

White Paper

Intent Based Campus Fabrics With Software Defined Access

What's the difference between a traditional campus LAN and Cisco's Software-Defined Access (SD-Access)?

Following on from an article that we wrote regarding the differences between Catalyst IOS and NX-OS, this is the first in a series that seeks to simplify the differences between traditional networking and SDN (software defined networking).

Automation and Orchestration

The first and most significant difference is automation and orchestration.

Traditionally a LAN implementation, for example, would require low-level configuration of each and every device that comprises the LAN, which can be a laborious process, and it also means that any subsequent changes need to be made to each device. With SD-Access, network configuration is automatic and based upon a single policy that governs how the network will behave. Changes and additions are also automatic, with the changes to the policy being passed down to the devices by a controller.

SD-Access unifies both wired and wireless access together

Another important point is that SD-Access unifies both wired and wireless access together, whereby the same policies are propagated to both wired and wireless access networks at the same time.

A network policy controller (3rd party provisioning tools notwithstanding) is not a brand new concept in networking; however, the way that it's implemented for Cisco SD-Access is. The main principles of SD-Access operation are as follows:

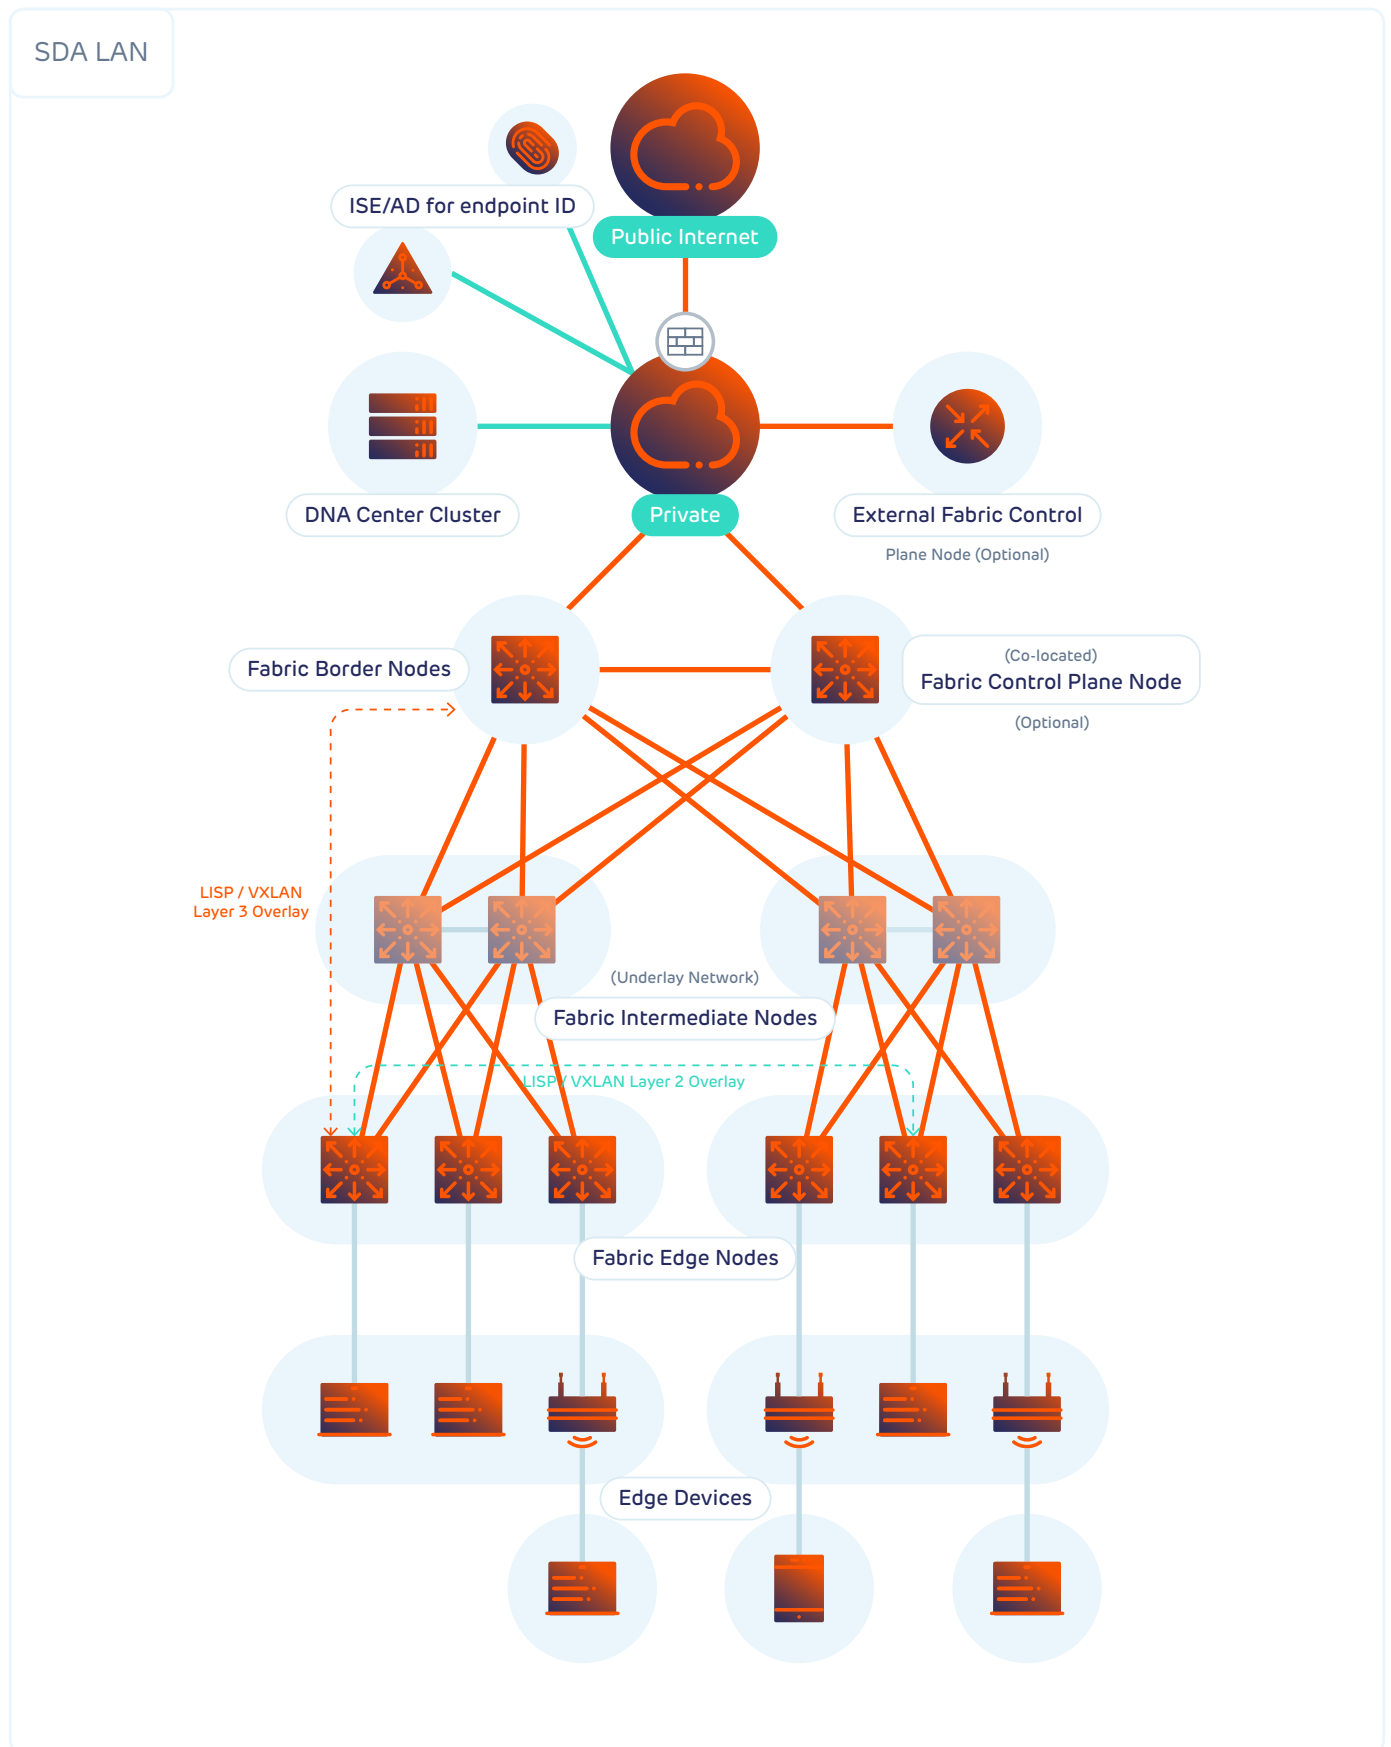
- SD-Access uses Cisco DNA Center to manage, monitor and implement network policies, the policies themselves are used to generate configurational changes, and then pass these changes to each of the devices either by direct SSH access or by a Restful API call.
- SD-Access uses an automatically built underlay network using IS-IS as its control plane, the network automatically builds the underlay based on the initial policy.
- Fabric border nodes are analogous to distribution / core switches, they allow the end hosts to talk to IP networks outside of the fabric, such as a data centre or the internet.
- Fabric edge nodes are analogous to access layer switches, they allow the end hosts to connect to the network.
- Security policies for the end hosts are defined in Active Directory (AD) or a combination of AD and Cisco Identity Services Engine (ISE). 802.1X is then used to authenticate devices wishing to join the network.
- Both wired and wireless access can utilise the same policies and be controlled by the same instance of DNA Center, allowing the two technologies to be controlled in the same way, with the same policies.

