ACOS 6.0.9
Configuring Overlay Networks

July, 2026

Information in this document is subject to change without notice.

PATENT PROTECTION

A10 Networks, Inc. products are protected by patents in the U.S. and elsewhere. The following website is provided to satisfy the virtual patent marking provisions of various jurisdictions including the virtual patent marking provisions of the America Invents Act. A10 Networks, Inc. products, including all Thunder Series products, are protected by one or more of U.S. patents and patents pending listed at: a10-virtual-patent-marking.

TRADEMARKS

A10 Networks, Inc. trademarks are listed at: a10-trademarks

CONFIDENTIALITY

This document contains confidential materials proprietary to A10 Networks, Inc. This document and information and ideas herein may not be disclosed, copied, reproduced or distributed to anyone outside A10 Networks, Inc. without prior written consent of A10 Networks, Inc.

DISCLAIMER

This document does not create any express or implied warranty about A10 Networks, Inc. or about its products or services, including but not limited to fitness for a particular use and non-infringement. A10 Networks, Inc. has made reasonable efforts to verify that the information contained herein is accurate, but A10 Networks, Inc. assumes no responsibility for its use. All information is provided "as-is." The product specifications and features described in this publication are based on the latest information available; however, specifications are subject to change without notice, and certain features may not be available upon initial product release. Contact A10 Networks, Inc. for current information regarding its products or services. A10 Networks, Inc. products and services are subject to A10 Networks, Inc. standard terms and conditions.

ENVIRONMENTAL CONSIDERATIONS

Some electronic components may possibly contain dangerous substances. For information on specific component types, please contact the manufacturer of that component. Always consult local authorities for regulations regarding proper disposal of electronic components in your area.

FURTHER INFORMATION

For additional information about A10 products, terms and conditions of delivery, and pricing, contact your nearest A10 Networks, Inc. location, which can be found by visiting www.a10networks.com.

Table of Contents

Introduction

An overlay network is a virtual network built on top of an existing physical network in order to introduce a service that is not available in that existing network. The overlay network uses virtual links to connect various nodes of the existing physical network. Virtual network traffic then tunneled across the network using technologies like VXLAN (Virtual Extensible LAN) and NVGRE (Network Virtualization using Generic Routing Encapsulation).

Both VXLAN and NVGRE serve to connect multiple Layer 3 networks and give the appearance that they all share the same Layer 2 domain.

Additional Implementation Notes

The following notes clarify the current implementation of SDN.

- This feature is supported on all A10 Thunder Series platforms and hypervisors (KVM, ESXi, Hyper-v and OCI, and OpenStack) except for the following:
  - A10 Thunder Series 5630
  - vThunder platforms
    - AWS
    - Azure
- Only IPv4 networks, hosts and tunnels are supported, both for the overlay and the underlay. All non-IPv4 traffic going across the overlay will need the corresponding logical interface to be configured as a Layer 2 interface.
- Any Layer 3 or Layer 4 features configured for the overlay traffic needs to be explicitly configured on the logical interface. For example if access-list is required to police incoming traffic, it will need to be bound to the logical interface. Configuring it and binding it on the underlay interface will not apply to the overlay traffic as the incoming traffic is de-encapsulated first and then the policy lookup is performed.
- For VRRP-A, the VTEP IP should be configured as a floating IP. With VRRP-A configured, all VTEP and logical interface partitions need to be part of the same VRRP group. The VTEP IP should be configured as the vrid-lead, while the vrid resources in the logical interface partition, will need to be configured to follow the VTEP IP.
- With VRRP-A configured, the provider partition that has the VTEP configured and all related logical interface (customer) partitions must be configured so that all of them are active or standby on the same ACOS device. VTEP-to-logical interface communication between different ACOS devices is not supported.
- Forward reference for the logical interface partition as well as the logical interface is supported for the overlay-tunnel configuration. For example, the overlay-tunnel configuration can be added without the need to first configure the overlay partition (if different from the underlay partition) or creating the logical interface.

