

Pluribus Netvisor Open Networking Edition (ONE) Network Operating System

Comprehensive Enterprise-Class Layer 2/3 Network Operating System for Open Networking Hardware Optimized for Enterprise, Service Provider and Distributed Cloud Data Center Deployments

Highlights

- Linux-based, virtualized, modular network OS built for ONIE- and OCP-compatible networking hardware
- High-performance, scalable, best-in-class, interoperable IPv4 and IPv6 L2/L3 switching and routing
- A rich set of distributed and highly automated overlay services such as distributed policies, distributed VRF routing, anycast gateway, bridge domains with flexible encapsulation (including Q-in-Q) and more
- Deep network slicing to enable customers' different business needs
- Extensive QoS, advanced security, policy and control plane protections
- Highly available and resilient network operating system for non-stop operations
- Open REST API enabling programmability and integration with orchestration platforms
- Extensive integration with Ansible to automate operational workflows
- Single pane of glass to manage a geographically distributed fabric


Moving to bare metal open networking switches brings significant operational and financial benefits to build next-generation networks with the capacity needed to achieve distributed cloud scale, agility, elasticity and adaptability. In the move to a disaggregated open networking strategy, selecting the right network operating system (OS) is a critical success factor as the OS unlocks the functionality, services and performance from open networking switches to realize the most value.

Netvisor® ONE is an open, secure and programmable next-generation network OS that is purpose built to optimize the power and performance of bare metal open networking hardware. Deployment-proven in production mission-critical enterprise and carrier networks, Netvisor ONE meets the most stringent performance requirements and delivers the maximum levels of reliability and flexibility at scale, without compromise.

Built on a Linux foundation, Netvisor ONE is a best-in-class, full-featured network OS that delivers an extensive range of network services, and Layer 2 and 3 switching and routing protocols supporting both IPv4 and IPv6 implementations. Netvisor is built using open networking standards to assure interoperability and operational consistency with existing traditional networking platforms, enabling seamless transition to a disaggregated open networking strategy.

The programmable, modular and containerized architecture of Netvisor ONE enables the rapid adoption and deployment of new networking features and advanced networking services as standards evolve and new capabilities become available.

Exposed APIs support automation and provide unrestricted access to the data, control and management planes for true agility to deliver exceptional operational flexibility. Netvisor ONE brings exceptional operational flexibility with extensibility to scale from deployment in a single device, to advanced services such as VXLAN, to enabling the distributed Adaptive Cloud Fabric™ architecture that federates many devices into a holistic peer-to-peer distributed network to meet diverse deployment requirements.



Netvisor OS provides a modular, resilient software design with advanced Layer 2 and 3 networking and services

As enterprises and regional cloud and colocation providers build their next generation of data centers spanning multiple locations to meet the application's requirements and the placement of workloads at the lowest cost, we will see an increase in the number of mini and micro data center locations, possibly by a factor of ten or more. With this comes increased complexity and a need for a counterbalance – a highly automated network fabric that can make multiple edge locations appear as one logical unit in order to simplify the management of multiple remote data center sites.

Built to Optimize Open Networking Hardware

Netvisor ONE runs on many Open Compute Project (OCP) and Open Network Install Environment (ONIE) hardware-compliant switches, including devices from Dell Technologies, Edgecore and Celestica, as well as the Pluribus Freedom™ Series network switches. This flexibility allows organizations the choice of open networking hardware to build scale-out networks with 10, 25, 40 or 100 Gigabit Ethernet interfaces. This allows an entire data center to be built with only a few physical switch models to improve operational consistency, lower costs and simplifying sparing strategies.

Deployment Flexibility

Netvisor ONE can be deployed as a single OS software image to support any mix of multiple vendor open networking switches in all deployment topologies, including the data center leaf and spine, metro rings, campus building aggregation and service provider. This allows building a network with multi-vendor hardware to flexibly support evolving physical interface requirements unified by a common and consistent OS to reduce operational complexity, improve efficiency and lower TCO.

