

**Open Compute  
Network Operating System  
Version 1.1**

**OcNOS™ Validated Solution Guide**  
**EBGP-based Data Center with OcNOS**

This documentation is subject to change without notice. The software described in this document and this documentation are furnished under a license agreement or nondisclosure agreement. The software and documentation may be used or copied only in accordance with the terms of the applicable agreement. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or any means electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's internal use without the written permission of IP Infusion Inc.

**IP Infusion Inc.**

3965 Freedom Circle, Suite 200, Santa Clara, CA 95054 (408) 400-1900- main (408) 400-1863 - fax  
For support, questions, or comments via E-mail, contact: support@ipinfusion.com

IP Infusion, ZebOS, VirNOS and OcnOS are trademarks or registered trademarks of IP Infusion. All other trademarks, service marks, registered trademarks, or registered service marks mentioned are the property of their respective owners.

## Table of Contents

<b>CHAPTER 1</b> .....	<b>3</b>
<b>Data Center Interconnect Overview</b> .....	<b>3</b>
Large-Scale Data Center Requirements.....	3
Large-Scale Data Center Topologies.....	3
Large-Scale Data Center Routing.....	3
EBGP-Routed Clos Topology-Based Data Center.....	4
EBGP Data Center Design using OcnOS .....	5
<b>CHAPTER 2</b> .....	<b>6</b>
<b>Configuration</b> .....	<b>6</b>
ToR (leaf node) .....	6
Tier-2 (spine node).....	7
Tier-3 (core node) .....	8
Tier-2 (border router).....	9
Tier-3 (WAN router) .....	11
Other Configurations.....	10
Validation .....	11
Conclusion.....	14
References .....	14

# Glossary

BGP – Border Gateway Routing Protocol

EBGP – External BGP

STP – Spanning Tree protocol

TRILL – Transparent interconnection of lots of links

SPB – Shortest path bridging

ECMP – Equal Cost Multipath

# CHAPTER 1

## Data Center Overview

This solution guide describes an approach to build data centers using Layer3 BGP routing protocol. It also summarizes on some design philosophies for data center and why E-BGP is better suited.

- Large-scale data center requirements
- Large-scale data center topologies
- Large-scale data center routing
- EBGp-routed large-scale Clos topology-based data center

### Large-Scale Data Center Requirements

The design of large-scale data centers is driven by operational simplicity and network stability. Operational simplicity and network stability ensures easier manageability and therefore reduced operational expenses. From the network design aspect, the requirements are:

- Ability to accommodate the variable-application bandwidth and strict latency requirements.
- Ability to handle the increased east-west (server-to-server) traffic within the data center due to massive data replication between clusters and virtual machine migrations.
- Traffic-Engineering with application load balancing. The network infrastructure should itself perform controlled per-hop traffic engineering.
- Minimize CAPEX and incorporate vendor diversity by using a simple, interoperable routing protocol with a minimal set of features.
- A design to minimize OPEX by keeping the failure domain at the lowest level in the network hierarchy.

### Large-Scale Data Center Topologies

A traditional tree-based (upside down) topology with a three-layer hierarchy of core, aggregation and access layer can be used in a data center design. This approach is suited if the majority of the traffic is entering and leaving (north-south) the data center. An increase in bandwidth requirements then can be addressed by upgrading the device line cards or port density. However with the current trend of increasing server-to-server (east-west) traffic, scaling these networks horizontally is expensive or impossible at times.

A Clos network (leaf and spine) is a horizontally scalable topology where every leaf node is connected to every other spine. The topology can be extended to different stages for scaling. Clos networks are fully non-blocking and load balancing is inherent in the topology itself as all available paths are ECMP. Clos networks are ideal for the current requirements of a large-scale data center.

### Large-Scale Data Center Routing

Layer 2-only routing was used in a traditional tree-based data center topology. Traditional layer-2 protocols such as STP do not give bi-sectional bandwidth, whereas recent developments such as TRILL, SPB have selected vendor support.