- Bridge VLANs with logical interfaces configured as Layer 2 interfaces are not supported.
- If the real servers are across an overlay tunnel, then the health check packets from both Active and Standby ACOS Thunder devices will go over the overlay tunnel. Responses from the servers will only be received by the Active ACOS device with the active Floating VTEP IP, so the real servers on the Standby will be in “down” state until switchover happens.
- Given that both Active and Standby ACOS devices will be sending health checks to the servers, if these were to go over the tunnel, the responses are directed only to the Active ACOS device on the VTEP IP address. On the Standby ACOS device, the Servers will be in “down” state until the switchover happens and health checks start passing.

Understanding Packet Encapsulation and Transport

The following topics are covered:

Understanding VXLAN Encapsulation

In a VXLAN network, Layer 2 packets are tunneled over a Layer 3 network using UDP/IP over IP frames, as shown in Figure 1.

Figure 1 : VXLAN Packet Encapsulation

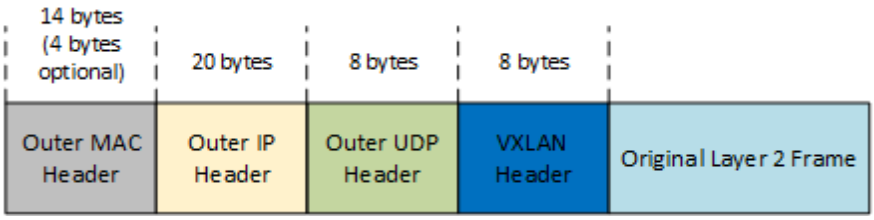

Table 1 describes the various fields in the VXLAN encapsulation.

Table 1 : VXLAN Packet Encapsulation

Field	Description
Outer MAC Header	Contains the following: • Destination MAC address (48 bits) • Source MAC address (48 bits) • VLAN type set to 0x8100 (16 bits) • VLAN ID (16 bits)

Table 1 : VXLAN Packet Encapsulation

Field	Description
	• Ethertype of 0x0800 to designate the payload as an IPv4 packet (16 bits)
Outer IP Header	Contains the following • IP header miscellaneous data (72 bits) • Protocol set to 0x11 to indicate UDP (8 bits) • Header chksum (16 bits) • Source IP address of the originating VTEP (32 bits) • Destination IP address of the target VTEP (32 bits)
Outer UDP Header	Contains the following: • Source port (16 bits) • VXLAN port (16 bits) • UDP length (16 bits) • UDP checksum (16 bits)
VXLAN Header	Contains the following • VXLAN flags Reserved bits set to zero except for the third bit, which is set to 1 to indicate a valid VNI. (8 bits) • VNI (24 bits) Two sets of fields, 24 bits and 8 bits, that are reserved and set to zero.

Understanding NVGRE Encapsulation

NVGRE packets are tunneled directly over IP (as opposed to UDP in VXLAN). The outer IP header protocol field is set to 47 (the protocol number for GRE).

Figure 2 : NVGRE Packet Encapsulation

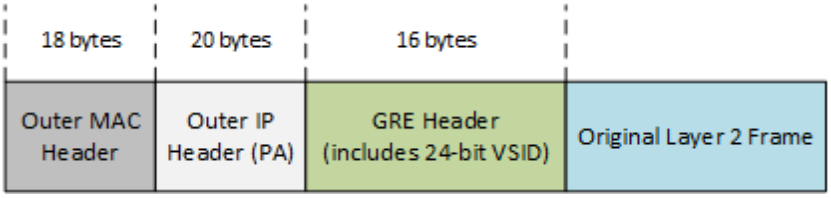

Table 2 describes the various fields in the NVGRE encapsulation.

