



WHAT'S UNIQUE ABOUT NVIDIA CUMULUS?





10 TECHNICAL REASONS WHY NVIDIA CUMULUS PROVIDES UNPARALLELED TECH

Building a web-scale network with NVIDIA® Cumulus® enables efficiency, scalability and cost reduction. The business benefits are clear, but it's the technical innovations that make them possible.

So what makes NVIDIA Cumulus really special from a technical standpoint?

Here are 10 items to consider.

1: CUSTOMERS HAVE ACCESS TO OUR SOURCE CODE

With traditional, legacy network providers, an organization has very limited access to the source code, meaning you're dependent on a third party for troubleshooting and support. Since NVIDIA Cumulus Linux is open and based in Linux®, you have full access to most of the code. 90% of the code we write is python so you can view, investigate and modify whatever you like directly on the switch. We believe this empowers the user to troubleshoot under his/her own expectations and get support from the community quickly. This means less downtime, quicker bug fixes and continually optimized code.

Plus, since you have access to the code, you can edit on the box natively, with any language you prefer. No more guessing about what's working because you're

stuck behind a black box, and no more long calls to support for something you could fix yourself if you only had access. You'll speed up delivery of fixes because minor changes can be made quickly, by you, without waiting for a vendor to release an entirely new binary release.

Nearly all NVIDIA networking features, from adapters to switches to data processing units (DPU's) can be accessed and managed with open interfaces. In fact, NVIDIA is one of the top three Linux kernel contributors and a top contributor to projects such as SONiC, the Data Plane Developers Kit (DPDK), and Free Range Routing (FRR).

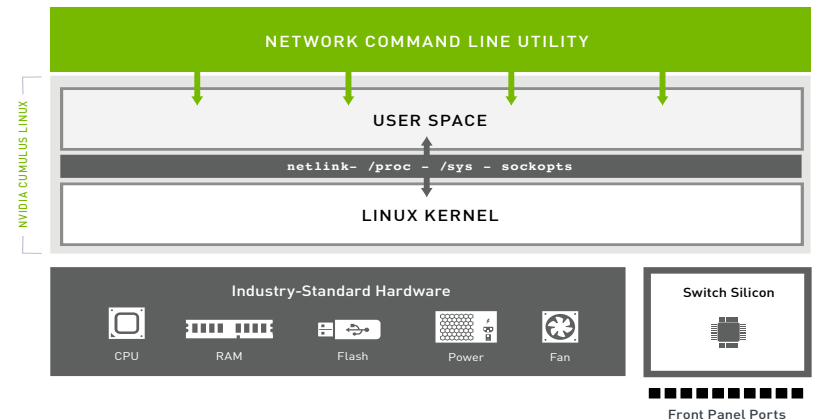


2: YOU GET THE BENEFITS OF LINUX WITHOUT HAVING TO LEARN IT

You don't need to know Linux to configure with NVIDIA Cumulus Linux (unless, of course, you want to). For those that know Linux, you already know that by having an operating system that is Linux, we completely unify the stack so that all apps and pieces of the data center are speaking the same language. Plus, you'll always have the power of the Linux community behind you. So you don't have to wait on support for bug fixes or security patches (again, unless you want to).

For those seeking a command line interface, we have a simplified CLI that takes 92 lines of code down to a very manageable 11. The Cumulus CLI is a modern CLI solution that resides in the Linux user space, as seen below. It provides consistent access to networking commands directly via bash, thereby making configuration and troubleshooting simple and easy — no need to edit files or enter modes and sub-modes. In addition our CLI offers way more than a traditional CLI by:

- > **Embedding help, examples and automatic command checking with suggestions in case you've entered a typo**
- > **Running directly from and integrating with bash, while being interoperable with the regular way of accessing underlying configuration files and automation**
- > **Providing audit trails so you can look back in time and see how a configuration ended up there in the first place**
- > **Automatically configuring dependent features so you don't have to**
- > **Giving you guard rails so you can get to know Linux networking without risk**



The wrapper utility is called net. net is capable of configuring layer 2 and layer 3 features of the networking stack, installing ACLs and VXLANs, rolling back and deleting snapshots and providing monitoring and troubleshooting functionality for these features. This includes full configuration of all interfaces and advanced routing functionality, including BGP and OSPF. Our simplified CLI provides consistent access to networking commands directly via bash, thereby making configuration and troubleshooting simple and easy — no need to edit files or enter modes and sub-modes.

