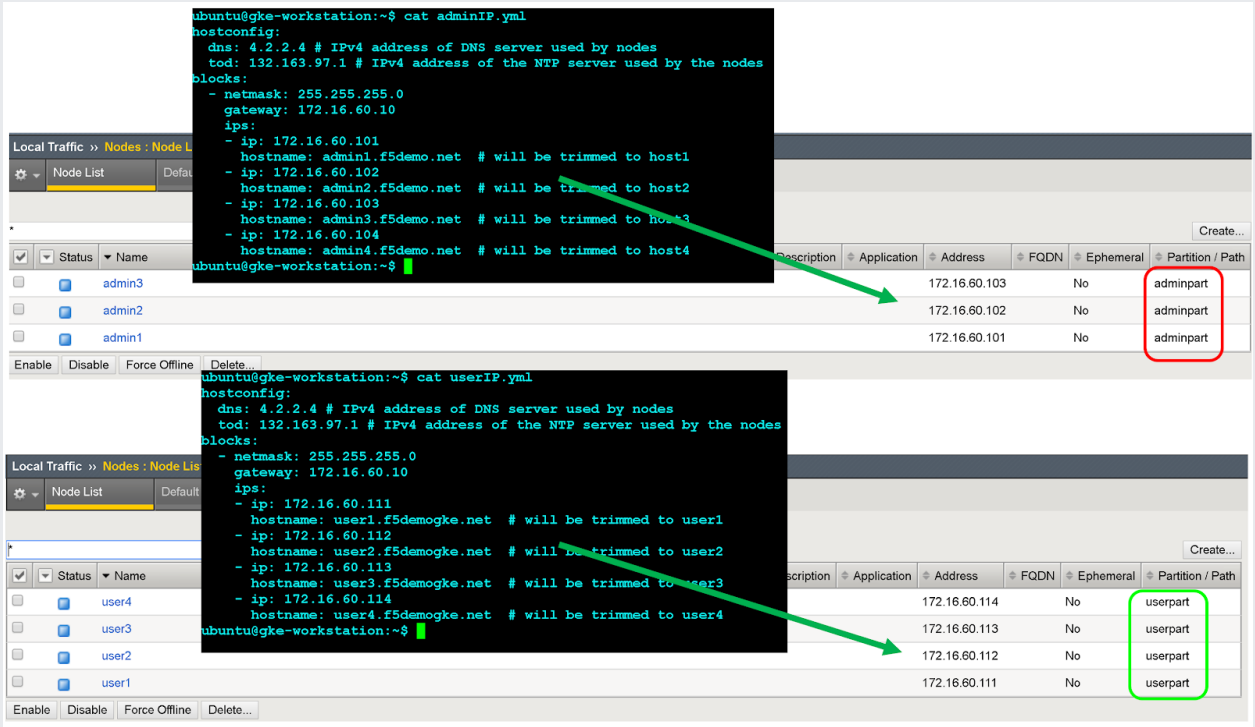

You create a backend pool for each required VIP, seven in total.

1. In the BIG-IP management console, click **adminpart** for the admin partiton that you previously created.
2. Go to **Local Traffic** > **Pools** > **Pool List**.
3. Click **Create**.
4. In the **Configuration** drop-down list, click **Advanced**.
5. In the **Name** field, enter `Istio-80-pool`.
6. To verify the pool member accessibility, under **Health Monitor**, click **tcp**. Optional: Because this is a manual configuration, you can also take advantage of more advanced monitors (https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-local-traffic-manager-monitors-reference-13-1-0.html) as appropriate for your deployment.
7. For **Action on Service Down,** click **Reject**.