However, a hybrid of layer 2 /layer 3 can be used to limit the size of failure domain and scale up the data center. Layer 3 routing can be used in tier 1 (core) and layer 2 in tier 3 (access). Tier 2 can be based on either layer 2 or layer 3. A hybrid model has the advantage of seamless Virtual Machine mobility and requires less IP subnets for the data center. Although this design can scale-up, it is difficult and complex to manage the different protocols.

A layer 3 only design simplifies the network and improves network stability and scalability, as well as localizing the failure domain (confined to the L2 broadcast domain). Seamless virtual mobility can be achieved in a L3 only based data center by using L2 overlay networks. From experiment and analysis, External BGP (EBGP) is considered ideal compared to IGP's due to the following [See Reference]:

- Less complex protocol, simple state machine
- Information flooding overhead is less, no frequent updates unlike IGP's
- Network failure recovery is very fast. Although BGP convergence is slower than IGP, in a Clos topology with ECMP links, the failure is masked as soon as an alternate path is found.
- Failure domain is minimized in a Clos topology with EBGP. Some of the failures are local/hidden/not propagated if the failed link was not selected/advertised as the best path among the ECMP paths by the BGP speaker. The failures, where all devices have to withdraw some prefixes or update the ECMP groups in the FIB, are very limited and in those failures the failed link/node does not impact the re-convergence process.
- Administrator can define the application traffic path. BGP provides services like prefix distribution, prefix filtering, traffic engineering, traffic tagging, and multi-vendor stability better than other IGP's.
- Easier to troubleshoot.

### **EBGP-Routed Clos Topology-Based Data Center**

EBGP-routed Clos topology is considered the best choice for laying the IP fabric in a data center because of the horizontal scalability feature of Clos topology and the ease of use and services provided by EBGP especially prefix-filtering, prefix distribution, and traffic engineering which are required extensively in a data center.

#### *Configuration Guidelines*

Configuration guidelines for laying IP fabric using EBGP efficiently are as follows:

- Run all EBGP sessions over single-hop point-to-point links.
- Use private Autonomous System Numbers (ASNs) (64512-65534) to avoid ASN conflicts.
- Give all tier 1 (core) devices a single ASN.
- Give all tier 2 devices in the same cluster the same unique ASN. A cluster or pod is a group of tier 2 (spine switches) + tier 3 switches (ToR/leaf) + servers.
- Give every tier 3 (ToR) device in a cluster a unique ASN.
- Reuse tier-3 ASNs across clusters. Configure tier-3 devices with the BGP "allowas-in" feature to allow route learning of prefixes from the same ASNs in other clusters.
- Announce server subnets on tier-3 devices via BGP without using route summarization on tier-2 and tier-1 devices.

- Use edge clusters (pods) for external connectivity. Each edge cluster consists of border routers (tier-2) and WAN routers (tier-3). Give each WAN router a unique public ASN to connect the data center to the external world.
- For border routers, remove private ASNs before sending the information to WAN routers by configuring border routers with the “remove-private-AS” BGP feature.
- To relax the BGP ECMP criteria for AS paths, configure BGP “as-path multipath-relax” on all routers/switches. This way, an equal cost path with a different AS PATH, but the same AS PATH length is also considered an equal cost path (ECMP).
- For faster failure detection, configure the BGP session with BFD.
- To avoid recurring BPG update/selection for a single failure through all peers or BGP update message dispersion on a particular speaker, use BGP update groups. The BGP update group feature processes an update once and sends it to a group of neighbors that share a common outbound policy. The BGP RIB is scanned every time for each peer to apply the outbound filter.
- To avoid micro routing loops, configure tier-2 and tier-1 with static discard or null routes rather than a default route. Routing loops can happen when a tier-2 device has lost all its learned prefixes, but has a default route to a tier-1 device and that tier-1 device still has a route back to the tier-2 device.

