

Arista EOS and Tigera Calico Technical Brief

Deploying Kubernetes with Arista EOS and Project Calico for high performance networking

Inside

Solution Architecture

Arista's Universal Spine architecture is a high performance layer 3 leaf spine network that provides IP connectivity to the Kubernetes nodes across the data center. Each Kubernetes node runs the Tigera 'calico/node' agent which works with the Linux network stack to establish a BGP session from the node to a leaf switch. The switch will act as a BGP Route Reflector for the BGP peers, and a single AS number is used for the entire rack.

Introduction

As customers deploy their applications onto Kubernetes clusters, many are looking to optimize the performance of network traffic while maintaining operational efficiency. By using Tigera's Calico network plugin for Kubernetes and Arista's Universal Cloud Network, customers are able to directly connect their container workloads to Arista network switches, removing the overhead of encapsulation protocols. In addition, Calico's support for Kubernetes network policy allows administrators to define which pods are allowed to communicate with one another.

This paper describes how to configure Kubernetes networking with Arista EOS and Tigera Calico. In particular, we will cover using Calico to connect either a single-homed or dual-homed server directly to an Arista top of rack (ToR) switch using border gateway protocol (BGP) peering.



Single homed server

In the first scenario a Kubernetes node running Calico is connected to an Arista ToR switch with a single network connection. BGP is established between the node and the ToR without an IP-in-IP overlay. Once the BGP peering is established, the ToR will see routes from the Kubernetes node.




Figure 1: Single homed Kubernetes node

To implement this configuration, the Arista switch must be configured to peer with the Calico node, and act as a iBGP route reflector.

Arista Switch Config - cs-lf13

```
router bgp 65130
  maximum-paths 32
  neighbor 172.20.13.153 remote-as 65130
  neighbor 172.20.13.153 maximum-routes 12000
  neighbor 172.20.13.153 route-reflector-client
```

The Kubernetes administrator will also need to apply two YAML files to configure Calico using the `calicoctl` command. These YAML files are specific to the node:

Calico Config - kube153

bgp-node-kube153-lf13.yaml

```
kind: Node
metadata:
  name: kube-153
spec:
  bgp:
    asNumber: 65130
    ipv4Address: 172.20.13.153/24
```

bgp-peer-node-kube153.yaml

```
apiVersion: projectcalico.org/v3
kind: BGPPeer
metadata:
  name: bgppeer-node-kube153
spec:
  peerIP: 172.20.13.1
  node: kube-153
  asNumber: 65130
```

After configuration the switch will have a route for the container subnet from the attached node:

```
cs-lf13#sh ip route 192.168.58.0
VRF name: default
Codes: C - connected, S - static, K - kernel,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type2, B I - iBGP, B E - eBGP,
       R - RIP, I L1 - ISIS level 1, I L2 - ISIS level 2,
       O3 - OSPFv3, A B - BGP Aggregate, A O - OSPF Summary,
       NG - Nexthop Group Static Route, V - VXLAN Control Service
B I    192.168.58.0/26 [200/0] via 172.20.13.153, Vlan13
```

Dual homed server

For customers looking for redundancy, it is possible to attach the server to two ToR switches. Following the above recommendation for an AS per rack, it is recommended to use BGP and ECMP to provide a simple layer 3 redundant topology from the data center network down to the kubernetes node. This requires configuring the same ASN and router-id on both Arista switches, establishing BGP peering from each NIC to its attached switch, and turning on ECMP.