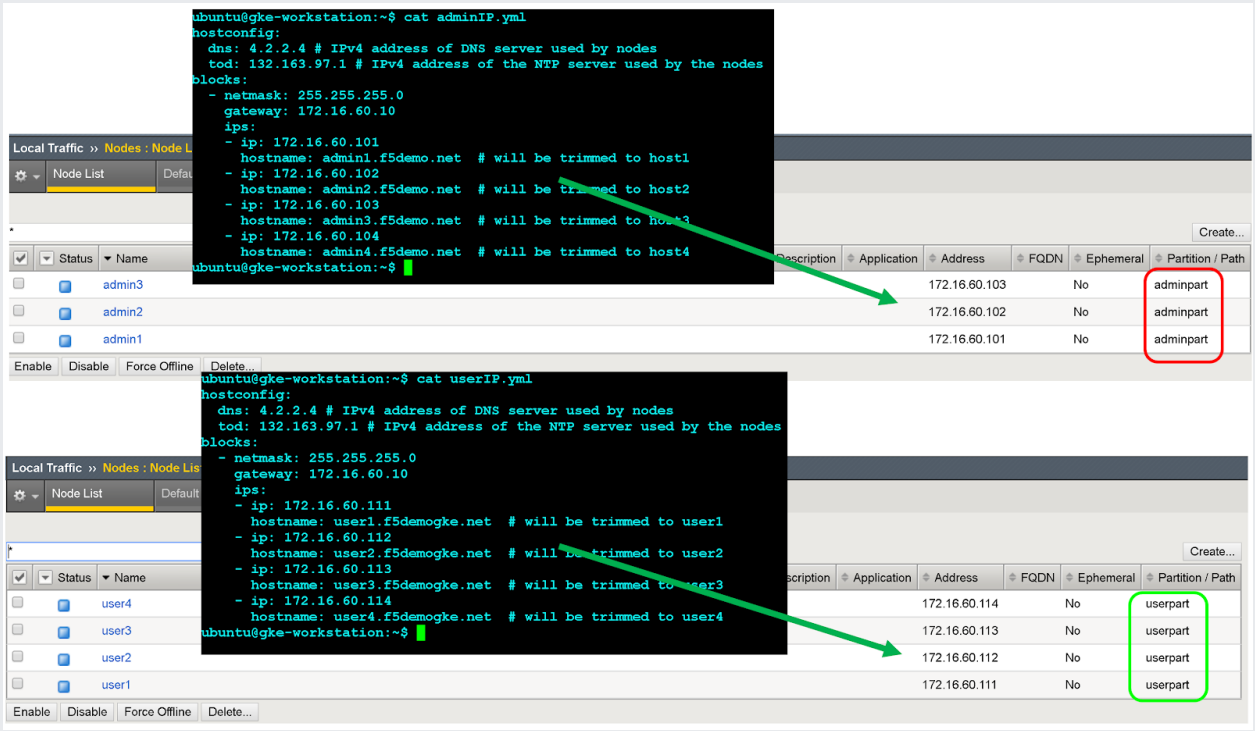

8. For this tutorial, in the **Load Balancing Method** drop-down list, click **Round Robin**.
9. In the **New Members** section, click **Node List** and then select the previously created node.
10. In the **Service Port** field, enter the appropriate `nodePort` from the gkectl configuration file (/gke-on-prem/docs/how-to/installation/manual-lb#reserve_nodeports).
11. Click **Add**.