Standards-Based Networking for Multi-Vendor Interoperability

The Netvisor ONE OS leverages the power of Broadcom switching silicon to deliver the same rich and reliable standards-based networking capabilities expected from traditional networking platforms. Consequently, Netvisor ONE-powered open networking switches can be seamlessly inserted into any existing network, enabling full interoperability with standards-based networking equipment, protocols or network topology. This enables graceful migration to next-generation, disaggregated open networking architectures while preserving existing technology investments and operational models. Integration of open source FRRouting brings in a rich IP routing protocol suite (OSPF, OSPFv3, BGP, MP-BGP, static routes etc.).

High Availability and Resiliency

To meet the stringent continuous availability requirements for mission-critical enterprise and service provider operations, Netvisor ONE provides a comprehensive array of standards-based redundancy protocols and services.

The fully distributed control plane architecture enables a highly available Layer 2/3 underlay and overlay infrastructure. Key high-availability capabilities include switch clustering, multi-chassis LAG (vLAG) VRRP, BFD, ECMP, redundant VTEPs and distributed anycast gateway. Netvisor ONE is built for ultra-fast convergence and sub-second failover to enable reliable deployment as either an underlay or overlay network. When deployed in the Adaptive Cloud Fabric distributed architecture, the unique peering architecture enables highly available, resilient operations across the fabric regardless of the physical location of devices.

Control Plane Traffic Protection

Netvisor ONE delivers exceptional high availability with its Control Plane Traffic Protection (CPTP) architecture. CPTP protects the CPU from excessive traffic volumes and provides fine-grained control and QoS over different types of control plane classes using 44 independent queues. The auto-quarantine host-hog prevention mechanism identifies and automatically quarantines offending host traffic in hardware. The offending host activity is monitored and traffic is resumed automatically when the offending flow ceases.

Manageability, Programmability and Automation

The programmability and automation of Netvisor ONE deliver rapid service provisioning through open programming interfaces enabling both NetOps and DevOps automation to accelerate provisioning, configuration changes and new service roll-outs. Netvisor ONE supports plug-and-play operations and zero-touch provisioning and configuration. Policy and services can be provisioned with a single command via RESTful APIs or command line interface (CLI) with complete functional parity, enabling IT organizations to overcome IT skill gaps and bridge the operational models between DevOps and NetOps.

Automation tools, such as Ansible or the Pluribus UNUM™ management platform, can be leveraged to streamline the provisioning of an entire deployment. In addition, Netvisor ONE supports a wide array of Linux tools for scripting and automation, as well as traditional NetOps interfaces for SSH, TLS 1.2, SNMP, Syslog and sFlow. As a result, workflow automation reduces configuration time by up to 90% over traditional box-by-box management, lowers the risk of configuration errors and dramatically improves service velocity and operational agility.

Secure, Role-Based Management Access

The Netvisor OS has extensive security mechanisms to protect access to OS commands through authentication, authorization, and accounting (AAA) access controls. Administrative user authentication is supported through standards-based mechanisms including TACACS+, Secure Shell (SSH) Version 2 and TLS 1.2. Granular permissions can be defined on a per-user, per-role and per-tenant basis, limiting command-level access for all commands performed for all configuration levels.


Advanced Multi-tenancy Design

Netvisor ONE is fully virtualized to decouple network resources from the underlying hardware. Similar to how a hypervisor virtualizes a bare metal server, this enables segmenting the operating system components and hardware resources into virtual network containers. Each virtual network container has its own dedicated network resources, including dedicated routing data and control planes and an independent management environment. This enables a much deeper degree of isolation among tenants and segmented services.

The unique network virtualization of Netvisor ONE allows any switch of the fabric to instantiate multiple virtual networks that can be dynamically allocated to a single device, or span across multiple physical devices enabling granular network segmentation and multi-tenant services at wire speed.

Adaptive Cloud Fabric Architecture

The Adaptive Cloud Fabric is a simple and secure next-generation software-defined peer-to-peer distributed network architecture that clusters all Netvisor ONE-powered switches into a symmetrical, unified operating domain. The Adaptive Cloud Fabric operates without a controller, so it can seamlessly interoperate with existing networks, enabling a non-disruptive, graceful migration to a software-defined data center or enterprise architecture.



