



# **Cumulus Quagga**

## **Bringing the routing from the switch to the host**

**Scott Suehle**

---

August 18, 2016

## Scott Suehle

- Community Manager for Cumulus Networks
- Board member for LOPSA
- Former customer support junkie
- [scotts@cumulusnetworks.com](mailto:scotts@cumulusnetworks.com)
- @8bitdad on the twitters
- kilted1 on #CumulusNetworks on Freenode



redhat

Duke




NICHOLAS SCHOOL OF THE  
ENVIRONMENT

EUCALYPTUS


- Cumulus Networks created a Debian-based (Wheezy and Jessie) network operating system
- The same L2 and L3 experience compared to legacy network vendors, but without proprietary APIs or controllers
- 41+ certified hardware platforms from 8 different vendors (Agema, Dell, Edgecore, HP, Mellanox, Penguin, QCT, Supermicro) with ONIE

# Transformation: First Servers, Now Networking



- Depend on legacy L2 networking for legacy applications.
- Layer 2 networking needed for:
  - Heartbeat
  - Discovery
  - Moving IP addresses around
- L2 domains segmented by use of VLANs

# Dual Attached Host Data Center Design



- Large broadcast and failure domains
- Spanning tree(STP) challenges
- Dependant on proprietary routing protocols for redundancy

That is great and all,

But that is not what you are  
here to hear about.




- Layer 3 all the way to the host and build a modern unified data center that is simple to scale and provides increased freedom
- Network Complexity Reduction
- Subnet Freedom and Mobility
- Increased Redundancy
- Stateless Anycast Load Balancing

Routing on the host gives you Layer 3 networking throughout the entire datacenter.

- This opens the data center to the usage of OSPF and BGP through the entire rack and makes it highly scalable.
- This makes troubleshooting network issues in the data center easier to diagnose.


# Advantage 1: Reduction in Network Complexity



With routing on the host, your IP address is independent of its location in the datacenter.

- Host can now advertise on a /32 network.  
Allowing for a more granular IP assignment.
- No VLANs or switched virtual interfaces are needed.
- No configuration changes needed for moving hosts and connecting to a different leaf switch.

## Advantage 2: Subnet Freedom and Mobility



Routing on the host give server greater network redundancy by connecting to more leaf switches in the data center.


- This helps reduce bandwidth drop off in case of hardware failure.
- Routing configuration can be set to allow the host and spine to avoid downed leaf hardware and avoid traffic loss.

## Advantage 4: Stateless Anycast Load Balancing Functionality

Routing on the host brings with it the ability to enable load balancing through the use of anycast addressing.

- Hashing on the leaf switches allow multiple hosts to broadcast the same IP address range, distributing connections from remote nodes among those hosts.

# Use Anycast for Load Balancing between Servers




\* More information on resilient hashing can be found here:  
<http://bit.ly/1N0Eogd>



# Layer-3 to the Host: Multiple Attach (Quagga on the Host)

- Overall: The best overall networking solution with OpenStack and Cumulus Networks in large configurations. 100% simple and flexible architecture with Layer-3 networking using Linux quagga package extendable to other software solutions.



- Easy Troubleshooting and Maintenance
- Subnet Freedom
  - Ability to move hosts with impunity
    - Route moves with host

## Improve Redundancy and Bandwidth

- Connect server to “n” leaf switches
- Free up 2 additional leaf ports per switch
- Anycast for Load Balancing

# Cumulus and Openstack



## Installing, Configuring, and Managing a 300+ OpenStack Node...

... network in 15 minutes  
... overcloud in 6 hours

Nolan Leake — Co-Founder, Cumulus Networks  
Brad Watkins — NFV Partner Engineer, Red Hat  
Jaiwant Virk — Senior Product Manager, Dell

April 28, 2016 @ OpenStack Summit Austin 2016

[cumulusnetworks.com](http://cumulusnetworks.com) | [redhat.com](http://redhat.com) | [dell.com](http://dell.com)

### Step 0: Find a lab, thank you to Dell!



[cumulusnetworks.com](http://cumulusnetworks.com) | [redhat.com](http://redhat.com) | [dell.com](http://dell.com)

# Cumulus and Openstack

- Can we quickly deploy the compute nodes?
- Can the entire pod be prototyped virtually first?
- Can the entire pod be entirely Layer 3 with VXLAN?
- Can Ansible and Git be the common “language” between Network Engineers and System Admins?
- Can RHOSP Director handle the bulk deployment?
- Can the entire project be done remotely with ease?

# Links and Resources

Git: Ansible playbooks, quagga package builds and deployment:

- <https://github.com/leifmadsen/quagga-config-deploy>
- <https://github.com/leifmadsen/quagga-rpm>
- [https://github.com/CumulusNetworks/dellrh\\_openstack\\_ansible](https://github.com/CumulusNetworks/dellrh_openstack_ansible)

I know what you are saying to yourself...

Where can I find out more about this?



- <https://github.com/CumulusNetworks/quagga>
- <https://cumulusnetworks.com/routing-on-the-host/>
  - Opensource project
  - Come join us and get involved



## Shapeways Case study

<https://cumulusnetworks.com/media/resources/case-studies/Cumulus-Networks-Case-Study-Shapeways-Routing-on-the-Host.pdf>


## Webinar

<http://go.cumulusnetworks.com/roh-wbr>



## So, what's next?

- Keep on learning
- Join the community <http://community.cumulusnetworks.com/>
- Participate in discussions




**Download Cumulus VX** to work on proofs of concept, automation building, and continued practice: <https://cumulusnetworks.com/cumulus-vx/>

**Use lab guides and demo modules** to learn at your own pace:  
<https://support.cumulusnetworks.com/hc/en-us/articles/201787686> and  
<https://support.cumulusnetworks.com/hc/en-us/sections/200398866>



## Bringing the Linux Revolution to Networking



# Thank You!

© 2014 Cumulus Networks. CUMULUS, the Cumulus Logo, CUMULUS NETWORKS, and the Rocket Turtle Logo (the “Marks”) are trademarks and service marks of Cumulus Networks, Inc. in the U.S. and other countries. You are not permitted to use the Marks without the prior written consent of Cumulus Networks. The registered trademark Linux® is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a world-wide basis. All other marks are used under fair use or license from their respective owners.