Software-Defined Access for Secure Networks

Contents

The modern secure network

Traditional enterprise networks are complex systems. For the Department of Defense (DoD), managing this complexity in a fast-changing digital landscape can be challenging. Securing these networks in the modern era is even more difficult. The network has become a weapons system, and the efficiency and security of this system are more critical than ever. This changing environment is driving the military's need for a new, modern network infrastructure that can deliver scale, agility, and security in a simplified way.

The enablement of such a modern secure network requires:

Consistent management: Consistent management is key to enabling scale, agility, and security. Using a centralized Software-Defined Network (SDN) controller, network administrators can turn up new network infrastructure faster than ever before. SDN enables network change in a secure, error-free manner, drastically reducing human error in configurations at scale.

Automated network segmentation and group-based policy: Modern networks require a zero-trust security model. Traditional methods of providing network segmentation are not sufficient when enabling zero-trust architectures at scale. In order to accomplish zero trust, the modern network must be able to classify and segment users based on group policy. This means we must be able to assign policy from a centralized location based on:

- Who: User identity
- What: Device type, for example, corporate laptop, cellphone, etc.
- When: The time of day access is requested
- Where: The location of the access request, for example, corporate office, coffee shop, home office, etc.
- How: Wired, wireless, or VPN

Contextual insights for fast issue resolution: The network must be able to proactively predict network-related and security-related risks by using telemetry to improve the performance of the network and reduce the mean time to resolution.

Open and programmable interfaces: Network devices and controllers must be programmable and open in order to enable interoperability within the infrastructure. For the DoD it is critical to be able to integrate between Military Departments (MILDEPs) and with coalition mission partners. In modern networks this is achieved with Application Programmable Interfaces (APIs).

An intent-based solution

In the era of multidomain operations, it is critical for the network to reflect the commander's intent. To achieve this, military leaders must adopt a radically new approach to networking. The current rigid, manual lifecycle management approach is no longer sustainable for deploying, maintaining, and updating networks, and it cannot scale to meet the growing complexity. For the DoD to dominate cyber, space, air, land, and sea, the network needs to be able to adapt quickly to changing mission requirements or "intent." The network needs to support an increasingly diverse and fast-changing set of users, devices, applications, and services. It needs to ensure fast and secure access to and between workloads wherever they reside. And for the network to work optimally, all this needs to be achieved from end to end, between users, devices, applications, and services in both the tactical and strategic environments.

Cisco® Software-Defined Access (SD-Access) is an Intent-Based Network (IBN) solution that can accomplish just that. SD-Access leverages intent to make IT operations more efficient, networks more secure, and the user experience more consistent.

The Cisco SD-Access advantage

Cisco SD-Access provides organizations with a broad set of capabilities via fabric overlays. SD-Access uses programmable overlays that can enable many logical networks over the top of a physical underlay network to enable design intent. Leveraging fabric overlays, Cisco SD-Access helps ensure user mobility, cloud and data center integration, and segmentation through a zero-trust security architecture.

Network virtualization

Taking segmentation a step further, Cisco SD-Access can create Virtual Networks (VNs) that reside on top of a single physical infrastructure. These VNs are created, and then encoded as VNID's into the Virtual Extensible LAN (VXLAN) header, while leveraging Location/ID Separation Protocol (LISP) as the control plane for end-point location within the SD-Access fabric. These virtual networks are completely isolated from one another and can communicate only via a secured policy control point, such as a fusion firewall or router. If implemented, the fusion device is the point of enforcement for communications flowing between VNs; otherwise, the VNs would remain isolated.