The Adaptive Cloud Fabric architecture clusters member switches into a symmetrical, unified operating domain

To enable distributed deployments, the Adaptive Cloud Fabric features an innovative distributed control plane that allows multiple Netvisor ONE-powered switches to be operated and managed as a single, distributed virtualized switch. Each physical switch maintains its own individual control and data plane to support scale, high performance and resiliency.

To meet stringent high availability requirements, the Adaptive Cloud Fabric architecture has no single point of failure and delivers fabric-wide sub-second failover and resiliency. Fabric automation provides a single point of management and control, distributes intelligence, integrates a broad range of advanced overlay network services and provides pervasive visibility for all traffic traversing the fabric.

The Adaptive Cloud Fabric can be deployed across a single data center, across a campus or geographically distributed to seamlessly interconnect dozens of data centers or aggregate the campus edge over any existing Layer 2 or Layer 3 core. The fabric VXLAN-based overlay can scale out to support many thousands of ports, with multi-terabit capacity, performance and latency predictability, and support millions of concurrent connections.


The Adaptive Cloud Fabric overlay is highly automated and eliminates 90+% of the configuration complexity associated with managing an overlay fabric across N switches. The result is a simple and dynamic architecture that brings a significant operational advantage over controller-based SDN architectures.

Advanced Fabric Services

In addition to enabling the distributed fabric architecture, the Adaptive Cloud Fabric license also enables high-value advanced network services, including:

- Advanced overlay services based on VXLAN
- Network overlay virtualization to enable traffic segmentation and multi-tenant services
- Network packet broker services, which can be provisioned as a stand-alone packet broker fabric or as a virtualized overlay service integrated in a production network fabric (refer to the Pluribus Network Packet Broker data sheet for details)
- Integrated network performance monitoring telemetry
- Flowtrace for troubleshooting packets across the fabric
- Integrations with RedHat OpenStack for automated provisioning of network resources

The Adaptive Cloud Fabric capabilities and advanced services of Netvisor ONE are optionally licensed at time of initial deployment, or may be unlocked as an in-place, non-disruptive upgrade for existing devices through a simple license key.



Integrated telemetry

Warranty and Support

Pluribus Networks offers a wide range of advanced services spanning the entire network lifecycle to protect investments and help accelerate success when deploying and optimizing the Netvisor operating system and next-generation network architectures. Multiple extended support options are available, including on-demand global support, on-site support, advanced hardware replacements and professional implementation services. Maintenance options include direct access to a team of expert network engineers with deep networking experience and our self-service online Customer Portal. For more information about Pluribus support options, visit <http://www.pluribusnetworks.com/support> or contact a Pluribus Networks authorized reseller.

	License Options				
	Enterprise License	Fabric License	vNET License	VirtualWire Single-switch License	VirtualWire+ License
Layer 2 Networking	●	●			●
Layer 3 Networking	●	●			●
Adaptive Cloud Fabric Distributed Architecture	● ¹	●			●
VXLAN services	● ¹	●			●
vNET Segmentation and Tenant Services			●		
Monitoring Telemetry	● ¹	●			●
VirtualWire™ and Packet Broker Services				●	●

¹ Enterprise License supports Adaptive Cloud Fabric features for fabrics up to 6 switches. For larger fabrics, the Fabric License is required.

Features and Specifications

Layer 2

- 802.3z Gigabit Ethernet
- 802.3ab 1000BASE-T
- 802.3ae 10 Gigabit Ethernet
- 802.3ba 40 Gigabit Ethernet
- 802.3ba 100 Gigabit Ethernet
- 802.1D Spanning Tree
- 802.1w Per VLAN Rapid Spanning Tree Protocol (RPVSTP) and RSTP PortFast
- 802.1s Multiple Spanning Tree Protocol (MSTP)
- 802.3ad Link Aggregation (LACP)
- Link Aggregation Group (LAG)
- Multi-Chassis LAG (vLAG)
- STP Cluster Awareness
- Port Fast, BPDU Guard, BPDU Filter, Root Guard
- 802.1q VLANs, VLAN Trunks
- 802.1ab Link Layer Discovery Protocol (LLDP)
- Storm control for Multicast and Broadcast
- 802.1/Qbb — Priority-based Flow Control
- IGMP v2/v3 snooping
- MLD snooping v1/v2
- Jumbo frames (9216 Bytes)
- Private VLAN Edge (for Cisco Interop)
- Fabric ARP and IPv6 Neighbor Discovery Optimization
- Fabric Guard

