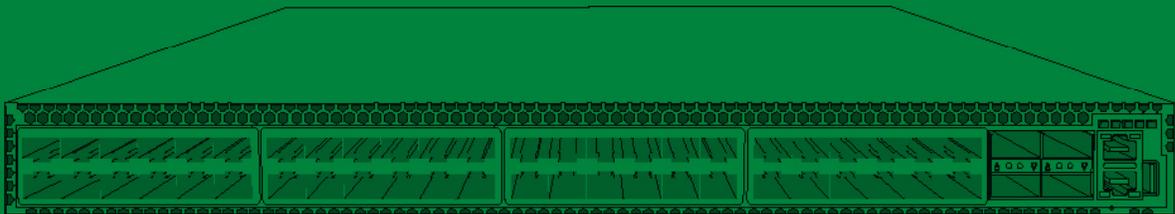


The Benefits of Open Networking and Software-Defined Storage



WHITE PAPER





Table of Contents

Introduction	1
Combined Open Networking and Storage Benefits	2
Scality RING Features	2
Cumulus Linux Features	3
Common Management and Monitoring Using Linux Native Tools	4
Convenient Locality	5
Load Balancing Achieved through ECMP with Anycast	6
Conclusion	7

Introduction

At the OpenStack Summit: Austin, the CTOs of Scality and Cumulus Networks discussed how Open Networking with Software-Defined Storage (SDS) enables use of the optimal hardware and system software to support the needs of a wide range of applications. Leveraging the open platform model on bare metal switches and servers, the unified network and storage solution also provides a common automation, management, and monitoring framework.

The OpenStack project was started to enable organizations build cloud-computing services leveraging standard hardware. Today, OpenStack has become the largest open source initiative since Linux. It is powered by a massive community of architects and engineers as well as a large number of companies that have been contributing to the project with code, applications, documentation, and usage feedback.

Two areas that have shown particular value to the OpenStack project are Open Networking and SDS as these technologies have enabled users to choose the best hardware and operating systems for their networking and storage needs. Open Networking and SDS benefit from an unlocked model that increases flexibility: Cumulus Linux gives users freedom in choice for network hardware while Scality lets users build geo-distributed, petabyte scale unstructured data storage infrastructure for any OpenStack application deployed on industry-standard servers. This paper covers the benefits of Cumulus Networks and Scality operating together in a unified data center.



Combined Open Networking and Storage Benefits

Combining Linux-powered Open Networking and SDS creates even more opportunities as this allows for data center unification through unified automation, management, and monitoring. This open platform enables common management and monitoring tools, server/storage placement flexibility, and automated load balancing. The combined model lets customers choose the best hardware and operating system, and enable well established features of the Scality RING and Cumulus Linux for their storage and networking infrastructure.

Cumulus Linux runs on a variety of bare metal switches while Scality runs on a variety of bare metal Linux servers. This paper covers a few additional benefits of Cumulus Networks and Scality operating together in a unified data center. A solution that includes Cumulus Linux as your network OS and Scality for storage provides:

Common management & monitoring tools

Convenient locality solutions

Automatic load balancing

Scality RING Features that Enable Open Storage Infrastructures

Scality's Software-Defined Storage solution provides two essential layers of flexibility that enable users to build open storage infrastructures: choice of hardware and architecture.

The Scality RING is fully hardware agnostic, without the need to certify specific hardware platforms. Customers have complete freedom from a wide selection of performance and capacity-optimized servers and can use multiple generations of servers, even simultaneously in the same RING. This way, customers can leverage hardware innovations immediately without the 9-12 month lag of appliance qualification for new technology. This means the latest in performance, capacity density, cost per GB, and energy savings can be adopted instantly.

To enable site-level disaster recovery solutions, Scality offers two geo-distribution options, whereby the RING can be deployed across

multiple sites (data centers) with a failure tolerance of one or more sites. The first — the Mirrored RING — involves independent RINGs, each within its own data center, with asynchronous mirroring employed to eventually manage synchronization between the RINGs. The second option is the Stretched RING, which makes use of a single logical RING deployed across multiple data centers and thus requires no synchronization.

On top of that, the RING provides data durability through a set of flexible protection methodologies and mechanisms optimized for distributed systems that allow applications to select the best data protection strategies for their data: Scality's data replication and erasure coded data protection methodologies provide extremely high levels of data durability, with the ability to fine tune performance and space characteristics for different data types.