## EBGP Data Center Design using OcnOS

Figure-1 shows a minimal representation that encompasses all the elements in a layer 3 data center.

The number of ECMPs in the data center is equal to the number of cores (tier-1 switches).

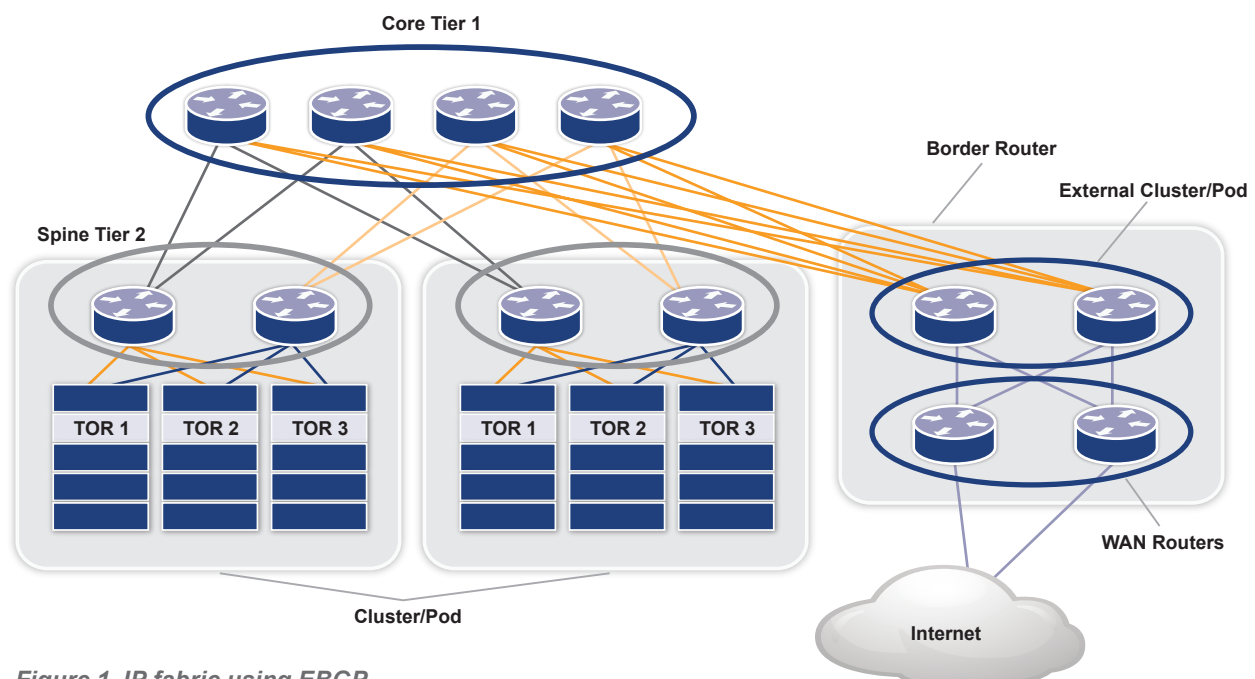


Figure 1. IP fabric using EBGP

Figure 2 shows the Autonomous System Number (ASN) allocation scheme used in the data center. The WAN routers are assigned a public ASN, which connects the data center to external world. The tier-3 ASNs per ToR are reused across the clusters.