Zero-trust security

Cisco SD-Access helps ensure a zero-trust security model in the network by implementing software-defined segmentation and policy enforcement based on user and/or device identity and group policy. Access is granted to identified users and devices only if they are properly authenticated. Once authenticated, they can be granted least-privilege access based on the desired authorization policy of the command. Software-defined segmentation is achieved using Cisco TrustSec®, which uses Scalable Group Tags (SGT) to classify traffic into groups, removing the need to segment by IP address, VLAN, etc. This means that a user can move around the campus, changing IP addresses along the way, and retain the same policy. The policy is centrally managed in Cisco DNA Center via an API integration to the Cisco Identity Services Engine (ISE). Cisco ISE pushes the policy to every network device in the fabric by way of security group ACLs, helping ensure policy enforcement on a hop-by-hop basis within the fabric.

Monitoring and troubleshooting

Cisco DNA Assurance provides unprecedented monitoring and troubleshooting capabilities through the use of machine learning, artificial intelligence, and streaming telemetry. It is able to save time, reduce work, and speed troubleshooting by reducing noise and false positives. When real network faults are detected, Cisco DNA Assurance can provide guided troubleshooting procedures, reducing the mean time to resolution for any organization.

Software Image Management

One of the biggest vulnerabilities in any network infrastructure is outdated software. The Cisco DNA Center Software Image Management (SWIM) solution allows for the centralized storage of software images for all devices in the infrastructure. One image per device type can be selected as "golden," indicating that it is the standardized image for the organization. SWIM manages software upgrades and controls image version consistency across the network, taking the software upgrade process from days and hours to minutes. Updates can even be scheduled to fit into an authorized service interruption.

Day-zero provisioning

Initial device deployment can be time-consuming and error prone. Network Plug and Play enables Cisco network devices to connect to Cisco DNA Center and download the appropriate software and device configuration. This process not only saves time but also helps ensure that the devices are running the approved software version along with a properly secured, error-free configuration.

Feature-rich APIs

Cisco SD-Access components all feature an open set of APIs that enable interoperability not only between the SD-Access components themselves, but also with third-party vendor ecosystem components. Using APIs between Cisco SD-Access and external systems enables closed-loop automation systems to be deployed in a programmatic way. Network intelligence can easily be passed to a third-party system, enabling the systems applications to analyze the data, or even respond back to the network. This type of closed-loop interoperability not only increases the organization's security posture but also reduces the mean time to detection and remediation.

Cisco SD-Access benefits

Table 1. Features and benefits of Cisco SD-Access

Features	Benefits
Secure, policy-based automation	Cisco SD-Access uses policy-based automated network provisioning across all network domains as well as simple segmentation constructs to build secure boundaries for users. • Networkwide policy enforcement regardless of location • Policy administered from a central dashboard • No IP address dependency with anycast gateway and SGTs • Policy defined once for LAN, WLAN, and WAN
Fast, easy service enablement	Reducing the number of manual configuration steps improves the efficiency of network operations and reduces human error. Cisco SD-Access also quickly enables services by using open APIs across a services ecosystem. • Devices deployed using best practice configurations • Simple user interface • Easy orchestration with objects and data models • Native third-party application hosting
Complete network visibility	The entire wired, wireless, and WAN network is managed on a single entity. Application visibility eliminates the complexity of managing separate policies for wired and wireless. • Consistent policy and management across wired and wireless • Optimal traffic flows with integrated roaming • Ability to instantly find any user or device • Enhanced visibility for troubleshooting user challenges
Enhanced analytics	Cisco SD-Access provides intelligent services for application recognition, traffic analytics, traffic prioritization, and steering. • Detailed analytics help you plan for future growth and diversification, as well as make more informed decisions • Track access point performance, heat maps, and channel information • Information includes data termination of wireless traffic, network forensics, and user-based application intelligence

Cisco SD-Access solution components

Figure 1 and Table 2 provide a high-level overview of the components used in the Cisco SD-Access solution. For a more comprehensive review of these components and their respective roles, refer to the Cisco Software-Defined Access Solution Design Guide. For compatibility of SD-Access components, please refer to the SD-Access Product Compatibility Guide.