Table 2 : NVGRE Packet Encapsulation

Field	Description
Outer MAC Header	The source and destination MAC addresses.
Out IP Header	The provider addresses (PA); the source and destination IP addresses for the tunnel.
GRE Header	Contains the following: • Flag field (8 bits) The third bit (the “key” flag) is set to 1, meaning the header contains a valid VSID. The remaining bits are set to zero. • Protocol Type This field is set to 0x6558 which is reserved for Transparent Ethernet Bridging. • 24-bit VSID: These 24 bits in the 32-bit key field in the GRE header. designate an individual NVGRE network on which virtual machines can communicate. Virtual machines in different NVGRE networks cannot communicate with each other. • Flow ID: These last 8 bits of key field are optional and can be used to add per-flow entropy within the same VSID so that entire key field (32 bits) can be used for Equal-Cost Multi-Path (ECMP) routing purposes by switches and routers. If Flow ID is not generated, it must be set to

Table 2 : NVGRE Packet Encapsulation

Field	Description
	zero.

Understanding GRE Encapsulation

ACOS GRE tunneling supports IPv4 GRE tunneling with IPv4 and IPv6 payload. The GRE KeepAlive feature is also introduced to continuously monitor GRE tunnel endpoints' health and bring down the corresponding logical interface (LIF) upon failure. The Encapsulation and Decapsulation support of the GRE packets is configured with the `encap GRE` CLI command.

The following features are provided:

- Send Syslog and SNMP trap alerts when the GRE tunnel is up or down
- View the GRE tunnel status using CLI and aXAPI
- Distribute routes based on GRE tunnel status

Known Limitations

The following functionalities are not supported in this release:

- Fragmentation for ISIS packets
- IPv6 GRE tunnels
- GRE supported for non-FTA models only

Understanding Fragmentation of Tunneled Packets

As a result of the additional header (50 bytes for VXLAN and 42 Bytes for NVGRE) added to tunneled packets, fragmentation of the packet can occur if the host is not aware of this condition and the maximum segment size (MSS) is exceeded. The ACOS device will automatically adjust the MSS exchanged for TCP flows to take into account this header.

To disable this automatic adjustment, use the **`overlay-tunnel options tcp-mss-adjust-disable`** command. For more information, refer to the **`overlay-tunnel`** command in the Command Line Interface Reference.

Packet fragmentation for tunneled packets is not supported on FTA platforms. If you are configuring overlay tunnels on an FTA device, make sure to enable this automatic adjustment to avoid fragmented packets being dropped. Use the following command: `no overlay-tunnel options tcp-mss-adjust-disable` .This command only works for TCP traffic. For non-TCP traffic, you must manually adjust the MTU on your device to avoid fragmentation.

Understanding IP Encap (IPinIP) Tunneling

The IP encap supports IPinIP, IPinIPv6, IPv6inIP and IPv6inIPv6 encapsulations. Payloads of IPv4 and IPv6 can be transported atop IPv4 or IPv6 encapsulation headers. Encapsulation and Decapsulation support of the IPv4 and IPv6 encapsulation is configured with the `ip-encap` VTEP encap.

Figure 3 : IP Encap (IPinIP) Tunneling

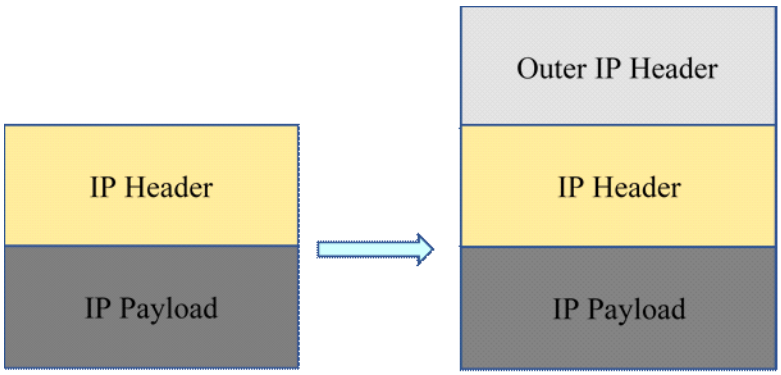

Configuring a Software Defined Network

This section describes how to configure the overlay tunnels to create a Software Defined Network (SDN) or Overlay Network.