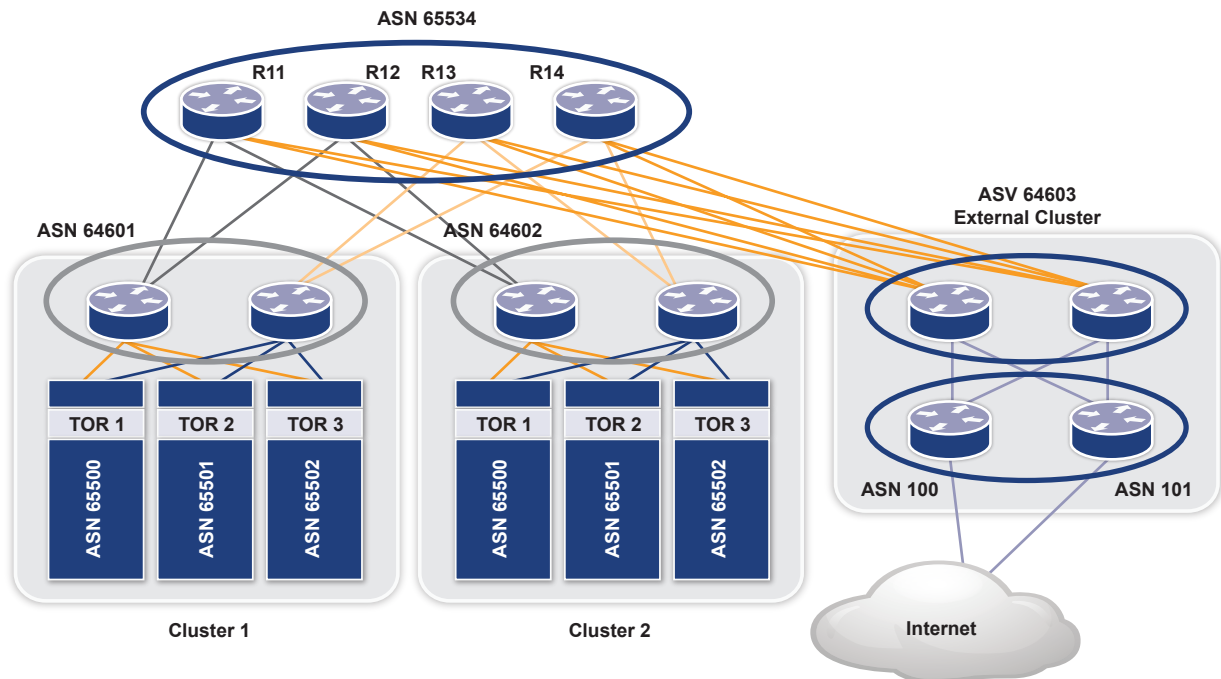


Figure 2: ASN allocation in an EBGP-based data center

## CHAPTER 2

### Configuration

#### ToR (Leaf node)

	Command	Purpose
Step 1	(config)#interface xe1	Enter interface mode.
Step 2	(config-if)#ip address 32.1.0.3/24	Configure ip address on the Interface
Step 3	(config-if)#exit	Exit interface mode.
Step 4	(config)#interface xe2	Enter interface mode.
Step 5	(config-if)#exit	Exit interface mode.
Step 6	(config)#router bgp 65500	Configure the EBGP routing process with private ASN
Step 7	(config-router)#max-paths ebgp 8	Exit interface mode.
Step 8	(config-router)#neighbor 32.1.0.2 remote-as 64601	Configure maximum EBGP ECMP that can be installed in BGP.
Step 9	(config-router)#neighbor 32.1.0.2 fall-over bfd	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote private ASN
Step 10	(config-router)#neighbor 32.1.0.2 allowas-in	Configure BFD for the BGP session for faster failure detection.
Step 11	(config-router)#neighbor 32.2.0.2 remote-as 64601	Configure "allowas-in" for the neighbor to accept routes with same ASN learned over this neighbor
Step 12	(config-router)#neighbor 32.2.0.2 fall-over bfd	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote private ASN

	Command	Purpose
Step 13	(config-router)#neighbor 32.2.0.2 allowas-in	Configure BFD for the BGP session for faster failure detection.
Step 14	(config-router)#exit	Configure “allowas-in” for the neighbor to accept routes with same ASN learned over this neighbor
Step 15	(config-router)#exit	Exit Router mode.