Cumulus Linux gives users ability to adopt web-IT methodology to build better, faster and easier networks. Scality lets users build geo-distributed, petabyte scale unstructured data storage infrastructure

Cumulus Linux Features that Enable Open Networking Infrastructures

Cumulus Linux provides three main aspects to allow customers to build open networks: a Linux operating model along with a separate virtual platform, a rich set of networking features in hardware, and a robust hardware ecosystem.

Since Cumulus Linux is Linux, all applications available with a Linux operating system are also available with Cumulus Linux. Linux is very mature, has a vibrant ecosystem, very open API, and fosters community involvement. The entire community can write and add applications to their network that are essential to them, without waiting for a vendor's roadmap. Additionally, Cumulus has developed Cumulus VX, a virtualized platform that allows customers to build and test entire networks right on their PC, and is available to customers for free.

A rich networking feature set is also available. Cumulus Linux supports features such as BGP, OSPF, VRF, MLAG, VLAN, VXLAN, and Access Lists, to name a few. The feature set supports modern data center architectures and allow for robust networks.

With Cumulus Linux, customers have a vast array of hardware options. Cumulus Linux supports customers' choice of cables, optics, silicon, and bare metal systems and CPUs — independently. Cumulus Linux currently runs on 40 hardware platforms from 8 different vendors. This gives customers a choice of hardware and avoids any potential hardware lock in or supply chain issues.



Common Management and Monitoring Using Linux Native Tools

Linux is the ultimate operating system for the data center. Since Linux has an open API, customers can deploy applications written by the community, and it's trivial to write custom tools that are applicable to their data center.

Scality Software-Defined Storage runs on Linux in the userspace without any modification to the kernel, while Cumulus Linux IS Linux. Using the Linux operating system for your network infrastructure, servers, and storage unifies the data center in terms of network configuration, monitoring, and management. The ability to use the same tools enables rapid availability of applications to the end user — making deployment and management simple and easy.

Open source-based common configuration management tools such as Ansible, Salt,

Puppet, or Chef can configure and manage the deployment of network, compute, and storage nodes simultaneously. Using one configuration management tool provides assurance of consistent configuration settings across the entire stack. They help deploy the network and the storage solutions quickly by automating the configuration, thus bringing up all aspects of the data center simultaneously if desired. Additionally, you can use common monitoring tools — such as Nagios or Sensu, among others — for both storage and networking. These tools monitor all the system infrastructure components including applications and services. Using all the same tools throughout the entire data center streamlines bringing new applications to market as both networking and storage can go online simultaneously. See Figure 1.