Layer 3

- Routing Protocols: OSPF, OSPFv3, BGP, MP-BGP
- Dynamic Routing for Virtual Routing and Forwarding (VRF)
- Static routes
- Dual-stack IPv4/IPv6
- VRRP for IPv4 and IPv6 with active-active forwarding
- Equal-cost multi-path routing (ECMP)
- Bidirectional Forwarding Detection (BFD)
- BGP unnumbered
- BGP Graceful shutdown
- Policy Based Routing (PBR)
- DHCP relay
- Route Maps

Fabric Overlay Network Virtualization

- VXLAN VTEP with high availability
- Automated tunnel management
- Automated VXLAN management
- L2 overlay segmentation services with VLANs and Bridge Domains
- Network Packet Broker
- E-Line, E-LAN, E-Tree services
- VLAN re-use
- VLAN aggregation

- Q-in-Q and .1Q
- No-MAC-learning mode
- L3 overlay segmentation with distributed VRF with anycast gateway (IPv4, IPv6, Unicast, Multicast)
- Virtual Service Group (vSG) for VRF leaking and service insertion
- Single-pass IoT (Broadcom Trident 3)
- Virtual Private Cloud (VPC) with per-tenant management segmentation (requires vNET license)
- BGP EVPN Gateway on Border Leaf (Broadcom Trident 3 switches)
- BGP EVPN Gateway with symmetric IRB

Security

- IPv4 Ingress/Egress ACL (vFlow)
- IPv6 Ingress/Egress ACL (vFlow)
- ACL statistics (vFlow stats)
- Control Plane Traffic Protection
- AAA authorization and accounting of all commands
- Full AAA switch control (shell, vtysh, CLI)
- Password protected management access, with role-based controls
- TACACS+ AAA
- 802.1X support on management interface
- BGPv4 MD5
- MAC Security
- SNMPv3 SHA (Authentication)
- Management ACL

Features and Specifications (continued)

QoS and Policy

- 802.1p Class of Service (CoS)
- Differentiated Services Code Point (DSCP)
- DSCP to CoS mapping
- Strict priority queueing
- QoS interface trust (CoS/DSCP)
- Egress per port rate limiting
- Per-port, per-CoS minimum egress bandwidth guarantee
- Per-port, per-CoS maximum egress bandwidth limit
- Weighted Round Robin (WRR) Scheduling
- ACL (vFlow) policing/rate limiting
- Priority-based Flow Control (PFC)

Monitoring and Visibility

- Port Mirroring
- RFC 3176 sFlow
- Traceroute
- Monitor logged-in users
- Fabric-wide embedded network and application traffic telemetry
- nvFlow for TCP connection visibility
- Flowtrace (packet tracing and capture)

VirtualWire Feature Set

(requires VirtualWire or VirtualWire+ license)

L1 VirtualWire

- One-to-One
- One-to-Many replication
- Many-to-Many aggregation and replication
- Many-to-One aggregation and replication
- Network Packet Broker Fabric filtering services
- Layer 2/3/4 traffic filtering
- VLAN tagging
- MAC re-write
- Jumbo frames pass-through
- Error pass-through (CRC) on Trident 3 switches
- Runt pass-through on Trident 3 switches

Management, Automation and Extensibility

- CLI
- RESTful API (with CLI parity)
- Pluribus UNUM for one-touch management and automation (see Pluribus UNUM datasheet)

- Ansible automation
- Zero-Touch Replacement for nodes in an HA cluster
- Configuration roll-back and roll-forward
- Syslog
- SNMP v1, v2, v3
- IPv6 SNMP
- SHA for SNMP authentication
- SSHv2
- TLS 1.2
- IPv6 for management
- Dual stack IPv4/IPv6 management
- Configuration export via SCP
- Fabric-wide single point of management eve
- Fabric-wide configuration simplification (single, atomic transactions for the entire fabric)
- Transaction Rollback and rollforward
- Fabric software upgrade