3. YOU CAN PROVISION A LARGE-SCALE DATA CENTER IN MINUTES

Not really into manually provisioning your network box-by-box? We're not either. While developing and testing the provisioning logic, you can use ZTP in NVIDIA Cumulus Linux to manually invoke your provisioning script on a device.

Zero Touch Provisioning (ZTP) enables network devices to be quickly deployed in large-scale environments. On first boot, Cumulus Linux will invoke ZTP, which executes the provisioning automation used to deploy the device for its intended role in the network.

The provisioning framework allows for a one-time, user-provided script to be executed. You can develop this script using a variety of automation tools and scripting languages, providing ample flexibility for you to design the provisioning scheme to meet your needs. You can also use it to add the switch to a configuration management (CM) platform such as Puppet, Chef, CFEngine or possibly a custom, proprietary tool.

ZTP in Cumulus Linux can occur automatically in one of the following ways, in this order:

- > Via a local file
- > Using a USB drive inserted into the switch (ZTP-USB)
- > Via DHCP

Although ZTP is offered by a variety of vendors, we believe we do it better. For starters, the provisioning files reside in Git, not on the box. If the box goes down, you can get the provisioning information downloaded from the cloud and bring a new box up the same way you had it before. Plus, we're just Linux. As long as the script runs, we don't care how you design the script. Write it however you want. Most other systems are much more bounded in terms of what is permitted. So in short, although other vendors offer ZTP, our functionality offers unmatched flexibility.





4. SETTING UP & TROUBLESHOOTING CABLING IS NO LONGER A DREADED TASK

Cabling is one necessary evil in networking that can quickly become exceptionally complex. In data center topologies, correct cabling is a time-consuming endeavor and is error prone. Prescriptive Topology Manager (PTM) is a dynamic cabling verification tool that helps you detect and eliminate such errors. It takes a Graphviz-DOT specified network cabling plan (something many operators already generate), stored in a topology.dot file, and couples it with runtime information derived from LLDP to verify that the cabling matches the specification. The check is performed on every link transition on each node in the network. Essentially, we took the complexity out of cabling.

To make matters even more magical, we also offer unparalleled fabric visibility and actionable insight with NVIDIA Cumulus NetQ. In addition to many other features, NetQ works together with PTM to locate a specific cabling or node issue within the data center. You no longer have to go box-by-box to find the issue. Troubleshooting becomes a breeze.

5. CONFIGURATION IS CRAZY SIMPLE WITH BGP UNNUMBERED AND AUTOMATION-READY CODE

Once again, NVIDIA Cumulus Linux is Linux So you can run any automation application or script to automate your network to your heart's content. You're immediately ready to configure Ansible, Puppet, Chef and more. This is a starkly different experience than working with traditional networking vendors, where you're dependent on proprietary automation software or proprietary scripts to make them work. When you add in NVIDIA Cumulus NetQ, you bring automation to the operations side of your network. You'll get fabric-wide insight from containers to the switch so you can troubleshoot, validate and diagnose issues quickly.

Automation gets even easier with BGP Unnumbered and its simplified IP approach. All you need is one IP template for leaf nodes and one for the spine nodes (unlike other vendors where you need a different configuration file for each device). Unnumbered interfaces are interfaces without unique IP addresses. In BGP, you configure unnumbered interfaces using extended next-hop encoding (ENHE), which is defined by RFC 5549. BGP Unnumbered interfaces provide a means of advertising an IPv4 route with an IPv6 next-hop. Prior to RFC 5549, an IPv4 route could be advertised only with an IPv4 next-hop.

BGP Unnumbered interfaces are particularly useful in deployments where IPv4 prefixes are advertised through BGP over a section without any IPv4 address configuration on links. As a result, the routing entries are also IPv4 for destination lookup and have IPv6 next-hops for forwarding purposes.