## Tier-2 (Spine node)

	Command	Purpose
<b>Configure the interfaces</b>		
Step 1	(config)#interface xe1	Enter interface mode.
Step 2	(config-if)#ip address 32.1.0.2/24	Configure ip address on the interface
Step 3	(config-if)#exit	Exit interface mode.
Step 4	(config)#interface xe46	Enter interface mode.
Step 5	(config-if)#ip address 32.3.0.2/24	Configure an IP address on the interface
Step 6	(config)#interface xe47	Enter interface mode.
Step 7	(config-if)#ip address 32.4.0.2/24	Configure an IP address on the interface
Step 8	(config)#interface xe48	Enter interface mode.
Step 9	(config-if)#ip address 21.1.0.2/24	Configure an IP address on the interface
Step 10	(config)#interface xe46	Enter interface mode.
Step 11	(config-if)#ip address 21.2.0.2/24	Configure an IP address on the interface
Step 12	(config)#router bgp 64601	Configure the eBGP routing process with private ASN
Step 13	(config-router)#bgp bestpath as-path multipath-relax	Configure “as-path multipath-relax” to relax the AS-PATH exact match (if AS-PATH length are same) criteria for BGP ECMP
Step 14	(config-router)#max-paths ebgp 8	Configure maximum EBGP ECMP that can be installed in BGP.
Step 15	(config-router)#neighbor 32.1.0.3 remote-as 65000	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 16	(config-router)#neighbor 32.1.0.3 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 17	(config-router)#neighbor 32.3.0.3 remote-as 65001	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 18	(config-router)#neighbor 32.3.0.3 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 19	(config-router)#neighbor 32.4.0.3 remote-as 65002	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 20	(config-router)#neighbor 32.1.0.3 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 21	(config-router)#neighbor 21.1.0.1 remote-as 65534	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 22	(config-router)#neighbor 21.1.0.1 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 23	(config-router)#neighbor 21.2.0.1 remote-as 65534	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 24	(config-router)#neighbor 21.2.0.1 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 25	(config-router)#exit	Exit Router mode.



### Tier-3 (Core node)

	Command	Purpose
<b>Configure the interfaces</b>		
Step 1	(config)#interface xe1	Enter interface mode.
Step 2	(config-if)#ip address 21.1.0.1/24	Configure ip address on the interface
Step 3	(config-if)#exit	Exit interface mode.
Step 4	(config)#interface xe46	Enter interface mode.
Step 5	(config-if)#ip address 21.5.0.1/24	Configure an IP address on the interface
Step 6	(config)#interface xe49	Enter interface mode.
Step 7	(config-if)#ip address 41.1.0.1/24	Configure an IP address on the interface
Step 8	(config)#interface xe50	Enter interface mode.
Step 9	(config-if)#ip address 41.5.0.1/24	Configure an IP address on the interface
<b>Configure BGP on the router</b>		
Step 10	(config)#router bgp 65534	Configure the eBGP routing process with private ASN
Step 11	(config-router)#bgp bestpath as-path multipath-relax	Configure “as-path multipath-relax” to relax the AS-PATH exact match (if AS-PATH length are same) criteria for BGP ECMP
Step 12	(config-router)#max-paths ebgp 8	Configure maximum EBGP ECMP that can be installed in BGP.
Step 13	(config-router)#neighbor 21.1.0.2 remote-as 64601	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 14	(config-router)#neighbor 21.1.0.2 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 15	(config-router)#neighbor 21.5.0.3 remote-as 64602	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 16	(config-router)#neighbor 21.5.0.3 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 17	(config-router)#neighbor 41.1.0.3 remote-as 64603	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 18	(config-router)#neighbor 41.1.0.3 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 19	(config-router)#neighbor 41.5.0.3 remote-as 64603	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 20	(config-router)#neighbor 41.5.0.3 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 21	(config-router)#exit	Exit BGP mode