SD-Access Diagram

A diagram detailing the components of SD-Access is shown on the next page, the main points to note are:


- Communication between hosts and outside networks is implemented by the use of overlays. When traffic from a source is permitted to talk to a destination an overlay is built to carry this traffic.
- The overlays are implemented using LISP encapsulated into VXLAN. LISP and VXLAN function to tunnel the traffic between fabric edge nodes and from edge nodes to border nodes.
- There are two types of overlays used, a layer 3 overlay when traffic needs to be routed and a layer 2 overlay when traffic needs to stay within the same broadcast domain.
- A layer 3 overlay is used to send traffic outside of the fabric, such as the internet or a data centre.
- A layer 2 overlay is used to send traffic between two hosts and is roughly analogous to a VLAN.

SD-Access Diagram




Non SD-Access

If we consider a typical (non SD-Access) topology from an endpoint perspective, a physical topology would look something like the following diagram.



However from a logical standpoint, the following diagram shows the overlay network running on top of the underlay network. As you can see the logical topology has been greatly simplified.



The overlay in this example could support any of the following protocols it doesn't just have to be LISP with VXLAN encapsulation (which are specific to Cisco's implementation of SD-Access);

- GRE or mGRE
- MPLS or VPLS
- VRF lite
- L2TPv3
- IPSEC or DMVPN

The control plane for LISP is effectively the equivalent of DNS, where an end host location is registered to a mapping server, and a lookup is performed against the mapping server to find the associated router when a host wants to communicate with it. The mapping server role can be co-located with a data plane LISP router but for scalability, it doesn't have to be.

Summary of LISP

There are multiple roles and terminologies for devices that participate in LISP, they are summarised below:

LISP infrastructure devices

ITR (Ingress Tunnel Router) is deployed as a CE (customer edge) device. It receives packets from site-facing interfaces, and either encapsulates packets to remote LISP sites or natively forwards packets to non-LISP sites.

ETR (Egress Tunnel Router) is deployed as a CE device. It receives packets from core-facing interfaces and either decapsulates LISP packets or natively delivers non-LISP packets to local EIDs at the site. The LISP specification does not require that a device perform both ITR and ETR functions, however.

xTR It is common for LISP site edge devices to implement both ITR and ETR functions. When this is the case, it is referred to as an xTR.

MS (Map-Server) is deployed as a LISP Infrastructure component. It configures the LISP site policy for LISP ETRs that register to it. Including the EID prefixes for which registering ETRs are authoritative. Map-Servers receive Map-Register control packets from ETRs.

MR (Map-Resolver) is deployed as a LISP Infrastructure device. It receives Map-Requests encapsulated to it from ITRs. Map-Resolvers also send Negative Map-Replies to ITRs in response to queries for non-LISP addresses.

LISP logical terminology

EID (Endpoint ID) an EID is used to represent the IP address of the end host, and EID and associated RLOC are included in the ETR map register request.

RLOC (Routing Locator) an RLOC is an IPv4 or IPv6 address of an ETR. An RLOC is the output of an EID-to-RLOC mapping lookup (ITR map-request and map-reply).


LISP Interworking Devices

PxTR (Proxy ITR / ETR) is a LISP Infrastructure device that provides connectivity between non-LISP sites and LISP sites. This not only facilitates LISP/non-LISP interworking, but also allows LISP sites to see LISP ingress traffic engineering benefits from non-LISP traffic.