Figure 4 shows a sample overlay tunnel topology that will be used as a sample configuration in this section.

Figure 4 : Sample Overlay Tunnel Topology

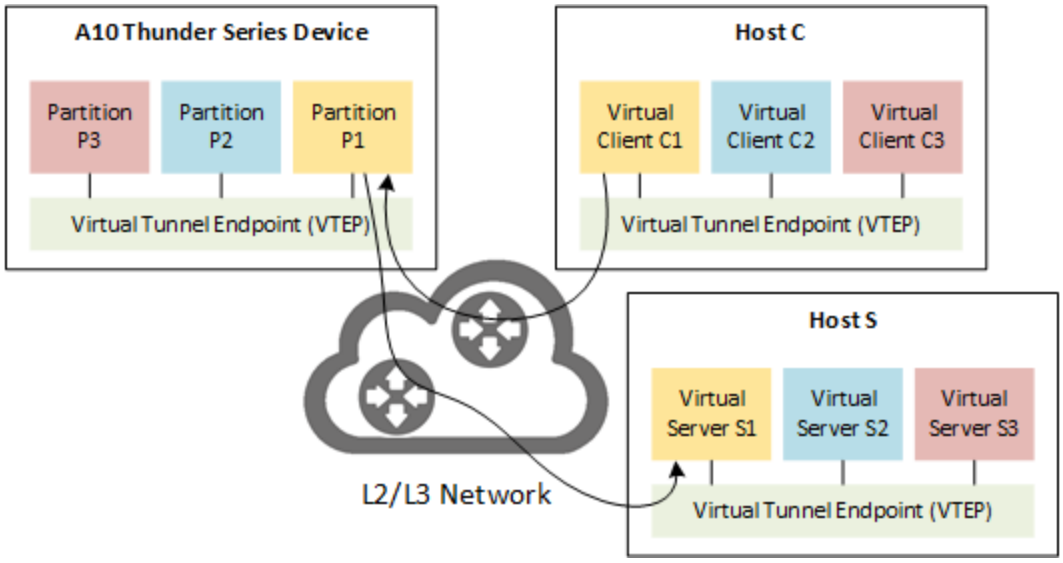

NOTE: The instructions assume that you have already created the necessary L3V partitions on your system. The overlay partition should be an L3V partition.

The following topics are covered:

Configuring the Underlay/Provider Partition

This section provides the commands used to configure the provider partition. More details about each command are available in the Command Line Interface Reference.

This example is configured using VXLAN encapsulation; replace **`encap vxlan`** with **`encap nvgre`** if you want to use NVGRE encapsulation.

The following topics are covered:

Beginning the Tunnel Configuration and Specifying the Encapsulation

Enter configuration mode on your ACOS device and enter the following commands.

```
ACOS# config
ACOS(config)# overlay-tunnel vtep 1
ACOS(config-overlay-tunnel:1)# encap vxlan
```

Configuring the Source IP Address

These commands create the binding between the user partitions (P1 and P2) and the tunnel endpoint VTEPa in the provider partition.

NOTE: The source or local IP address used in the configuration (in this case, 120.1.1.10) needs to be present in the underlay partition as an interface IP or a floating IP before any traffic will pass through.

```
ACOS(config-overlay-tunnel:1)# local-ip-address 120.1.1.10
ACOS(config-overlay-tunnel:1-src-vtep)# vni 100 partition p1 lif 10
ACOS(config-overlay-tunnel:1-src-vtep)# vni 200 partition p2 lif 20
ACOS(config-overlay-tunnel:1-src-vtep)# exit
ACOS(config-overlay-tunnel:1)#
```

Creating the Remote VTEP Configuration for the Clients (VTEPc)

