Deploying Dante on Cisco SD-Access Networks

CISCO SD-ACCESS OVERVIEW 1

DEPLOYING DANTE ON OVERLAY NETWORKS WITH LAYER 2 FLOODING ENABLED 2

DEPLOYING DANTE ON LAYER 3 OVERLAY NETWORKS 3

ADDITIONAL RESOURCES 5

Cisco SD-Access Overview

Cisco SD-Access Networks consist of both an underlay and overlay network. The underlay network consists of the physical switches and routers that are used to deploy the SD-Access network. The underlay network establishes IP connectivity via the use of routing protocol. In SD-Access the underlay switches and routers support the physical connectivity for users and endpoints, but subnets and endpoints are not part of the underlay network – they are part of the automated overlay network.

The fabric overlay network is created on top of the underlay network through virtualization (virtual networks).

There are four basic planes of operation in the fabric overlay:

- **Control Plane—**Messaging and communication protocol between infrastructure devices in the fabric.
- **Data Plane—**Encapsulation method used for the data packets.
- **Policy Plane—**Used for security and segmentation.
- **Management Plane—**Orchestration, assurance, visibility, and management.

The primary technology used for the fabric control plane is based on the **Locator/ID Separation Protocol (LISP)**. (IETF standard protocol RFC-6830) LISP is based on a simple endpoint ID (EID) to routing locator (RLOC) mapping system, to separate the “identity” (address) from its current “location” (attached router). LISP simplifies traditional routing environments by removing the need for each router to process every possible IP destination address and route. It does this by moving remote destination information to a centralized map database that allows each router to manage only it’s local routes and the query the map system to locate destination endpoints.

The primary technology used for the fabric data plane is based on **Virtual Extensible LAN (VXLAN).** (IETF stand protocol RFC-7348) VXLAN encapsulation is IP/UDP based, meaning that it

can be forwarded by any IP-based network and effectively creates the "overlay" aspect of SD-Access fabric. VXLAN encapsulation is used because VXLAN includes the source Layer 2 (Ethernet) header and provides special fields for additional information (such as virtual network [VN] ID and group and group [segment] ID). This technology allows for both Layer 2 and Layer 3 virtual topologies (overlays). Layer 2 overlays are identified with a VLAN to VN ID correlation, and Layer 3 overlays are identified with VRF (Virtual Routing and Forwarding) to VN ID correlation.

The Control plane uses Cisco TrustSec, which decouples access that is based strictly on IP addresses and VLANs by using logical groupings in a method known as Group-Based Access Control (GBAC).

The management plane is enabled and powered by Cisco DNA Center, the centralized management system used to design, provision, and apply policy across the wired and wireless SD-Access network.

Deploying Dante on Overlay Networks with Layer 2 Flooding Enabled

When deploying Dante on a Cisco SD-Access network where the overlay–virtual network allows Layer 2 flooding, very few additional Dante network design considerations must be taken. VXLAN encapsulation allows the Dante devices to act and communicate as if they are in the same VLAN. QoS and IGMP management policies can be set up as outlined in the ***Adding Dante to Your Network*** document. Multicast packets (Dante discovery, clocking, control, and user defied multicast audio [audio is unicast by default]) are handled and routed through the underlay network "automatically."

The above said there are some precautions to take and known issues to be aware of.

First, you'll want to make sure the hardware and software components of your Cisco SD-Access Network are capable of and have the proper software/firmware versions to support Layer 2 Flooding, multicast routing, networking policies needed, etc. Cisco support can help with this.

Secondly, Cisco treats Precision Time Protocol (PTP) traffic different than other multicast traffic. Dante uses primarily PTP version 1. Any PTP settings available on Cisco network hardware/software are assuming PTP version 2. Unless needed by other devices on your network, PTP should be set into "*forward*" mode which will treat it like "normal" multicast traffic.

There is a known issue with PTP version 1 traffic outside of the default PTP domain (0) and Cisco SD-Access networks. **This affects devices only when enrolled into a domain using Dante Domain Manager.** When you enroll a device into a domain it puts the PTP data packets into a multicast group other than the default (*224.0.1.129)*, which causes them to not be forwarded properly. The outcome is a leader clock gets elected on a per edge node ("switch") basis. *(You*

can verify this is the issue if when unenrolled the devices sync properly.) Cisco is aware of the issue and will be pushing an update to fix it. If you are experiencing this issue it may just require a software update to your Cisco DNA Center, check with Cisco support for any software/firmware update(s) & version(s) needed.