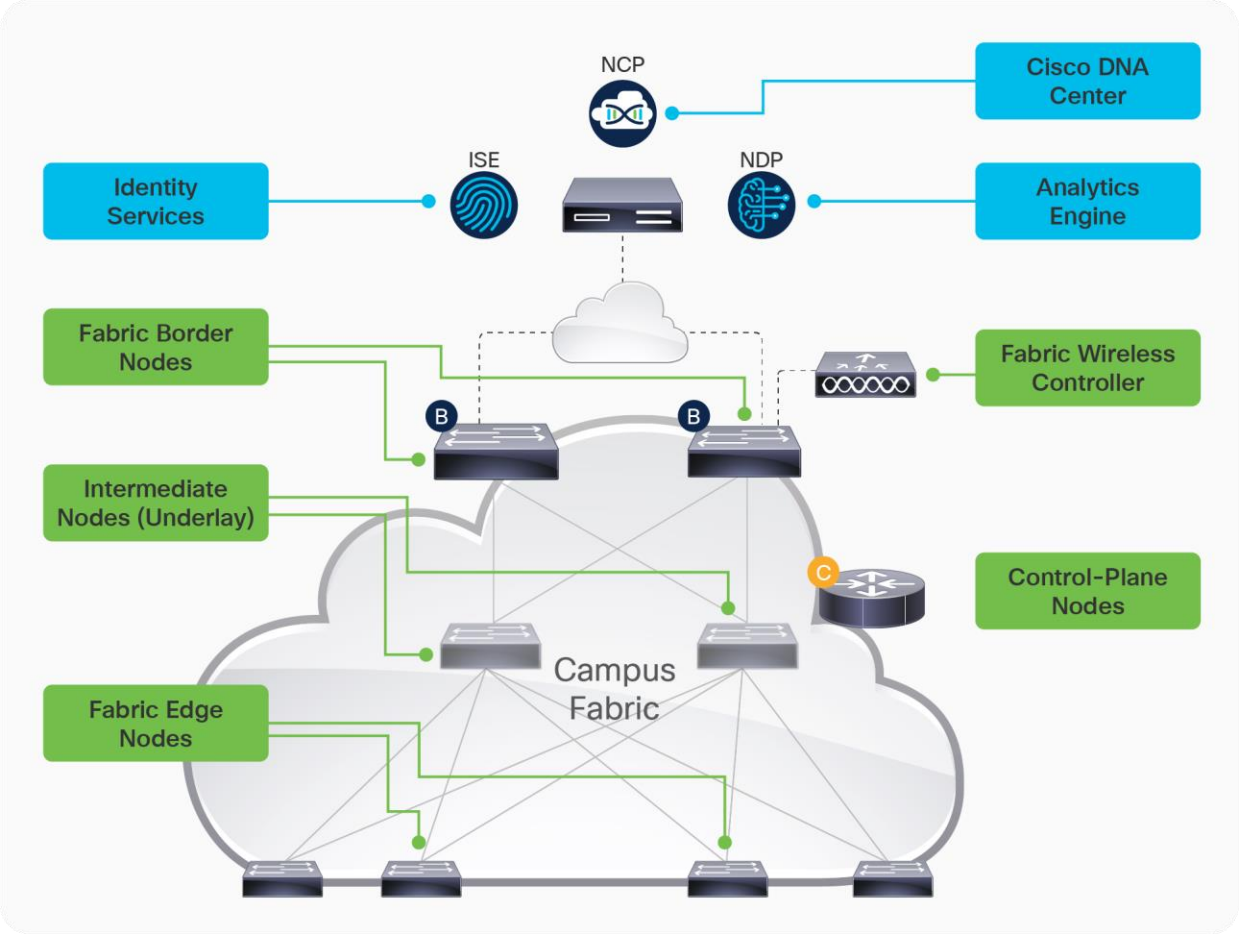

Figure 1.
Solution components

Table 2. Cisco SD-Access components

Component	Description
Cisco DNA Center	At the heart of automating the SD-Access solution is Cisco DNA Center. SD-Access is enabled with an application package that runs as part of the Cisco DNA Center software for designing, provisioning, applying policy, and facilitating the creation of an intelligent wired and wireless campus network with assurance.
Cisco Identity Services Engine	Cisco ISE is a secure network access platform enabling increased management awareness, control, and consistency for users and devices accessing an organization's network. ISE is an integral part of SD-Access for policy implementation, enabling dynamic mapping of users and devices to scalable groups and simplifying end-to-end security policy enforcement.
Fabric border node	The fabric border nodes serve as the gateway between the SD-Access fabric site and the networks external to the fabric.
Control plane node	The SD-Access fabric control plane node is based on the LISP map server and map resolver functionality combined on the same node. The control plane database tracks all endpoints in the fabric site and associates the endpoints to fabric nodes, decoupling the endpoint IP address or MAC address from the location (closest router) in the network.
Fabric edge node	The SD-Access fabric edge nodes are the equivalent of an access layer switch in a traditional campus LAN design. The edge nodes implement a Layer 3 access design with the addition of the following fabric functions: • Endpoint registration • Mapping of user to virtual network • Anycast Layer 3 gateway • LISP forwarding • VXLAN encapsulation and decapsulation
Fabric intermediate node	The fabric intermediate nodes are part of the Layer 3 network used for interconnections among the edge nodes to the border nodes.
Fabric wireless controller	The fabric WLC integrates with the fabric control plane. Both fabric WLCs and non-fabric WLCs provide AP image and configuration management, client session management, and mobility services.

Legacy architecture review

In most secure environments today, customers are operationally required to overlay a secure encrypted network over a nonsecure encrypted network. As with Cisco SD-Access, this is accomplished by using an overlay. These encrypted overlays are created by using IPsec VPN devices (IVDs). IVDs can be configured in a point-to-point configuration or in a hub-and-spoke configuration, as seen in Figure 2.