And now NVIDIA Cumulus offers production-ready automation — allowing you to automate with the click of the button. No other vendor out there delivers such a robust, production-ready automation solution that is validated and ready to go out of the box. This suite of automation and testing includes:

- > **Ready-to-go automation using Ansible. It's essentially a cut and paste.**
- > **A fully populated variables file object model**
- > **Complete Jinja2 templates**
- > **Ansible playbooks and a full battery of network validation tests enabling continuous integration (CI).**
- > **The complete and ideal framework for a tangible "infrastructure-as-code" deployment**
- > **Utilizes Ansible core modules without the need for any add-ons or plugins**
- > **Open source and available on GitLab.**

NVIDIA Cumulus Linux was designed for automation first. With unmatched flexibility, how and how much you automate your network is completely up to you.

6. MOST ADVANCED CAPABILITIES AVAILABLE FOR EVPN, LAYER 3 CONNECTIVITY

NVIDIA Cumulus Linux offers cutting edge features that move our customers beyond legacy layer-2 networks and into Layer 3 virtualization with EVPN.

Networks can provide connectivity via bridging (L2) or routing (L3). Thus, virtual networks can be either virtual L2 networks or virtual L3 networks. EVPN is a technology that connects L2 network segments separated by an L3 network. EVPN accomplishes this by building the L2 network as a virtual Layer 2 network overlay over the Layer 3 network. It typically uses Border Gateway Protocol (BGP) as its control protocol.

An EVPN-VXLAN-based architecture lets you deploy a common set of policies and services across data center and campus networks with support for L2

and L3 VPNs. Using a Layer 3 IP-based underlay coupled with an EVPN-VXLAN overlay, network operators can deploy much larger networks than would otherwise be possible with traditional L2 Ethernet-based architectures.

By using Cumulus EVPN-BGP as the control-plane solution for VXLAN tunnels, data center operators gain numerous, significant advantages including scalability, simplicity, vendor interoperability, controller-less VXLAN tunnels, more efficient traffic flows, fast convergence and host mobility, VXLAN active-active mode, multi-tenancy and a variety of cost efficiencies.

NVIDIA Cumulus Linux offers cutting edge features that move our customers beyond legacy layer-2 networks and into Layer 3 virtualization with EVPN.

7. NETWORKING BUILT FOR THE FUTURE, FIRST

One of the best things about NVIDIA Cumulus is that we're not simply adjusting to customer demands for a future-proofed network — we built our software for the future from the very beginning. Although many vendors claim to support things like automation, layer 3 networking and NetDevOps, we believe we do it better because we incorporated these things into the development of NVIDIA Cumulus Linux from day one. The future of networking is disaggregation, and we are the very organization that pioneered the movement. Although specific features may be replicated by other vendors or in other capacities, we've been doing this much longer and thus have a significant headstart on all of these concepts — and you can leverage them in your network today and we don't upcharge you for them.

Software that we have developed and released to the Linux community, like FRRouting and ifupdown2, have helped take Linux networking to a new level both in functionality and in hardening. Ifupdown2 is best-in-class interface configuration software that fixes many of the challenges in legacy network interface management. It handles network interface configuration dependencies, minimizes surface disruption, uses an enhanced network interface configuration file to simplify configuration and provides options to the user to query and validate running interface configuration with user provided config. Thus, it provides the ability to detect and debug problems with interface configuration.

It stands head and shoulders above other such interface configuration offerings elsewhere in Linux and in the proprietary vendor realm. Plus, it's open source and you can install it anywhere now, not just on NVIDIA Cumulus Linux switches.

Similarly, FRRouting is an IP routing protocol suite for Unix platforms, which includes protocol daemons for BGP, IS-IS, LDP, OSPF, EIGRP, PIM and RIP, and the community is working to make this the best routing protocol stack available and is already widely adopted including Cumulus Linux and Sonic. It's projects like these that define our future-first philosophy.

In addition to innovative protocols, we've also designed our network operating system for both layer 3 and layer 2 architectures. We support FRRouting, BGP and OSPF for modern age layer 3 networks. And although we believe the future of networking is in layer 3 solutions, we know that many customers aren't ready to make such a significant operational change. This is why we also offer a robust EVPN solution so that you can seamlessly extend layer 2 VLANs over a modern layer 3 fabric.