If a software update is not yet available there are two possible workarounds:

- If the number of Dante devices in your network or within each virtual network ("VLAN") is less than 250 *(safe value)* and it's ok if those devices share a clock master:
    - Set one domain to ***Dante - Custom Clocking*** on the **Domain - Advanced Options** page and set the PTP domain number to ***0***.
    - Create a ***Shared Audio Group*** that contains all your domains.
    - If your Dante network contains multiple virtual networks (VLANs) a unicast boundary clock must be elected in each.
- If your network contains no or very little multicast traffic (< 25 MB – *safe value, can be up to 70MB depending on other network traffic*):
    - IGMP Snooping can be disabled allowing multicast to flood the Dante Virtual Network(s) ("VLAN(s)").

Deploying Dante on Layer 3 Overlay Networks

As Dante uses multicast traffic for several key components of operation special care must be taken when designing Dante networks in which all traffic is routed (Layer 3) between fabric edge nodes ("switches"). Unmanaged Dante uses multicast for device discovery (mDNS), clocking (PTP version 1) and some control messages. **Dante Domain Manager** (DDM) is a **REQUIREMENT** for routed Dante networks.

Devices need Dante firmware version 4.0+ in order to take full advantage of the DDM features including being able to send audio to devices in different subnets.

Dante devices and controllers can be discovered by the DDM server using DNS-SD (Service Discovery), when mDNS is not a viable option for discovery such as multi-subnet Dante systems or in a High Availability cluster. Each DNS-SD entry consists of an SRV record describing how to connect to the DDM Server and a TXT record with additional information (empty in this case). For Dante systems that span multiple subnets where discovery with DNS-SD is not an option the devices can be manually enrolled into DDM with their IP address. This can be done by either typing in device IP addresses or uploading a .csv file of device IP addresses. (See *Device & Controller Discovery Configuration* documentation for details.*)* Once devices are enrolled in DDM, instances of Dante Controller on the network are not involved in the discovery of those devices and instead reference the DDM server.

To accomplish clocking across multiple subnets, Dante devices capable of becoming unicast boundary clocks are required on each edge node (switch) or subnet. This can be configured in the **Domain – Advanced Options** page in DDM.

Devices **not** capable of becoming unicast boundary clocks:

- Devices associated in domains in legacy mode (devices with pre 4.0 Dante firmware)
- Computers running DVS or Via or applications running Dante Application Library (DAL)
- Devices running Dante Embedded Platform (DEP)
- Legacy Ultimo chipset devices
    - Chipset type can be determined in Dante Controller in ***Device View*** under the ***Status Tab***
        - Legacy Ultimo devices are listed as Ultimo or Ultimo4
        - Ultimo X devices are listed as UltimoX or UltimoX4
        - *Note: **Ultimo X** and **AVIO** adapter module chipsets **can** act as unicast PTP boundary clocks though more powerful chipsets (i.e. Brooklyn II, Broadway, HC) should be preferred and are required if more than 30 devices are to be sync'd to that device.*

Dante devices cannot pass multicast audio between subnets. Special care should be taken when designing the system that devices have the number of transmit/receive flows required to transmit/receive all channels unicast. (Number of flows available can be found in the Transmit/Receive tabs in Device view in Dante Controller.)

Control messages between devices, the DDM server, and Dante Controller on the network will be sent unicast.

Additional Resources

Cisco SD-Access Solution Design Guide:
https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html

Cisco Software-Defined Access 1.0 White Paper:
https://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise-networks/software-defined-access/white-paper-c11-740585.pdf

Cisco SD-Access Multicast:
https://community.cisco.com/t5/networking-documents/cisco-sd-access-multicast/ta-p/4068110#toc-hId-323918503

Adding Dante to Your Network:
https://my.audinate.com/sites/default/files/PDF/adding-dante-to-your-network-audinate.pdf

Dante Domain Manager User Guide:
https://dev.audinate.com/GA/ddm/userguide/1.1/pdf/latest/AUD-MAN-DDM-v1-1-1_User_Guide-v1.0.pdf

Which Network Ports does Dante Use:
https://www.audinate.com/learning/faqs/which-network-ports-does-dante-use