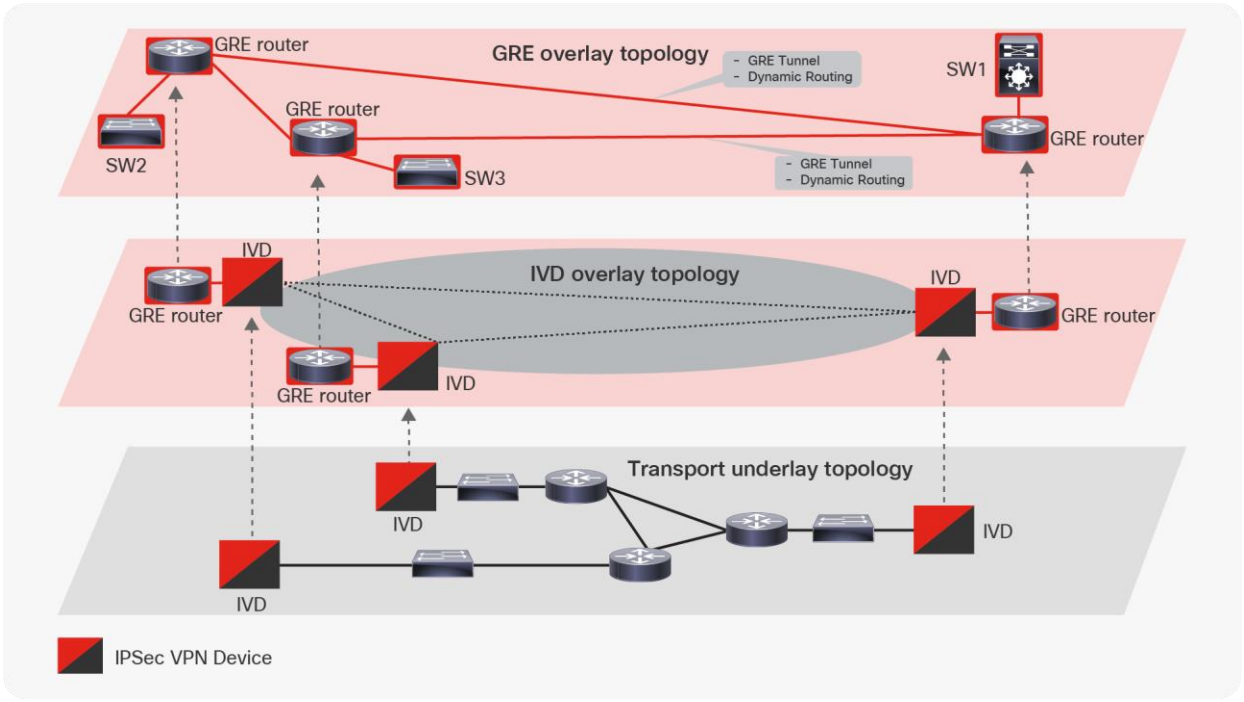

Figure 2.
Legacy LAN extension topology using IVDs

Most IVDs in use today lack critical features required for DoD networks, such as certain dynamic routing protocols, multicast, etc. Because of these limitations, it is very common to see a secondary overlay implemented. This is commonly accomplished by using Generic Routing Encapsulation (GRE). GRE is used to create a tunnel over the top of the IVD overlay, hiding the IVDs from the topology. When this type of network is implemented, the IVD overlay becomes the underlay network for the new GRE overlay. All traffic flows between IVD encrypted sites are encapsulated in GRE, enabling dynamic routing, multicast, and other required features. This can be seen in the GRE overlay topology in Figure 2. Traffic flows originating from SW2 and SW3 would be encapsulated in GRE, encrypted by the IVD, forwarded to the main site for decryption and decapsulation, and then forwarded normally to SW1.

The use of GRE, however, is not without its challenges. In many instances secure networks are being extended on a campus LAN to a small set of users. Under normal circumstances a switch would suffice, but when using IVDs this extension would require a switch and router. A router is often required due to no or limited GRE capability on most industry switches. This leads to increased CapEx and operational complexity, as even a simple LAN extension requires a router. Scale can also be a challenge when using GRE in this manner. If the network administrators want to add a new network or Virtual Routing and Forwarding (VRF) instance, a separate GRE tunnel needs to be built. In some cases, this can lead to dozens of GRE tunnels for a simple LAN extension.

Improving the design with Cisco SD-Access

SD-Access encapsulates traffic natively using VXLAN and LISP. This encapsulation eliminates the need for a manually built GRE tunnel overlay. As a result, OpEx is reduced greatly by eliminating the manual configuration of GRE tunnels between each of the routing elements, including the necessary routing protocol creation over the GRE tunnels (more state on the routers). CapEx in the organization is reduced because locations using IVDs now require only a switch to extend the campus LAN. In addition, since SD-Access encapsulation works by forwarding traffic between a source and destination IP address, no additional configuration is required in the IVD underlay. The use of SD-Access also creates the opportunity to arbitrarily add VNs without any additional configuration of the tunneling technology.

SD-Access technical terms

Since using these features requires knowledge of how SD-Access encapsulates traffic, a review of the related terms is necessary.

VN: Virtual network.: VNs provide macro-segmentation between networks by containing routing information for each network in a separate VRF instance. For each VN provisioned in Cisco DNA Center, a corresponding VRF will be created on each device belonging to that VN.

LISP: Location/ID Separation Protocol. LISP separates a location from an identification on the network by using two separate identifiers for an entity. The routing locator, or RLOC, assumes "ownership" of certain networks by registering that information with a map server. To send traffic to a certain network, you need to send the traffic only to its corresponding RLOC address.

xTR: Each router inside a LISP network can act as an Ingress Tunnel Router (ITR), an Egress Tunnel Router (ETR), or both (xTR). For the purposes of this design, every router has both roles and thus is an xTR.