## Tier 2 (Border router)

	Command	Purpose
<b>Configure the interfaces</b>		
Step 1	(config)#interface xe1	Enter interface mode.
Step 2	(config-if)#ip address 41.1.0.2/24	Configure an IP address on the interface
Step 3	(config-if)#exit	Exit interface mode.
Step 4	(config)#interface xe46	Enter interface mode.
Step 5	(config-if)#ip address 41.2.0.2/24	Configure an IP address on the interface
Step 6	(config)#interface xe47	Enter interface mode.
Step 7	(config-if)#ip address 41.3.0.2/24	Configure an IP address on the interface
Step 8	(config)#interface xe48	Enter interface mode.
Step 9	(config-if)#ip address 41.4.0.2/24	Configure an IP address on the interface
Step 10	(config)#interface xe49	Enter interface mode.
Step 11	(config-if)#ip address 51.1.0.2/24	Configure an IP address on the interface
Step 12	(config)#interface xe50	Enter interface mode.
Step 13	(config-if)#ip address 51.3.0.2/24	Configure an IP address on the interface
<b>Configure BGP on the router</b>		
Step 14	(config)#router bgp 64603	Configure the eBGP routing process with private ASN
Step 15	(config-router)#bgp bestpath as-path multipath-relax	Configure “as-path multipath-relax” to relax the AS-PATH exact match (if AS-PATH length are same) criteria for BGP ECMP
Step 16	(config-router)#max-paths ebgp 8	Configure maximum EBGP ECMP that can be installed in BGP.
Step 17	(config-router)#neighbor 41.1.0.1 remote-as 65534	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 18	(config-router)#neighbor 41.1.0.1 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 19	(config-router)#neighbor 41.2.0.1 remote-as 65534	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 20	(config-router)#neighbor 41.2.0.1 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 21	(config-router)#neighbor 41.3.0.1 remote-as 65534	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 22	(config-router)#neighbor 41.3.0.1 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 23	(config-router)#neighbor 41.4.0.1 remote-as 65534	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 24	(config-router)#neighbor 41.4.0.1 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 25	(config-router)#neighbor 51.1.0.3 remote-as 100	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN

Configure BGP on the router		
Step 26	(config-router)#neighbor 51.1.0.3 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 27	(config-router)#neighbor 51.1.0.3 remove-private-AS	Configure “remove –private-AS” to remove the private ASNs for the routes advertised to this neighbor.
Step 28	(config-router)#neighbor 51.3.0.3 remote-as 101	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 29	(config-router)#neighbor 51.3.0.3 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 30	(config-router)#neighbor 51.3.0.3 remove-private-AS	Configure “remove –private-AS” to remove the private ASNs for the routes advertised to this neighbor.
Step 31	(config-router)#exit	Exit BGP mode

### Tier-3 (WAN router)

This is a partial and does not contain the Internet configuration.

	Command	Purpose
<b>Configure the interfaces</b>		
Step 1	(config)#interface xe1	Enter interface mode.
Step 2	(config-if)#ip address 51.1.0.3/24	Configure an IP address on the interface
Step 3	(config-if)#exit	Exit interface mode.
Step 4	(config)#interface xe46	Enter interface mode.
Step 5	(config-if)#ip address 51.2.0.3/24	Configure an IP address on the interface
<b>Configure BGP on the router</b>		
Step 6	(config)#router bgp 100	Configure the eBGP routing process with public ASN
Step 7	(config-router)#max-paths ebgp 8	Configure maximum EBGP ECMP that can be installed in BGP.
Step 8	(config-router)#neighbor 51.1.0.2 remote-as 64603	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 9	(config-router)#neighbor 51.1.0.2 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 10	(config-router)#neighbor 51.3.0.2 remote-as 64603	Configure the EBGP neighbor over the connected interface using the neighbor IP and remote ASN
Step 11	(config-router)#neighbor 51.3.0.2 fall-over bfd	Configure BFD for the BGP session for faster failure detection.
Step 12	(config-router)#exit	Exit BGP mode

#### Other Configurations

You must repeat similar configurations for all ToR, spine, core, border, and WAN devices as well.