Integration with RedHat OpenStack

- ML2-plugrin for end-to-end automated overlay provisioning across the fabric

Supported RFCs

- RFC 7C8:D26 User Datagram Protocol (UDP)
- RFC 791 IP
- RFC 792 Internet Control Message Protocol (ICMP)
- RFC 793 TCP
- RFC 826 ARP
- RFC 854 Telnet
- RFC 959 FTP
- RFC 1305 Network Time Protocol (NTP) Version 3
- RFC 1519 Classless Interdomain Routing (CIDR)
- RFC 1591 Domain Name System (DNS) Client
- RFC 1724 RIPv2 MIB Extension
- RFC 1812 IPv4 Routers
- RFC 2236 Internet Group Management Protocol
- RFC 2328 OSPF Version 2
- RFC 2453 RIP Version 2
- RFC 7432 BGP MPLS-Based Ethernet VPN
- RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
- RFC 2519 A Framework for Inter-Domain Route Aggregation
- RFC 3101 OSPF Not-So-Stubby-Area (NSSA) Option
- RFC 3137 OSPF Stub Router Advertisement
- RFC 3376 Internet Group Management Protocol

- RFC 3509 Alternative Implementations of OSPF Area Border Routers
- RFC 4271 BGPv4
- RFC 4443 Internet Control Message Protocol (ICMPv6) for IPv6 specification
- RFC 4456 BGP Route Reflection
- RFC 4486 Subcodes for BGP Cease Notification Message
- RFC 4861 Neighbor Discovery for IP Version 6 (IPv6)
- RFC 4893 BGP Support for Four-Octet AS Number Space
- RFC 5549 Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop

SNMP MIBs

- RFC 3635 EtherLike-MIB
- RFC 3418 SNMPv2-MIB
- RFC 2863 IF-MIB
- RFC 2096 IP-FORWARD-MIB
- RFC 4363 Q-BRIDGE-MIB
- RFC 4188 BRIDGE-MIB
- RFC 4273 BGP4-MIB
- RFC 4750 OSPF-MIB
- RFC 2787 VRRPV2MIB
- RFC 4293 IP MIB
- STP MIB

Ordering Information

Software only, requires compatible switch hardware. License does not include maintenance, order desired maintenance separately.

Netvisor ONE Enterprise Edition Perpetual License (licensed per switch device)

- ONVL-10G-ENT-LIC — Pluribus Open Netvisor Linux Enterprise Edition for 10G switch
- ONVL-25G-ENT-LIC — Pluribus Open Netvisor Linux Enterprise Edition for 25G switch
- ONVL-40G-ENT-LIC — Pluribus Open Netvisor Linux Enterprise Edition for 40G switch
- ONVL-100G-ENT-LIC — Pluribus Open Netvisor Linux Enterprise Edition for 32x100G switch
- ONVL-100GXL-ENT-LIC — Pluribus Open Netvisor Linux Enterprise Edition for 64x100G switch

Netvisor ONE Fabric Edition Perpetual License (licensed per switch device)

- ONVL-10G-PLEX-LIC — Pluribus Open Netvisor Linux Fabric Edition for 10G switch
- ONVL-25G-PLEX-LIC — Pluribus Open Netvisor Linux Fabric Edition for 25G switch
- ONVL-40G-PLEX-LIC — Pluribus Open Netvisor Linux Fabric Edition for 40G switch
- ONVL-100G-PLEX-LIC — Pluribus Open Netvisor Linux Fabric Edition for 32x100G switch
- ONVL-100GXL-PLEX-LIC — Pluribus Open Netvisor Linux Fabric Edition for 64x100G switch

Netvisor ONE vNET Segmentation License Add-on (licensed per switch device)

- VNV-4-VNET — vNET license for independent per-tenant management of dedicated fabric resources

Netvisor ONE VirtualWire Single-switch License (licensed per switch device)

VirtualWire feature set enables network packet broker and Layer 1 lab automation capabilities. VirtualWire license is for a single switch (not part of a fabric)..