These commands create the mapping for packets between VTEPa to VTEPc, identifying the destination or remote IP address for VTEPc and the fact that there are two VNIs behind VTEPc:

```
ACOS(config-overlay-tunnel:1)# remote-ip-address 100.1.1.10
ACOS(config-overlay-tunnel:1-dst-vtep)# vni 100
ACOS(config-overlay-tunnel:1-dst-vtep)# vni 200
ACOS(config-overlay-tunnel:1-dst-vtep)# exit
ACOS(config-overlay-tunnel:1)#
```

Creating the Remote VTEP Configuration for the Servers (VTEPs)

These commands create the mapping for packets between VTEPa to VTEPs, identifying the destination or remote IP address for VTEPs and the fact that there are two VNIs behind VTEPs:

```
ACOS(config-overlay-tunnel:1)# remote-ip-address 110.1.1.10
ACOS(config-overlay-tunnel:1-dst-vtep)# vni 100
ACOS(config-overlay-tunnel:1-dst-vtep)# vni 200
ACOS(config-overlay-tunnel:1-dst-vtep)# exit
ACOS(config-overlay-tunnel:1)#
```

Configuring the Source Port Range

You can also configure UDP source port range option for VXLAN overlay tunnels with the ability to control the VXLAN source port range globally or per VTEP.

NOTE: This feature is not supported on multi-PU platforms.

Global VXLAN Configuration for SRC Port Range

```
ACOS(config)# overlay-tunnel options src-port-range min-port <1-65535>
max-port <1-65535>
```

Per VTEP Configuration for SRC Port Range

```
ACOS(config)# overlay-tunnel vtep 10
ACOS(config-overlay-tunnel:10)# src-port-range min-port 47000 max-port
48000
```

ACOS selects VXLAN UDP source (SRC) ports based on the following conditions:

- If the source port is not configured globally, then the minimum source port defaults to 1 and the maximum to 65535.
- If you do not configure a VTEP with a SRC port range, the VTEP uses the global SRC port range configured.
- If you change the global VXLAN UDP SRC port range, then existing VTEPs do not inherit the new SRC port range; only new VTEPs configured after the change will inherit the new SRC port range to prevent disruptions to existing VTEP sessions.

NOTE: For all existing VTEPs to use a configured global UDP SRC port range, save the configuration using the **`write memory`** command and reboot the device using the **`reload`** or **`reboot`** commands.

CAUTION: Changing a VTEP UDP source port range is effective immediately, removing existing sessions in the VTEP.

CLI Configuration

This overlay VXLAN tunnel example configures VTEP 20, UDP SRC-port range 2000 to 3000, virtual network ID, VNI 10, and the remote endpoint IP address. Configure the

LIF interface before configuring the VNI.

```
ACOS(config)# overlay-tunnel vtep 20
ACOS(config-overlay-tunnel:20)# encap vxlan
ACOS(config-overlay-tunnel:20)# src-port-range min-port 2000 max-port 3000
ACOS(config-overlay-tunnel:20)# local-ip-address 10.1.1.1
ACOS(config-overlay-tunnel:20-src-vtep:10....)# vni 10 lif 1
ACOS(config-overlay-tunnel:20)# remote-ip-address 203.0.113.10
```

Manually Configuring the Host-to-Destination VTEP Mapping to Avoid Dynamic Learning

By default, the ACOS device will attempt to dynamically learn the host-to-destination VTEP mapping by using unicast IP and ARP response packets coming over the tunnel.