## Validation

Use the show ip bgp command to validate the output at each node.

Consider the following case: for application load balancing and high availability/reliability, two similar application servers can be placed at two clusters. For users accessing the application server through the Internet, the access to the server is load balanced and failure of one of the application servers does not impact the accessibility. The following is the output at various nodes for a subnet, such as:

- 70.70.70.1 (application server) at ToR1 in cluster 1 and cluster 2
- 80.80.80.1 at ToR 1 cluster 1
- 90.90.90.1 at ToR 1 cluster 2

### ToR 1 Cluster 1

```
Tor-cluster1#show ip bgp
BGP table version is 2, local router ID is 34.34.34.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 70.70.70.0/24	0.0.0.0	0	100	32768	?
*> 80.80.80.0/24	0.0.0.0	0	100	32768	?
*> 90.90.90.0/24	32.1.0.2	0	100	0	64601 65534 64602 65500
>	32.4.0.2	0	100	0	64601 65534 64602 65500

Same AS\_PATH accepted

Total number of prefixes 3

### ToR 1 Cluster 2

```
Tor1-cluster2#show ip bgp
BGP table version is 2, local router ID is 31.31.31.31
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 70.70.70.0/24	0.0.0.0	0	100	32768	?
*> 90.90.90.0/24	0.0.0.0	0	100	32768	?
*> 80.80.80.0/24	32.7.0.2	0	100	0	64602 65534 64601 65500
>	32.10.0.2	0	100	0	64602 65534 64601 65500

Total number of prefixes 3

## R1 Cluster 1

```
r1-cluster1#show ip bgp
BGP table version is 4, local router ID is 35.35.35.35
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 70.70.70.0/24	31.1.0.3	0	100	0	65500
*	21.1.0.1	0	100	0	65534 64602 65500
*	21.5.0.1	0	100	0	65534 64602 65500
*> 90.90.90.0/24	21.1.0.1	0	100	0	65534 64602 65500
*	21.5.0.1	0	100	0	65534 64602 65500
*> 80.80.80.0/24	31.1.0.3	0	100	0	64602
					65534 64601 65500

Total number of prefixes 3

## R11 Tier 1

```
R11 #show ip bgp
BGP table version is 5, local router ID is 37.37.37.37
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 70.70.70.0/24	21.1.0.2	0	100	0	64601 65500
*>	21.5.0.2	0	100	0	64602 65500
*> 90.90.90.0/24	21.5.0.2	0	100	0	64602 65500
*> 80.80.80.0/24	21.5.0.2	0	100	0	64602 65500

Total number of prefixes 3

## Border Router B1 Tier 2 External Cluster

```
br1-ex-cluster #show ip bgp
BGP table version is 5, local router ID is 37.37.37.37
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 70.70.70.0/24	41.1.0.1	0	100	0	65534 64601 65500
*>	41.5.0.1	0	100	0	65534 64601 65500
*>	41.9.0.1	0	100	0	65534 64602 65500
*>	41.13.0.1	0	100	0	65534 64602 65500
*> 80.80.80.80/24	41.1.0.1	0	100	0	65534 64601 65500
*>	41.5.0.1	0	100	0	65534 64601 65500
*>	41.9.0.1	0	100	0	65534 64601 65500
*>	41.13.0.1	0	100	0	65534 64601 65500
*> 90.90.90.90/24	41.1.0.1	0	100	0	65534 64602 65500
*>	41.5.0.1	0	100	0	65534 64602 65500
*>	41.9.0.1	0	100	0	65534 64602 65500
*>	41.13.0.1	0	100	0	65534 64602 65500

Total number of prefixes 3

Multipath-relax:  
different ASPATH  
with same length  
marked as ECMP

Multipath-relax:  
different ASPATH  
with same length  
marked as ECMP

## WAN Router WR1 External Cluster