8. AUTOMATION AS A FORCE MULTIPLIER

We already discussed how ZTP allows you to handle initial configurations in minutes, but initial configuration is just one facet of automation. For Day 1 and Day 2 operations, having robust automation is an absolute must to get the most out of your data center infrastructure and maximize the utilization of your human infrastructure.

Here we reap significant advantages from being a Debian-based Linux distribution. We can effortlessly utilize common automation tools like Ansible, Puppet and Chef as well as uncommon ones or custom scripts to better support your application. You can even install Ansible directly on a NVIDIA Cumulus Linux switch and use your out-of-band switch to deploy Ansible playbooks against the rest of the infrastructure — how's that for doing double duty?

In addition to having advantages in the breadth of tooling that is supported, we also have benefits in the method that automation can be performed. Let's consider the case where you have 50,000 lines of configuration that need to get deployed to your switch — that's a LOT of config by almost any standard.

With automation tooling and traditional vendors, your automation is left to deploy configuration line-by-line using mostly the CLI as the delivery mechanism for all 50,000 lines. This will obviously take some time and you're forced to sacrifice some agility as a result. With NVIDIA Cumulus, you have no such restrictions on your automation.

You can leverage any of the thousands of modules that have been written to control and automate your Linux devices on your NVIDIA Cumulus switch. Take for example Ansible's "template" module. Using this module, you can write your entire 50,000 interface configuration in 1 second and apply all the changes at once. This will shave lifetimes off of your automation changes and potentially off of responding to a change required to remediate an outage somewhere in the environment. In our [NVIDIA Cumulus Linux and Ansible blog](#), we teased out some of the places where using the NCLU module makes good sense. However, when you're automating large environments, there's no substitute for the uninhibited access you get on a Linux system like NVIDIA Cumulus.



9. GET HOLISTIC AND REAL-TIME VISIBILITY, TROUBLESHOOTING AND LIFECYCLE MANAGEMENT

NVIDIA Cumulus NetQ is a highly-scalable, modern network operations tool set that leverages fabric-wide telemetry data to provide visibility and troubleshooting of your overlay and underlay network in real time. NetQ is the leading product for network telemetry, validation and one-of-a-kind network CI/CD integration capabilities. Included in the updated software are Lifecycle Management (LCM) capabilities that enable NetQ to deliver upgrade and configuration management with push button simplicity helping to remove complexity and open the door to adoption of open networking.

Managing network assets correctly is critical to avoiding application-impacting network outages, performance degradation and security incidents. Network operators often struggle with operational challenges such as network disruption caused by maintenance and configuration changes. To help solve these issues, the new LCM capabilities offer a simple, GUI-driven workflow for deploying, configuring, maintaining and retiring network switches running NVIDIA Cumulus Linux.

With the addition of full Lifecycle Management functionality, NetQ now combines the ability to easily upgrade, configure and deploy network

elements with a full suite of operations capabilities, such as visibility, troubleshooting, validation, trace and comparative look-back functionality. The addition of LCM complements customer deployments of Cumulus Linux and helps to further reduce operations costs.

NetQ provides a push-button solution for software upgrades and configuration changes, leveraging existing Cumulus Linux automation flows such as Zero Touch Provisioning (ZTP). To minimize risk of network disruption, NetQ also captures and compares the live state of the network before and after the maintenance (or configuration change) with an easy-to-use graphical user interface.

Key benefits of LCM include streamlining and automating the upgrade process, substantial reductions in network operations costs, reduced security risks, more efficient use of software licenses and maximizing the value of network infrastructure. Other benefits include lower overall TCO, fewer unexpected outages and more predictable control throughout the software lifecycle, eliminating the need to hire Linux experts to run network operations.

Key benefits of LCM include streamlining and automating the upgrade process, substantial reductions in network operations costs, reduced security risks, more efficient use of software licenses and maximizing the value of network infrastructure.