The `host` command is used to manually map the source VTEP to the logical interface behind the destination VTEP. For example, the following command maps VTEPa to Client C1 (IP address 10.1.1.10 and MAC address 001c.011c.111c) on VNI 100 behind VTEPc (IP address 100.1.1.10):

```
ACOS(config-overlay-tunnel:1)# host 10.1.1.10 001c.011c.111c vni 100
remote-vtep 100.1.1.10
```

Additional mappings between VTEPa to Client C2, Server S1, and Server S2 may be similarly configured:

```
ACOS(config-overlay-tunnel:1)# host 20.1.1.10 002c.022c.222c vni 200
remote-vtep 100.1.1.10
ACOS(config-overlay-tunnel:1)# host 10.1.1.20 001s.011s.111s vni 100
remote-vtep 110.1.1.10
ACOS(config-overlay-tunnel:1)# host 20.1.1.20 002s.022s.222s vni 200
remote-vtep 110.1.1.10
```

Configuring the Overlay/Tenant Partition P1

The following topics are covered:

Configuring a Standalone LIF as a Layer 3 Interface

The following commands configure a standalone Layer 3 logical tunnel interface to be used for outbound traffic from partition P1 to VTEPa:

```
ACOS[p1](config)# interface lif 10
ACOS[p1](config-if:lif10)# ip address 10.1.1.1 /24
ACOS[p1](config-if:lif10)# exit
ACOS[p1](config)#
```

Configuring a Pure Layer 2 LIF

Similarly, you can configure an untagged Layer 2 logical interface in a VLAN; an example of this is shown with partition P2:

```
ACOS[p2](config)# vlan 20
ACOS[p2](config-vlan:20)# untagged lif 20
ACOS[p2](config-vlan:20)# tagged ethernet 5
ACOS[p2](config-vlan:20)# exit
ACOS[p2](config)#
```

In such a configuration, when a packet is received on VLAN 20, a copy is tunneled over to the VTEP, and a second copy is sent to ethernet interface 5.

Configuring a LIF as a Layer 2 Interface Inside a VLAN with a Layer 3 Virtual Ethernet Endpoint

Alternatively, you can make the Layer 2 LIF into a Layer 3 interface by adding the IP address to a virtual ethernet interface configured on the VLAN:

```
ACOS[p1](config)# vlan 20
ACOS[p1](config-vlan:20)# untagged lif 20
ACOS[p1](config-vlan:20)# tagged ethernet 5
```

```
ACOS[p1](config-vlan:20)# router-interface ve 20
ACOS[p1](config-vlan:20)# exit
ACOS[p1](config)# interface ve 20
ACOS[p1](config-if:ve20)# ip address 20.1.1.1 /24
ACOS[p1](config-if:ve20)# exit
ACOS[p1](config)#
```

Verifying the Configuration

To verify your configuration, you can use the following **show** commands. Additional information for each command is available in the *Command Line Interface Reference*.

To verify the configuration of the VTEP, use the show running-config overlay-tunnel command. For more information, see , see *Command Line Interface Reference.*

```
ACOS# show running-config overlay-tunnel
!
overlay-tunnel vtep 1
  local-ip-address 120.1.1.10
   vni 100 partition p1 lif 10
  remote-ip-address 100.1.1.10
   vni 100
!
```

Glossary

A

ACL

An Access Control List (ACL) restricts management-plane access by permitting only authorized source IP addresses.

ARP

Address Resolution Protocol (ARP) used by ACOS in scaleout deployments to resolve IP addresses to MAC addresses. It enables proper Layer 2 communication between clustered devices and connected network elements.

ASN

ASN (Autonomous System Number) is a unique identifier used in BGP routing for enabling BGP peering, route advertisement, multi-homing, and integration with ISP-grade routing domains.

B

Bare Metal

The Thunder Bare Metal product line offers high-performance software to manage the high demands of today's application networking and security workloads.

BFD

The Bidirectional Forwarding Detection (BFD) is a lightweight protocol that detects routing path failures between adjacent forwarding nodes.

BGP

Border Gateway Protocol (BGP) supports within scaleout that enables dynamic route advertisement and exchange. It allows ACOS devices in a cluster to participate in routing decisions and maintain reachability

during failover and scale events.

D

DSCP

Differentiated Services Code Point. A field in the IP header used to classify and manage network traffic by assigning different levels of service priority.

