

Multi-Chassis Link Aggregation

Historically, the capabilities of networking technologies have often missed the network designer's performance requirements by as much as an order of magnitude. For many years, networks have been oversubscribed. Such oversubscription was tolerable in enterprise networks as the applications were not bandwidth intensive. However, recent changes in technology, such as rich-media applications and low-cost compute servers, increased the bandwidth requirements of the network and exposed a glaring problem in datacenter networks.

In a datacenter, the primary reason for this oversubscription is insufficient uplink bandwidth from each rack. Each top-of-rack switch is typically connected to two aggregation switches for redundancy. However, half the uplinks are blocked by spanning tree in order to avoid loops in the network; this reduces the available bandwidth between the rack and aggregation layer of the network by 50%.

Arista's Multi-Chassis Link Aggregation (MLAG) feature removes this bottleneck and allows the utilization of all interconnects in an active/active mode.

Legacy Network Designs

The network design in Figure 1 shows legacy architecture. Half of the Layer 2 uplinks are in STP blocked state, resulting in increased oversubscription in the network and also a higher price premium when you look at usable bandwidth per server. While the Multiple Spanning Tree Protocol (MSTP) has been used to manually divide the traffic, such designs still do not offer the flat network that many applications require for all nodes to communicate to each other in a single VLAN.




Figure 1: Legacy Design

Arista MLAG

In order to utilize all interconnects in an active/active manner, Arista EOS now supports the MLAG feature. This allows one to interconnect two Arista switches and use them as one logical switch for the purpose of L2 protocols such as STP or LACP.




Figure 2: Active/Active Network with MLAG

Key Benefits of MLAG:

- No wasted bandwidth with uplinks in Spanning Tree Blocking state
- Allows you to design non-blocking networks
- Maintains same level of resiliency with redundant paths available at all times
- Two Arista switches can be in an MLAG pair, increasing your network scalability by 2X
- No proprietary protocol is required to connect an MLAG pair to servers or other switches. Interconnect to other switches or servers can use static LAG or IEEE 802.3ad LACP
- Up to 16 1GbE or 10GbE ports can be aggregated together and spread across the two switches
- Spanning Tree Protocol can still be used in conjunction with MLAG to detect and handle any misconfigurations

Applications of MLAG

MLAG can be used at various places in the network to eliminate bottlenecks and provide resiliency.

Leaf layer

At the leaf of the network, servers with multiple interfaces connect to the switches using NIC bonding, or standard Link Aggregation. With two of your leaf switches in an MLAG pair, you can now have active/active redundant connections from the server to the switches. The LAG groups can be formed using static link aggregation or LACP based negotiation methods.

Spine layer

At the spine layer, two switches can form an MLAG pair and aggregate all of the uplinks from the datacenter racks. This eliminates Spanning Tree blocked ports and allows you to use the interconnect bandwidth on all ports. MLAG at the spine layer goes a long way in improving the Layer 2 scalability of your network, without increasing your cost.

Summary

Data center and high performance cloud computing networks require both the highest level of network bandwidth and reliability. MLAG takes the benefits of link aggregation and spreads them across a pair of data center switches to deliver system level redundancy as well network level resiliency.

Arista's MLAG feature allows you to scale at Layer 2 without wasting bandwidth in Spanning Tree Blocked mode. Any networking device supporting static LAG or IEEE 802.3ad LACP can connect to an MLAG pair. Now you can scale your network, without changing your topology.

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. 01/12