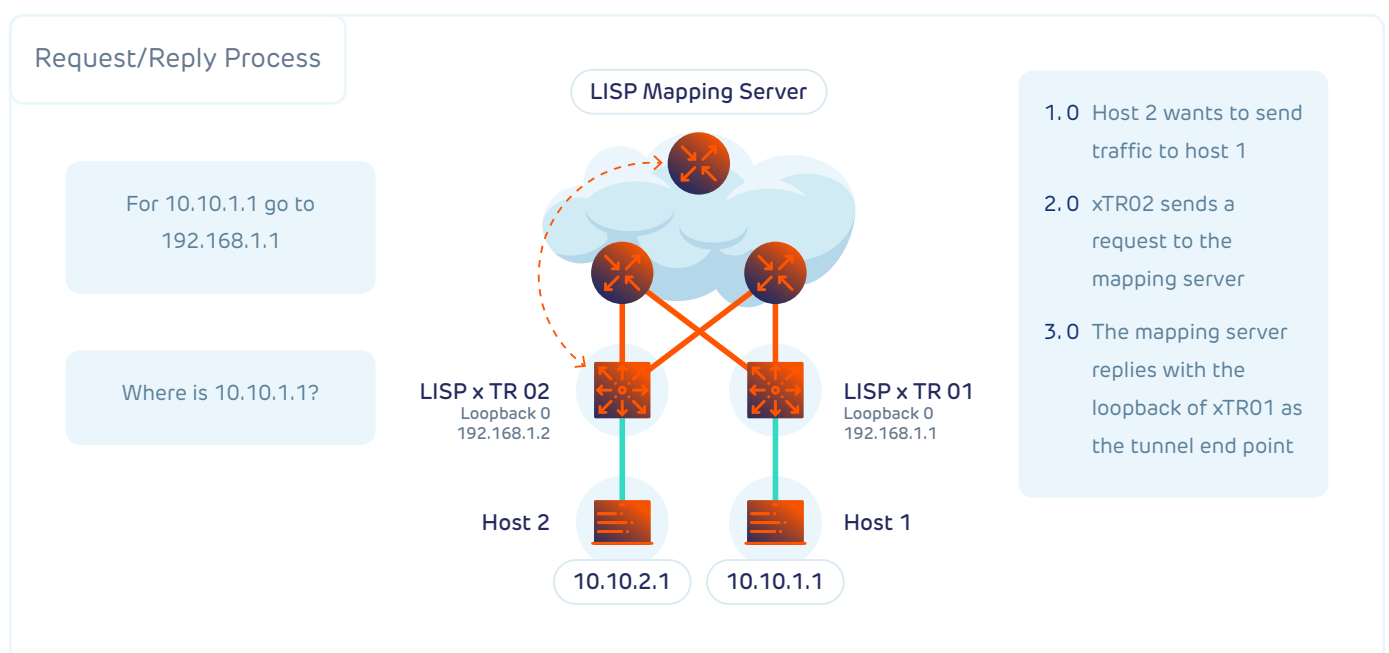
4-Steps to LISP Tunnel Creation

There are effectively 4-steps to LISP tunnel creation, and these are summarised below:

01 — LISP Registration



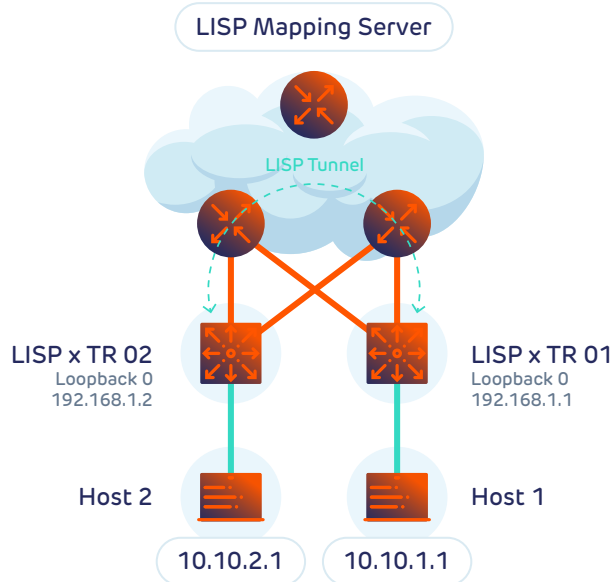
02/03 — LISP Lookup/Reply



4-Steps to LISP Tunnel Creation

04 — LISP Tunnel

Tunnel Process



1.0 xTR02 sets up a tunnel from itself to xTR01 allowing the hosts to communicate



Marc Mawson
Senior Consultant
Forfusion

Author Bio

Marc Mawson is a Senior Consultant for Forfusion in the pre-sales team. Marc is an IT professional, boasting over 15 years experience within the industry.

With over 10 years experience designing, supporting and implementing enterprise networks. Marc's role gives him the opportunity to work with some cutting edge vendor technologies such as SD-WAN / LAN.

Marc specialises in SD (Software Defined) networking, for the Campus, WAN and Datacentre. However, his overarching specialisation is in the datacentre and cloud space.

He holds a CCIE (Cisco certified internetworking expert) qualification in Datacentre. The CCIE is a coveted top tier networking qualification, and is seen as the de facto for consultants/network architects.

Marc finds the fact that most network engineers have very similar interests fascinating and believes that each piece of knowledge you acquire teaches you how much you still don't know.

Forfusion can offer a complete "service wrapper" for software defined access, which includes design and implementation, carried over into a support cycle. This helps businesses leverage these seemingly complex environments so they can reap the benefits without taking on the risk.

Thank you

**I hope that this white paper
has been informative and
useful to you.**

To receive notification of our next exclusive white paper,
sign up to our mailing list today at www.forfusion.com