E

ECMP

The Equal-Cost Multi-Path (ECMP) refers to routers or switches distributing traffic across multiple ACOS paths or appliances that share the same routing cost, providing higher throughput and redundancy.

F

FTA

The Failed Authentication Attempts (FTA) track unsuccessful login attempts to support lock-out enforcement and security monitoring.

I

ICMP

Internet Control Message Protocol (ICMP) is used for network diagnostics, error reporting, NAT handling, health checks, and supporting path MTU discovery to ensure proper traffic flow and troubleshooting.

IGMP

Internet Group Multicast Protocol (IGMP) ensures efficient handling of multicast traffic by allowing the system to participate correctly in multicast group membership and preventing unnecessary multicast flooding.

L

L3V

Layer 3 Virtualization (L3V) is a virtualization layer that allows organizations to utilize the same IP address ranges for ensuring that the multi-tenant data center architecture gets the flexibility similar to that of a independently-deployed device.

LACP

The Link Aggregation Control Protocol (LACP) dynamically bundles multiple physical network interfaces into a single logical link to increase bandwidth and provide link-level redundancy.

LIF

A Logical Interface (LIF) represents a virtual management interface used for device communication and authentication traffic.

LLDPDU

Link Layer Discovery Protocol Data Unit (LLDP) allows ACOS devices to discover directly-connected LAN neighbors and allows these neighbors to discover the ACOS devices. Configure LLDP only in the shared partition.

M

MSTP

MSTP (Multiple Spanning Tree Protocol) enables loop-free Layer-2 networks by allowing VLANs to be grouped into spanning-tree instances, improving scalability, interoperability, and traffic distribution in bridged deployments

MTU

The Maximum Transmission Unit (MTU) is the maximum size of a packet that can be trans-

mitted over a network interface.

N

NAT

Network Address Translation (NAT) is a method of real-locating one IP address space into another by changing the network address information in the IP header when the packets are still being transmitted across a traffic routing device.

NSM

NSM (Network Service Monitor) is an internal health-monitoring component that supervises ACOS processes and can restart or terminate them if they become unresponsive, ensuring system stability, and high availability.

NVGRE

NVGRE (Network Virtualization using Generic Routing Encapsulation) is an overlay tunneling method that encapsulates Layer-2 traffic inside GRE to create scalable, isolated virtual networks, enabling ACOS to deliver ADC, CGN, and security services within modern SDN environments.

O

OSPF

The Open Shortest Path First (OSPF) is a routing protocol that run over IPsec tunnels.

R

RSTP

Real Time Streaming Protocol (rtsp). A network control protocol designed for use in entertainment and communications systems to control streaming media servers. It is used to establish and control

media sessions between endpoints.

S

SNAT

A Source Network Address Translation (SNAT) process that rewrites outbound source addresses to firewall-assigned values.

T

TWAMP

The Two-Way Active Measurement Protocol (TWAMP), defined in RFC 5357, is an IP QoS network measurement protocol that provides QoS scrutiny of circular-tour performance between two network endpoints.

U

UDLD

UDLD (Unidirectional Link Detection) in ACOS detects and protects against one-way physical link failures—especially in LACP trunks—ensuring link integrity, preventing loops, and improving reliability of Layer-2 connectivity

V

VCS

The Virtual Chassis System (VCS) allows multiple A10 devices operate together as one unified, scalable, highly available platform with a single management point.

VIP

The Virtual IP (VIP) address used by ACOS to provide a single entry point for load-balancing.

VRID

The Virtual Router Identifier(VRID) used by VRRP-A to group redundant firewalls that share virtual IPs.

VRRP

The Virtual Router Redundancy Protocol (VRRP) provides gateway high availability by allowing multiple ACOS devices to share a virtual IP so that traffic continues even the primary router fails.

VTEP

The VXLAN Tunnel Endpoint (VTEP) is the endpoint that performs VXLAN encapsulation or decapsulation, enabling ACOS devices to operate within VXLAN-based overlay networks.

Contact Us