- ONVL-10G-VW-LIC –VirtualWire service license for 10G switch
- ONVL-25G-VW-LIC –VirtualWire service license for 25G switch
- ONVL-40G-VW-LIC –VirtualWire service license for 40G switch
- ONVL-100G-VW-LIC –VirtualWire service license for 32x100G switch
- ONVL-100GXL-VW-LIC –VirtualWire service license for 64x100G switch

Netvisor ONE VirtualWire+ License (licensed per switch device)

VirtualWire+ license includes fabric, network packet broker and Layer 1 lab automation capabilities.

- ONVL-10G-VW+-LIC –VirtualWire+ service license for 10G switch
- ONVL-25G-VW+-LIC –VirtualWire+ service license for 25G switch
- ONVL-40G-VW+-LIC –VirtualWire+ service license for 40G switch
- ONVL-100G-VW+-LIC –VirtualWire+ service license for 32x100G switch
- ONVL-100GXL-VW+-LIC –VirtualWire+ service license for 64x100G switch

Netvisor ONE OS Scalability



Supported Hardware Devices					
Pluribus Freedom	9373-X 9372-T		9572L-V 9532L-V 9532-C	9432-C 9476-V 9460-X	F9664-C
Dell EMC PowerSwitch	S4048-ON S4048T-ON S6010-ON	S4148F-ON S4148T-ON S4128F-ON S4128T-ON S4112FF-ON S4112T-ON	Z9100 S5048-ON	S5232F-ON S5248F-ON S5224F-ON	Z9264F-ON
Edge-Core	AS-5812-54X AS-5812-54T		AS-7312-54XS AS-7712-32X Wedge 100-32X	AS-7726-32X AS-7326-56X AS-5835-54X	AS7816-64X
Celestica				Edgestone™ DX032 Seastone2E™ DX031	
MAC Addresses*	200,000	207,000	102,000	229,000	200,000
IPv4 LPM Routes*	16,000	16,000	16,000	32,000	32,500
IPv4 host table size*	16,000	32,767	32,767	38,500	48,500
Rapid per-VLAN Spanning Tree (RPVST) instances	256	256	256	256	510
VLANs (with MST)*	4,096	4,096	4,096	4,096	4,096
VLAN with Bridge Domain	8K per cluster (switch HA pair)	8K per cluster (switch HA pair)	8K per cluster (switch HA pair)	8K per cluster (switch HA pair)	8K per cluster (switch HA pair)
Bridge Domain	4000	4000	4000	4000	4000
MST Instances	20	20	20	20	20
Link Aggregation Group (LAG)	40	40	40	40	40
Number of links per LAG	16	16	16	16	16
Multi-chassis LAG (vLAG)	32	32	32	32	32
Member ports per vLAG	16	16	16	16	16
IGMP Groups	4,000	8,000	4,000	16,000	8,000
MLD Groups	4,000	8,000	4,000	12,000	4,700
Mirroring Sessions	2 bidir/4 unidir	2 bidir/4 unidir	4 unidir	2 bidir/4 unidir	2 bidir/4 unidir
Jumbo Frames	9416B	9416B	9416B	9416B	9416B
Ingress ACL (vFlow)	8,000	4,096	1,000	3,072	1,024
L3 Interfaces	64	64	64	64	64
Distributed VRFs	1,000	1,000	2,000	4,000	1,000
VRRP Groups*	256	256	256	256	256
BGP Neighbors	64	64	64	64	64
OSPF Neighbors	64	64	64	64	64
BFD Sessions	64	64	64	64	256
ECMP Paths	16	16	16	16	16
VXLAN Tunnels	1,000	1,000	1,000	1,000	1,000
QoS Queues	8	8	8	8	8
Telemetry: TCP Connection Track Rate	5,000/sec	5,000/sec	5,000/sec	5,000/sec	5,000/sec
Number of vRouters	2	2	2	2/32***	2

Notes:

- Functionality and scalability varies based upon underlying open networking hardware capabilities

* = Limited by hardware

** = Check Netvisor Release Notes for availability

*** = Only for Celestica switches