12. Repeat steps 8-9 and add each cluster node instance.

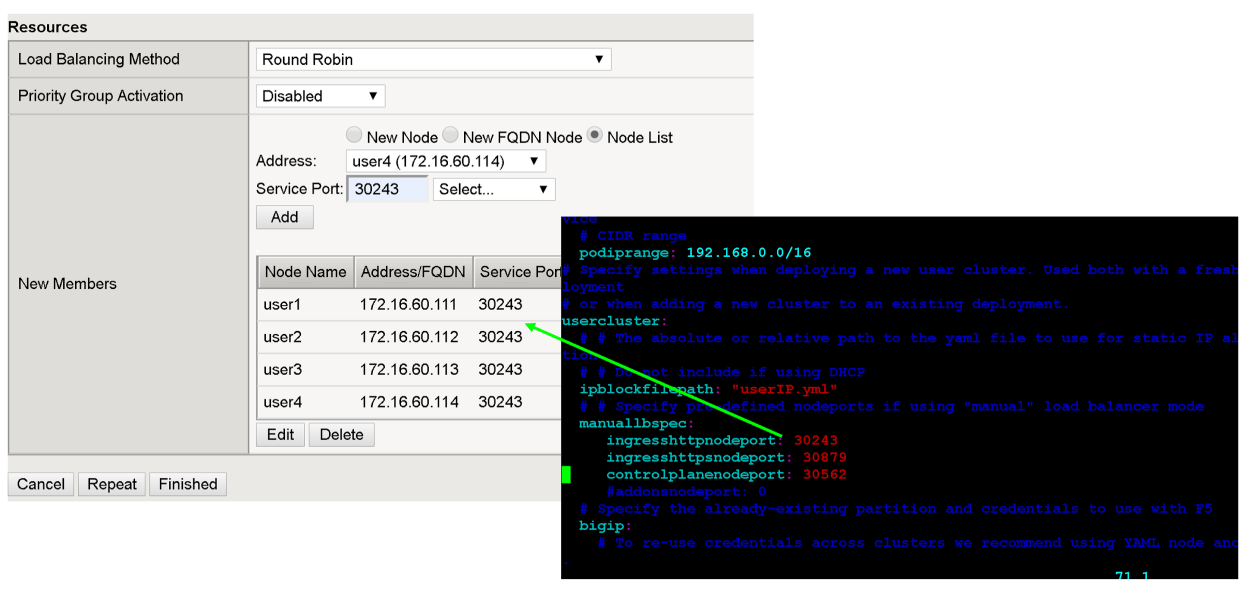

13. Click **Finished**.
14. Repeat all of these steps in this section for the remaining required admin cluster VIPs (/gke-on-prem/docs/how-to/installation/manual-lb).
15. Repeat all of these steps in this section for each user cluster pool, except in step 1, click **userpart** instead of **adminpart**.

You create a total of seven virtual servers on the BIG-IP with five for the admin clusters and two for the user clusters. The virtual servers correspond to the VIPs required to deploy GKE On-Prem.

1. In the BIG-IP management console, click the **Admin partition** that you previously created.
2. Go to **Local Traffic** > **Virtual Servers** > **Virtual Server List**.
3. Click **Create**.
4. In the **Name** field, enter `istio-ingress-80`.
5. In the **Destination Address/Mask** field, enter the IP address for the VIP. For this tutorial, use the HTTP ingress VIP in the `gkectl configuration file` (/gke-on-prem/docs/how-to/installation/manual-lb).