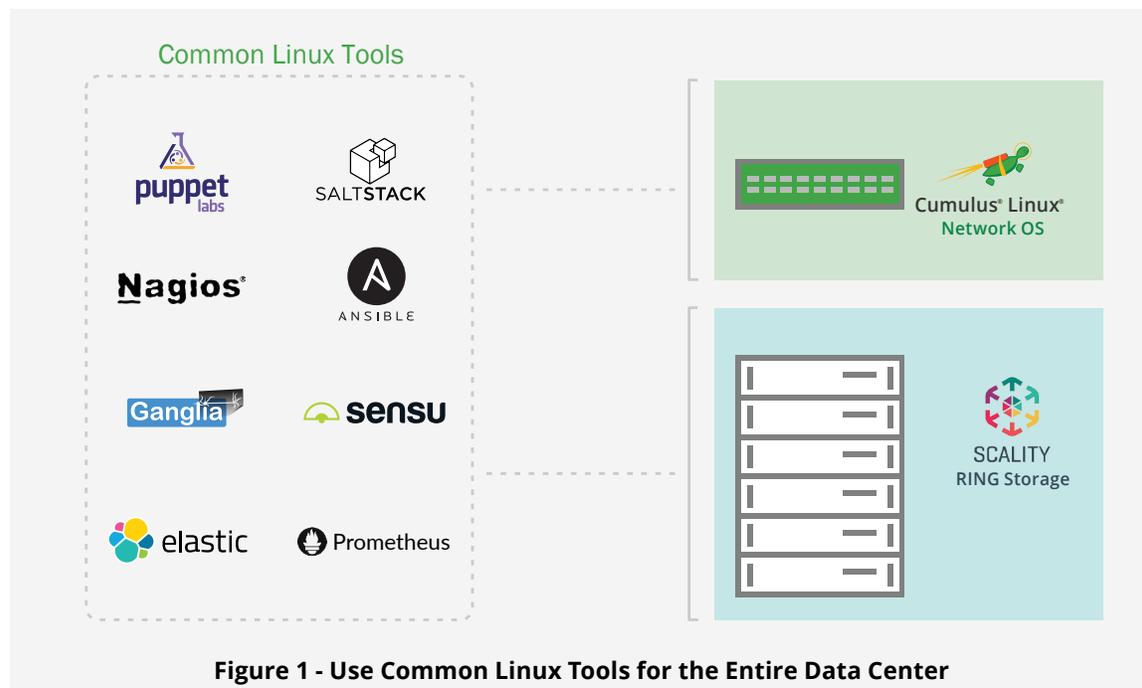


Figure 1 - Use Common Linux Tools for the Entire Data Center



Further, an operator can run various Linux commands to access system level statistics, gather critical information about the system, and establish a converged platform for storage and network nodes alike. System performance-associated tools based on Linux such as iostat and iotop provide important insights. Data I/O flows to the storage nodes require strong understanding of the network system, and Linux-based tools such as tcpdump help provide insights into the network traffic that affect storage. Finally, since Linux is an open environment, customers can run any tool that is supported by Linux — they can use tools created by the community or even write their own in languages such as Python and Java.

The ability to use the same Linux tools for the entire data center enables rapid availability of applications to the end user - making deployment, management and monitoring simple and easy.

Convenient Locality

Cumulus Linux running in a leaf/spine Clos network architecture with Scality servers in the rack eliminates the problem of required storage locality and provides flexibility in server/storage location within the data center. The storage servers or nodes may be spread between different racks or placed on the same rack.

Additionally, several storage servers hosting the same data are no longer required to be in the same rack to talk with each other to achieve top

performance, as the simple Clos design allows easy bandwidth-rich connectivity. Using a layer 3 fabric in the data center along with Equal Cost Multipath Routing (ECMP) enables all links to be used concurrently without the hassle of spanning tree.

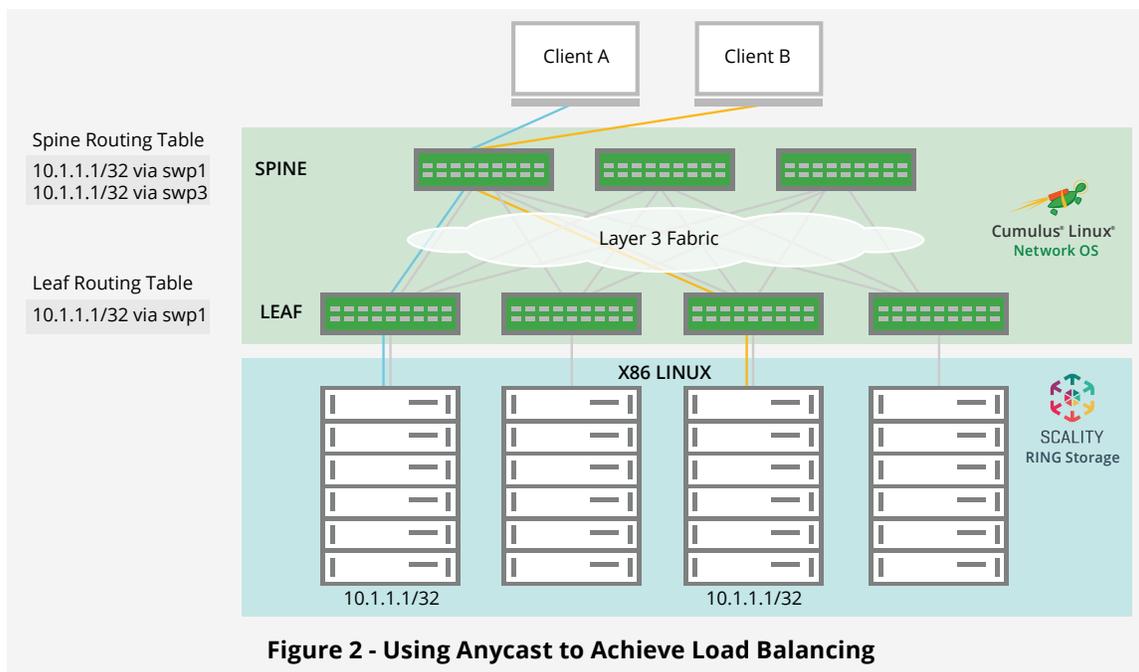
It is desirable to host the storage servers on different racks from each other to enhance resiliency — if one rack goes down, the data is still available.