```
wr1-ex-cluster #show ip bgp
BGP table version is 5, local router ID is 37.37.37.37
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
                S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

	Network	Next Hop	Metric	LocPrf	Weight
*>	70.70.70.0/24	41.1.0.1	0	100	0
*>		41.5.0.1	0	100	0
*>	80.80.80.80/24	41.1.0.1	0	100	0
*>		41.5.0.1	0	100	0
*>	90.90.90.90/24	41.1.0.1	0	100	0
*>		41.5.0.1	0	100	0

Path
65500
65500
65500
65500
65500

Private AS  
sequences  
removed

## Conclusion

OcNOS with EBGP routing is highly scalable, simple and flexible way of laying IP fabric in a data center. The data center can be easily scaled for:

- Higher computing needs by adding more clusters.
- Higher performance and redundancy by adding more cores
- Higher uplink speeds by adding more external/edge clusters.

## References

Use of BGP for routing in large scale data centers:

<https://tools.ietf.org/html/draft-ietf-rtgwg-bgp-routing-large-dc-09>

## WAN Router WR1 External Cluster

```
wr1-ex-cluster #show ip bgp
BGP table version is 5, local router ID is 37.37.37.37
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
                S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight
*> 70.70.70.0/24	41.1.0.1	0	100	0
*>	41.5.0.1	0	100	0
*> 80.80.80.80/24	41.1.0.1	0	100	0
*>	41.5.0.1	0	100	0
*> 90.90.90.90/24	41.1.0.1	0	100	0
*>	41.5.0.1	0	100	0

Path
65500
65500
65500
65500
65500
65500

Private AS  
sequences  
removed

## Conclusion

OcNOS with EBGp routing is highly scalable, simple and flexible way of laying IP fabric in a data center. The data center can be easily scaled for:

- Higher computing needs by adding more clusters.
- Higher performance and redundancy by adding more cores
- Higher uplink speeds by adding more external/edge clusters.

## References

Use of BGP for routing in large scale data centers:

<https://tools.ietf.org/html/draft-ietf-rtgwg-bgp-routing-large-dc-09>

## Please contact us to learn more

Phone: +1 877-MYZEBOS  
Email: [sales@ipinfusion.com](mailto:sales@ipinfusion.com)  
Web: [www.ipinfusion.com](http://www.ipinfusion.com)

U.S. (Santa Clara), +1 408-400-1912  
Japan (Tokyo), +81 03-5259-3771  
Korea (Seoul) +82 (2) 3153-5224

India (Bangalore), +91 (80) 6728 7000  
China (Shanghai), +86 186 1658-6466  
EMEA (Stockholm), +46 8-566 300 42

## About IP Infusion

IP Infusion is a leading provider of intelligent network software for enhanced Ethernet and IP services. Tier one and two OEMs rely on IP Infusion's ZebOS software and global professional services to bring products to market faster, and to differentiate them from competitors with less cost. Products built on IP Infusion technology are deployed in networks with five-9s reliability across five continents—as well as a growing number of enterprises—to improve network performance, decrease network infrastructure costs, and grow revenue. IP Infusion is headquartered in Sunnyvale, Calif., and is a wholly owned and independently operated subsidiary of ACCESS CO., LTD., of Tokyo, Japan.

© 2016 IP Infusion, Inc. All rights reserved. OcNOS and IP Infusion are registered trademarks and the ipinfusion logo is a trademark of IP Infusion, Inc. All other trademarks and logos are the property of their respective owners. IP Infusion assumes no responsibility for any inaccuracies in this document. IP Infusion reserves the right to change, modify, transfer, or otherwise revise this publication without notice.