

VXLAN Routing with Arista 7050X2 Series

The Arista Networks 7050X2 series of fixed configuration products expands the broader Arista 7050X product portfolio. The 7050X2 products provide high-density systems, with line-rate VXLAN routing without the need for additional ASICs, recirculation or the reduction in external front panel ports. The 7050X2 products, in combination with the 7050X systems, provide a comprehensive set of VXLAN solutions that scale to the full forwarding capacity of the 7050X system.

Background

The VXLAN protocol is an RFC (7348) standard that was co-authored by Arista Networks and defines a MAC in IP encapsulation protocol allowing the construction of layer 2 domains across a physical layer 3 IP infrastructure. VXLAN decouples the physical network topology from a 'logical' or 'virtual' topology often referred to as the

'Overlay Network.' The benefits to VXLAN are the construction of layer 2 and layer 3 logical topologies over a physical infrastructure that takes advantage of proven, standards-based IP ECMP routing topologies. The logical topologies, identified by virtual network IDs (VNIs), can be used as a way to build multi-tenant environments across a shared infrastructure and/or increase deployment agility by providing layer 2 connectivity between racks, halls or physical data centers.

7050X VXLAN Routing

Virtual Extensible LAN (VXLAN) allows the extension of an L2 segment or subnet over an L3 interconnect network via an overlay or tenant VNIs. VXLAN tunnel end-points (VTEPs) encapsulates packets with a VXLAN header identifying the appropriate tenant VNI, as well as an additional IP header with a local source VTEP IP and a destination address of a remote VTEP (where the subnet has been extended to). Intermediate devices forward the packet based entirely on the IP information contained within this external provider header. VXLAN routing is needed when hosts in different logical topologies (VNIs) need to communicate with one another.

Because VXLAN routing is performed based on the inner tenant IP packet rather than the source/destination IPs of the VXLAN encapsulation, a VXLAN router must be capable of both encapsulating and decapsulating VXLAN headers as well as routing the internal packets. Although a VXLAN Routing capability was not offered as a feature of the silicon, through their extensive understanding of the underlying Packet Processor, Arista Networks have been able to support VXLAN routing on Trident-II chipsets through the 7050X Series.

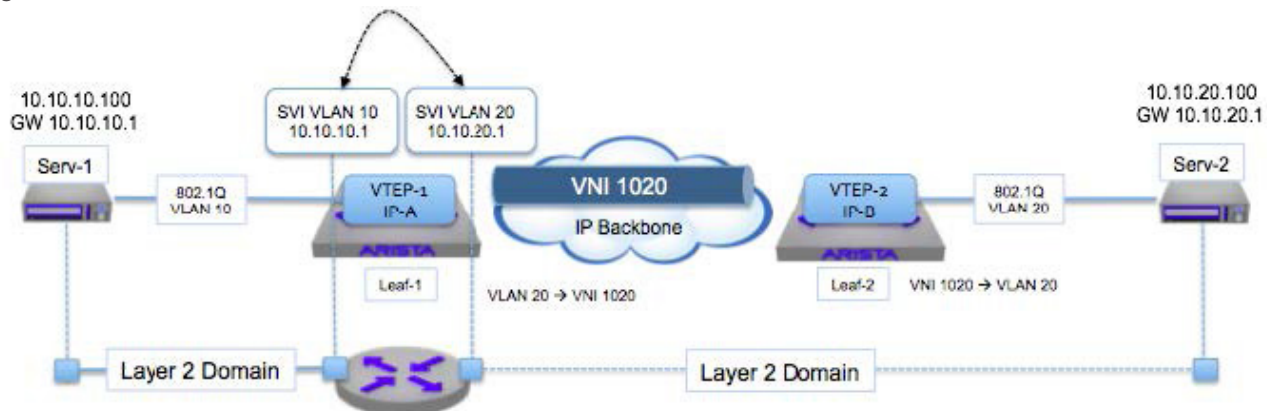


Figure 1: VXLAN Routing in a Network

This allowed Arista to bring VXLAN routing to all 7050X platforms without adding complexity, increasing power, or adding cost associated with adding a secondary packet processor. The Arista implementation enabled line rate VXLAN bridging and routing concurrently on all 7050X platforms using all the available internal chip bandwidth. There is spare internal bandwidth available on all 7050X models with the exception of the three highest density fixed configuration models: 7050SX-128, 7050TX-128 and the 7050QX-32/QX-32S.




Figure 2: 7050X Series VXLAN Route and Encapsulation Model

In order to bring VXLAN routing to the three higher density platforms, external bandwidth was needed to recirculate packets through the packet processor for the encapsulation or decapsulation of the VXLAN headers. High performance, line-rate VXLAN routing on these platforms requires disabling some of the front panel ports for recirculation bandwidth. The table below summarizes the current 7050X VXLAN support and available recirculation bandwidth.

Table 1: Arista 7050X VXLAN Routing

Models	VXLAN Routing	Front panel ports	Available internal bandwidth
7050SX-64 7050TX-64	Supported	48 x 10G + 4 x 40G	16 x 40G
7050SX-72 7050SX-72Q 7050TX-72 7050TX-72Q	Supported	48 x 10G + 6 x 40G	14 x 40G
7050SX-96 7050TX-96	Supported	48 x 10G + 12 x 40G	8 x 40G
7050QX-32 7050QX-32S	Supported	32 x 40G	N/A
7050SX-128 7050TX-128	Supported	96 x 10G + 8 x 40G	N/A

With the evolution of the Trident-II+ chipset the packet processor enables a single-pass, line-rate VXLAN bridging and routing capability. The Arista 7050X2 models of the 7050X Series enable line-rate VXLAN bridging and routing, without the need to disable front panel interfaces in a single-pass.



Figure 1: VXLAN Routing in a Network

For customers that do not require VXLAN routing or require VXLAN routing on lower port count switches, there is adequate internal bandwidth available for recirculation. The 7050X2 models are for customers requiring line rate VXLAN routing in combination with higher port counts.

Table 2: Arista 7050X2 VXLAN Routing

Models	VXLAN Routing	Front panel ports	Available internal bandwidth
7050QX2-32S	Supported	32 x 40G	Not required
7050SX2-72Q	Supported	48 x 10G + 6 x 40G	Not required
7050SX2-128 7050TX2-128	Supported	96 x 10G + 8 x 40G	Not required

Summary

VXLAN routing is an open solution to extend L2 segments and remove the traditional L2 boundaries of the physical infrastructure. The Arista 7050X Series enable VXLAN bridging and routing in combination with support for active-active fully redundant topologies, VXLAN control services, network wide services and VXLAN aware network telemetry. The Arista 7050X2 continue these capabilities and deliver full bandwidth support for VXLAN routing, in a single pass with no requirement to assign either internal or external bandwidth.

For more information on Arista VXLAN capabilities, designs and best practices refer to the VXLAN Design Guide at www.arista.com.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office

1390 Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2016 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. May 2016 02-0066-01