Proxy-xTR (PxTR): A proxy XTR is a router that has the capability of assuming "ownership" of networks that reside outside of the LISP network. In SD-Access, this includes any shared services, such as DHCP, DNS, ISE, etc. Any traffic destined to those services will be encapsulated and sent to the PxTR for processing.

VXLAN: Virtual extensible local area network. VXLAN works by encapsulating an entire Ethernet frame inside of a User Datagram Protocol (UDP) packet. This allows networks to cross Layer 3 boundaries, which is crucial to SD-Access. Traffic in the SD-Access overlay networks is encapsulated in VXLAN packets.

Network design recommendation

The purpose of this design recommendation is to provide guidance on Cisco's preferred method for implementing SD-Access when external IVDs are required for the purpose of network encryption. This paper also serves as a solution to the often-required use of a router configured with GRE to extend campus switching between buildings when IVDs are used.

Technical overview: Next-generation secure network with SD-Access

With SD-Access technologies, the switch is functioning as both an access switch, with switched virtual interfaces (SVIs) for various segments, as well as a router, by forwarding traffic throughout the campus fabric using the native SD-Access forwarding. This is accomplished through the following steps:

Step 1. The border/control plane node for the fabric is configured as an external border. This enables Proxy-xTR functionality on the border nodes, allowing them to encapsulate and decapsulate traffic between the fabric and external networks.

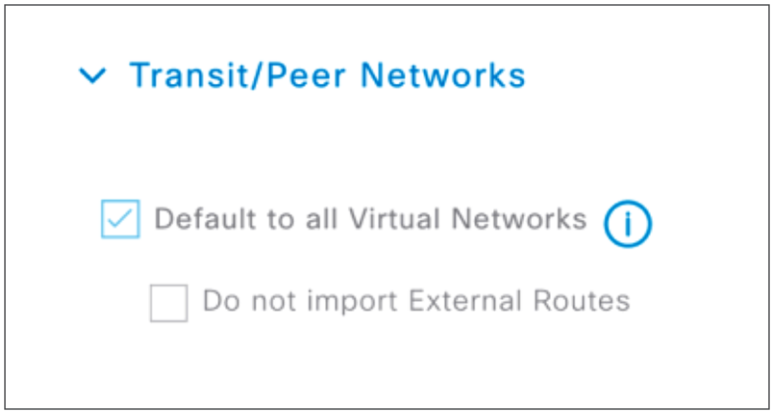

Step 2. The INFRA_VN is activated for the fabric, which will create a LISP instance for the Global Routing Table (GRT).