6. In the **Service Port** field, enter the appropriate listener port for the VIP. For this tutorial, use port `80`.

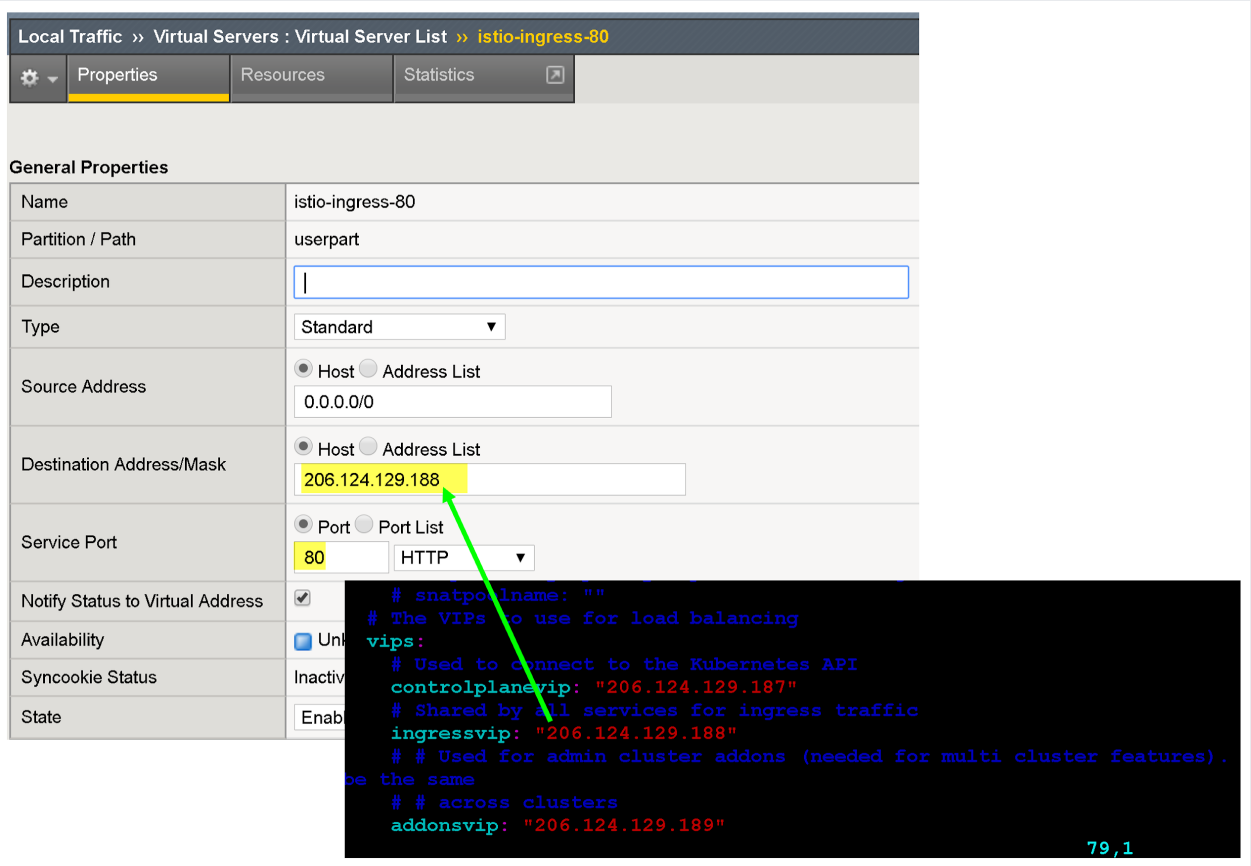

   There are several configuration options for enhancing your app's endpoint, such as associating protocol-specific profiles, certificate profiles (https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/ltm-profiles-reference-13-1-0/6.html#guid-cc146765-9237-474b-9faf-0e18a208cb23), and WAF policies (https://support.f5.com/csp/article/K85426947).

7. For **Source Address Translation** click **Auto Map**.
8. For **Default Pool** select the appropriate pool that you previously created.
9. Click **Finished**.
10. Repeat these to create the remaining required admin cluster VIPs (/gke-on-prem/docs/how-to/installation/manual-lb).
11. Repeat these steps to create the user cluster virtual servers, but select the **User partition**.
12. Create and download an archive of the current configuration (https://support.f5.com/csp/article/K4423).

- To further enhance the security and performance of the external-facing VIPs, consider the following:
    - F5 Advanced WAF (https://www.f5.com/products/security/advanced-waf)
    - F5 Access Policy Manager (APM) (https://www.f5.com/products/security/access-policy-manager)
    - Caching & Compression (https://www.f5.com/products/big-ip-services/local-traffic-manager)
    - Advanced health monitoring (https://techdocs.f5.com/kb/en-us/products/big-iq-centralized-mgmt/manuals/product/big-iq-centralized-management-monitoring-and-reports-6-1-0.html)
    - Advanced Load-Balancing Methods (https://www.f5.com/products/big-ip-services/local-traffic-manager)
- Learn more about F5 BIG-IP Application Services (https://www.f5.com/products/big-ip-services).
- Learn more about BIG-IP configurations and capabilities:
    - Certificate profiles (https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/ltm-profiles-reference-13-1-0/6.html#guid-cc146765-9237-474b-9faf-0e18a208cb23)
    - WAF policies (https://support.f5.com/csp/article/K85426947)
- Try out other Google Cloud features for yourself. Have a look at our tutorials (/docs/tutorials).