Load Balancing Achieved through ECMP with Anycast

Another advantage of the converged solution is the ability to provide simplified and automatic load balancing through the use of anycast. By assigning the same IP address to several servers/connectors, and advertising each connector's /32 IP address into the routing domain, several routes to the same "destination" are advertised and seen by the routers. Each router then has an equal cost multipath (ECMP) route to the same IP address, although in reality each IP host route in the routing table equates to a different server. Naturally, each connector hosting the same IP address should be running the same applications and services.

When a client tries to connect to the connector IP address, the first hop router with the ECMP routes performs a hashing algorithm based on attributes such as source/destination IP address, port, and so forth. Since these attributes are the same for each flow, the packets from a specific client to a specific port on a connector take the same path through the network, and thus reach the same connector each time. Enabling resilient hashing ensures that if an ECMP link goes down, the hashing calculations do not change and the existing flows do not get recalculated.

Since there are several clients with different source IP addresses, they end up reaching several different connectors, thus distributing the load across the data center. See Figure 2.





Conclusion

Using Cumulus Linux enables you to operate a switch as if it were a Linux server, and Scality to operate storage as a Linux server, so it's like a developer's paradise: full of customizations for enhancing your specific needs beyond a typical API handoff. The ability to manage a network and storage infrastructure that scales linearly is critical to meeting the needs of various application workloads and the SLAs required for today's digital business.

Cumulus Linux along with Scality Software-Defined Storage enables a consistent experience between the network, compute, and storage while leveraging existing investments in orchestration tools and driving the next wave of scale, collaboration, and innovation in networking and storage.

The solution provides vast flexibility and resiliency by enabling the distribution of servers and storage throughout the data center. Careful placement of servers and storage can help protect against outages due to rack power loss. Features such as using anycast for load balancing adds more value by eliminating additional hardware in the network.

Combining Cumulus Linux with Scality Software-Defined Storage provides the foundation for modern data centers. This combined solution can enable tremendous advantages in economics versus the traditional non-standard hardware centric approaches, more flexibility in management, and efficiency in operations. We invite you to try the solutions and learn more by following the links below.

Get Started!

Try Cumulus Linux on Cumulus VX:
cumulusnetworks.com/cumulus-vx/

Try the Scality RING:
www.scality.com/trial/

Get the OpenStack and Cumulus Linux Validated Design Guide:
go.cumulusnetworks.com/OpenStack-Design-Guide

Get the Integrating Scality RING into OpenStack white paper:
www.scality.com/resources/

About Scality

Scality is the industry leader in software-defined storage for the information age, serving over 500 million users worldwide. Scality RING provides 100% reliable Software-Defined Storage with unmatched performance, and is perfect for capacity-driven workloads such as cloud services, high definition video, and enterprise archiving. It runs on standard x86 servers powered by Linux, (e.g. HP, Dell, Cisco, SuperMicro, or Seagate), and creates an unlimited storage pool for file, object and OpenStack applications. Thanks to its underlying object storage architecture, the RING scales to exabytes of data and trillions of files. Seven of the top twenty telecommunication companies, petabyte-scale cloud services, and Global 2000 enterprises rely on the RING for mission-critical applications.

For more information visit scality.com, or follow [@scality](https://twitter.com/scality).

About Cumulus Networks®

Cumulus Networks helps customers realize cost-effective, high capacity networking for modern data centers. Linux transformed the economics and innovation for data center compute, and Cumulus Linux is doing the same for the network. It radically reduces the costs and complexities of operating modern data center networks for businesses of all sizes. Cumulus Networks has received venture funding from Andreessen Horowitz, Battery Ventures, Sequoia Capital, Peter Wagner and four of the original VMware founders.

For more information visit cumulusnetworks.com, or follow [@cumulusnetworks](https://twitter.com/cumulusnetworks).

©2016 Cumulus Networks. All rights reserved. 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.

04252016 - 1C