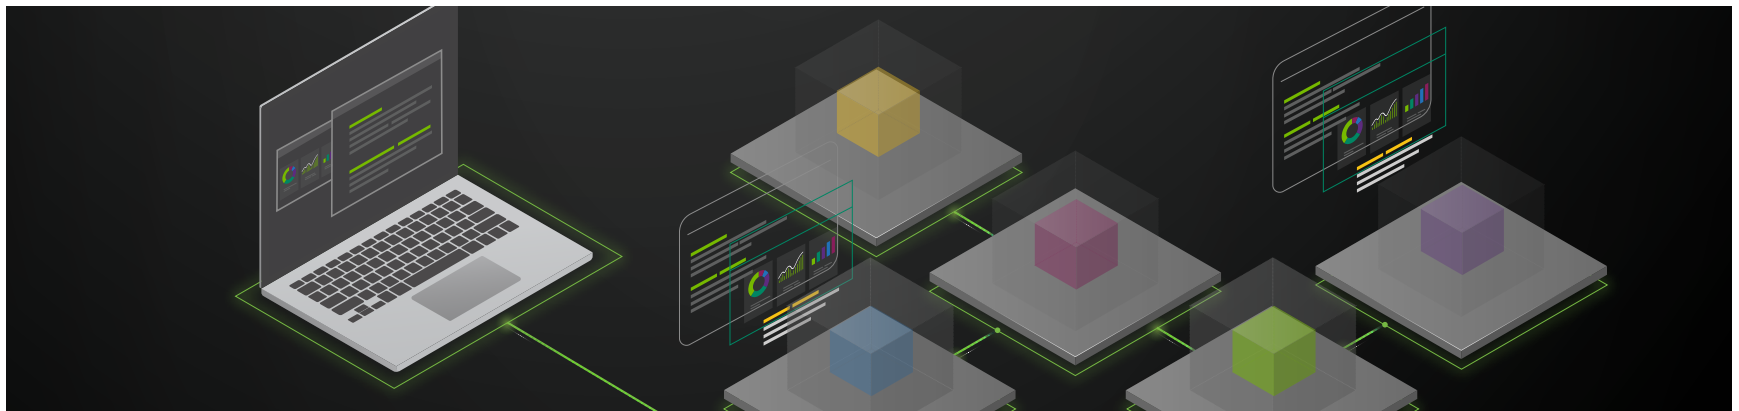
10. CONTAINER NETWORKING TRACKING AND VISIBILITY

In addition to receiving information from all the switches and hosts in the network, NetQ can directly access the Kubernetes API. NetQ taps into the k8s API server master to view the health, status and connectivity of the nodes and the workloads in the cluster, including their networking characteristics. By tapping into the k8s and Docker, NetQ can easily determine and display the new pod/container locations and how they are connected to the network. With Kubernetes, NetQ also supports the Container Networking Interface (CNI) ecosystem, specifically the popular CNI's of Calico or Flannel, giving you more choice in your container interfaces.

NVIDIA Cumulus NetQ is uniquely positioned to help networking teams embrace containers since it integrates directly with both container orchestrators and the Netlink interface into the Linux kernel, to provide real-time access to Linux networking events across the data center fabric. NetQ shows the container state, what changed, which containers have been re-deployed from one worker to another, shows the connectivity matrix, and shows impact of a potential ToR outage.

NetQ gives container service insight and provides the tools engineers need to design, update, manage and troubleshoot the container network. It is the only network based telemetry agent which directly integrates with Kubernetes. With this deep insight, you can now monitor containers as they change and trace connectivity to their ports.

In summary, NetQ is the only major telemetry product that interfaces with Netlink for real-time access to Linux kernel networking events. NetQ is also the only major network-based telemetry product that interfaces Kubernetes to tie together container and networking visibility. In addition, NetQ also timestamps telemetry, which allows time machine capabilities to examine the state of the network at a point in time in the past.





There are a variety of ways NVIDIA Cumulus brings unmatched flexibility, agility, efficiency and savings to the data center network.

For a look at how web-scale networking with NVIDIA Cumulus can lower your networking total cost of ownership, head to our [TCO calculator and white paper](#).

If you'd like to compare configurations and command lines of NVIDIA Cumulus Linux and other vendors, check out [our how-to video series](#).

Learn more about NVIDIA Cumulus's ethernet switching solutions at:
www.cumulusnetworks.com

© 2020 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, and Cumulus are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. All other trademarks and copyrights are the property of their respective owners.