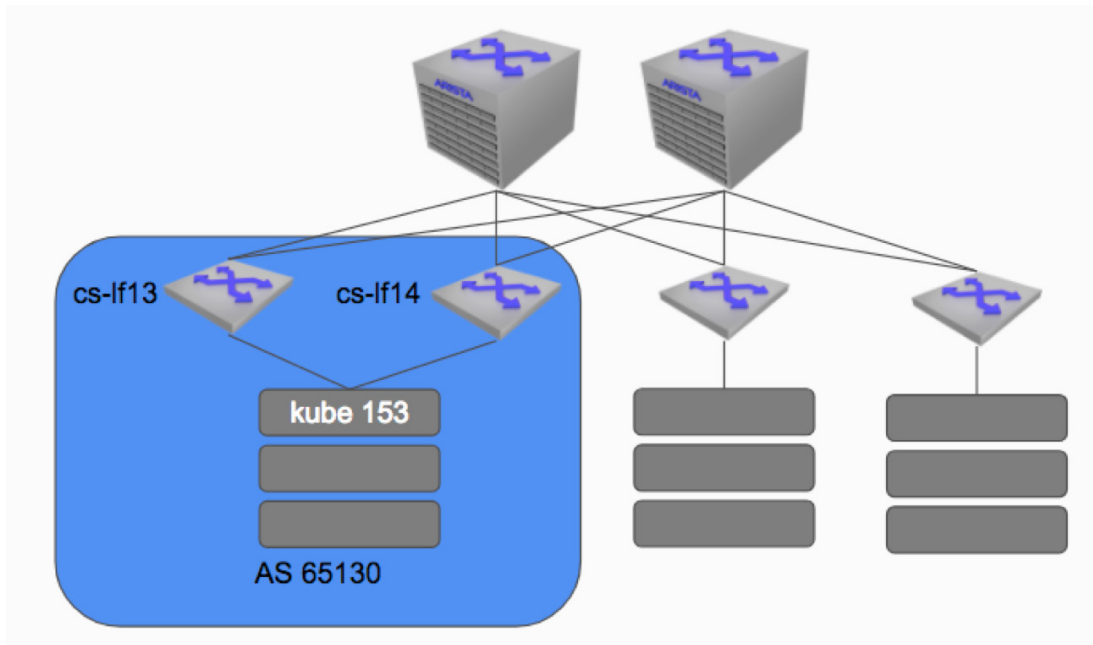


Figure 2: Dual homed Kubernetes server

Arista Config - cs-lf13

```
router bgp 65130
  router-id 192.168.100.13
  maximum-paths 32
  neighbor 172.20.13.153 remote-as 65130
  neighbor 172.20.13.153 next-hop-self
  neighbor 172.20.13.153 route-reflector-client
```

Arista Config - cs-lf14

```
router bgp 65130
  router-id 192.168.100.13
  maximum-paths 32
  neighbor 172.20.14.153 remote-as 65130
  neighbor 172.20.14.153 next-hop-self
  neighbor 172.20.14.153 route-reflector-client
```

Calico Config - kube153

```
bgp-node-kube153.yaml
```

```
kind: Node
```

```
metadata:
```

```
  name: kube-153
```

```
spec:
```

```
  bgp:
```

```
    asNumber: 65130
```

```
bgp-peer-node-kube153-lf13.yaml
```

```
apiVersion: projectcalico.org/v3
```

```
kind: BGPPeer
```

```
metadata:
```

```
  name: bgppeer-node-kube153
```

```
spec:
```

```
  peerIP: 172.20.13.1
```

```
  node: kube-153
```

```
  asNumber: 65130
```

```
bgp-peer-node-kube153-lf14.yaml
```

```
apiVersion: projectcalico.org/v3
```

```
kind: BGPPeer
```

```
metadata:
```

```
  name: bgppeer-node-kube153
```

```
spec:
```

```
  peerIP: 172.20.13.1
```

```
  node: kube-153
```

```
  asNumber: 65130
```

Note: to apply the dual-homed configuration requires some changes to the Calico BGP template which go beyond the scope of this document. Readers are encouraged to contact their Arista or Tigera solutions architect for assistance in this case.

About Arista

Arista Networks was founded to deliver software-driven cloud networking solutions for large data center storage and computing environments. Arista's award-winning platforms, ranging in Ethernet speeds from 10 to 100 gigabits per second, redefine scalability, agility and resilience. Arista has shipped more than 15 million cloud networking ports worldwide with CloudVision and EOS.

About Tigera

Tigera (www.tigera.io) delivers solutions for secure application connectivity for the cloud native world by addressing the connectivity and network security challenges that arise in at-scale production. Tigera's flagship product, CNX, meets enterprise needs for zero trust security, multi-cloud and legacy environment support, organizational control and compliance, and operational simplicity. CNX leverages open source projects including Tigera's Project Calico (www.projectcalico.org), listo, and Kubernetes

Conclusion

By using standard network protocols such as BGP, and best in class data center networking switches, Arista and Tigera provide a high performance networking solution for Kubernetes clusters. This solution provides operational simplicity and visibility to networking teams deploying Kubernetes in the data center, while supporting features such as Kubernetes network policy. With Arista EOS and Tigera Calico, customers can deploy scalable Kubernetes clusters without sacrificing features or performance.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390

Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2016 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. May 11, 2018