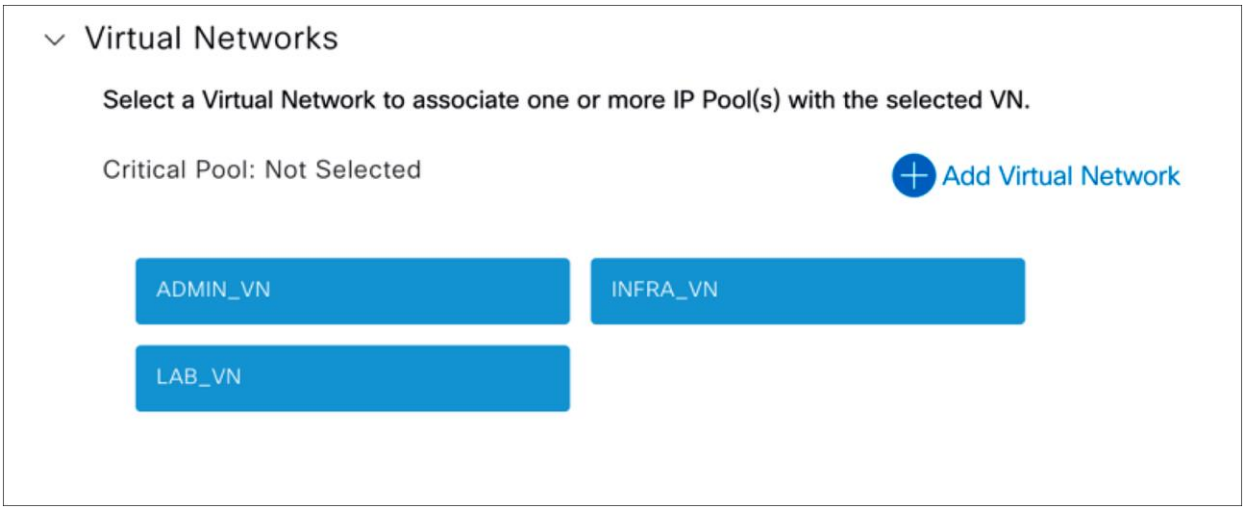

Tech note: CLI commands depicted in step 4 are for informational purposes only. Actual switch configuration will be done automatically using Cisco DNA Center.

Step 3. Since the border is now able to advertise prefixes that exist outside the fabric, we need to bring those prefixes in to the GRT on the border. Since BGP handoffs will have been created for the overlay networks, we will need to create one for the underlay network as well. This will allow Fabric Edge devices to query the Border/Control Plane node and forward their encapsulated traffic appropriately.

Tech note: A route-map is recommended to control the flow of routing information and reduce the possibility of loop formation.

Step 4. We need to tell the fabric edge devices behind IVDs that GRT traffic can be found by querying the LISP database. We do this by adding the following line under the LISP configuration:

```
!
instance-id 4097
remote-rloc-probe on-route-change
dynamic-eid INFRA_90-IPV4
database-mapping 10.64.90.0/24 locator-set rloc_c3bf491f-a0d0-40b0-b074-f824406ecd8e
exit-dynamic-eid
!
service ipv4
eid-table default
map-cache 0.0.0.0/0 map-request
exit-service-ipv4
!
exit-instance-id
```

Tech note: At the time of this writing, a static route is necessary in the GRT on each fabric edge device to direct traffic through the LISP process for encapsulation. A fix for this is in development.

Example: `ip route 0.0.0.0 0.0.0.0 LISP0.4097 <Border RLOC Address>`

Tech note: Since VXLAN and GRE are both tunneling technologies, in both cases the size of the packet containing the tunneled traffic will be more than 1500 bytes. This will result in fragmentation if the Maximum Transmission Unit (MTU) isn't controlled throughout the fabric. To overcome this, the SVIs that are created by Cisco DNA Center in the fabric can be configured to have a lower MTU, ensuring that the size of the tunneled packets stays within limits. Additionally, the TCP Maximum Segment Size (MSS) will need to be lowered to 40 bytes less than the configured MTU on the SVIs. This can also be done on the SVIs created on the border node to prevent fragmentation of traffic destined for hosts within the fabric.

Tech note: Configuration of IVDs is out of the scope of this document, but there are some recommendations.

- **This solution works in either a hub-and-spoke or full-mesh configuration of the IVD tunnels. Careful planning of the physical paths of traffic is required in cases such as VOIP and multicast.**
- **The only routes required in the IVDS are between the different RLOC addresses in the fabric.**

Since IVDs cannot be configured automatically, static routing must be used to ensure reachability between the various RLOC addresses in the fabric. Additionally, each fabric device will require static routing to reach the other RLOCs in the fabric. Routes can be summarized from fabric edge nodes to other fabric edge nodes, but not to the border or control plane nodes.

Because the same IP address is used on the fabric edge switch for both management and traffic sourcing, the RLOC addresses are no longer locally significant and must be reachable by the rest of the network. This is done by redistributing the static routes from the border and control plane devices to the fabric edge devices back into the campus network. It is recommended to use a summary route to limit routing configuration.

In summary, we have created an additional VN for the GRT in a manner identical to the creation of other VNs that exist within the fabric. This has several useful consequences. First, since VXLAN is now encapsulating all traffic to and from the switch, the routing information required on the IVDs is for the encapsulated traffic that has been "wrapped" in a VXLAN header and sent across the network. This is **identical** to what is accomplished with GRE, but with the added benefits of SD-Access. Additionally, since the switch is still managed through Cisco DNA Center, new networks and services can be provisioned on the fly, without any manual configuration required in the fabric or IVDs.

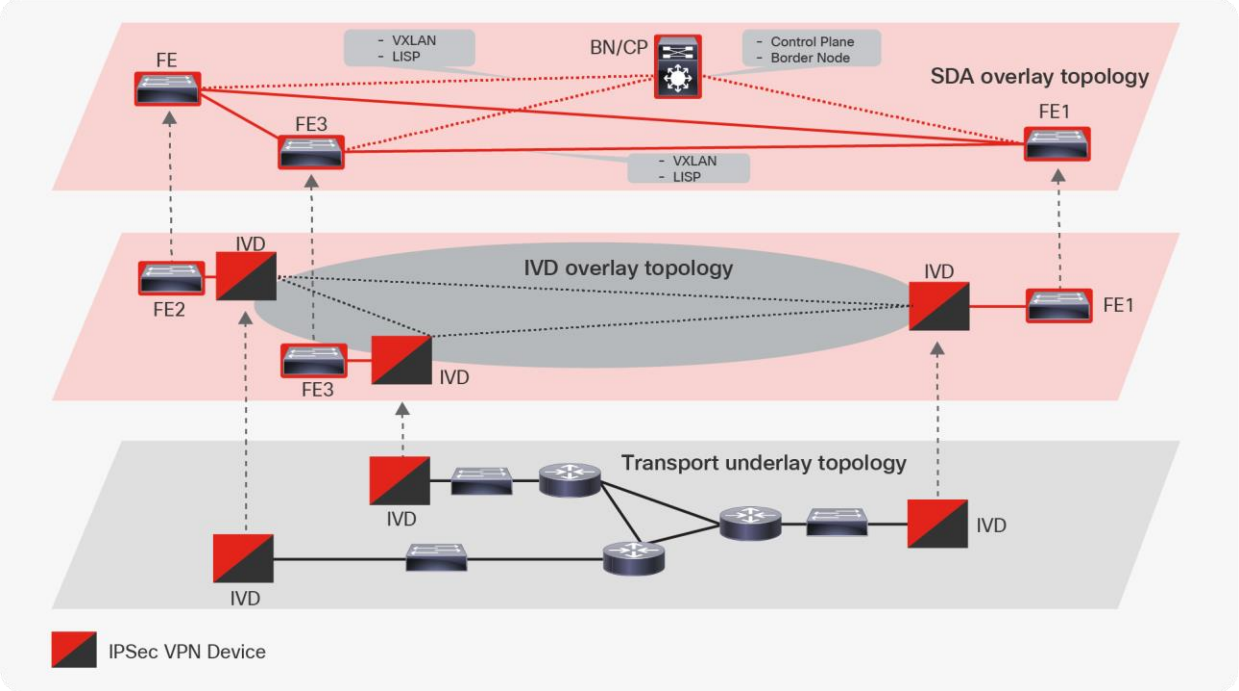

Figure 3.
Cisco SD-Access design

In this example, there are three fabric edge devices and one border/control plane in the fabric. IVDs are placed between the various LAN extensions and the main site. Traffic from a user on FE2 destined for a user on FE3 will be encapsulated and sent to FE3 for processing. The only routing information required in the IVDs is the addresses of the corresponding RLOCs that traffic will be sent from/to. The role of the border node in this scenario is to act as the control plane device, telling the various fabric edge devices which RLOC address to send their traffic to. If there is another fabric edge device residing in another place on the LAN, the only requirement is that the RLOC addresses are reachable. This may require redistributing the static routes from the border node back into the campus LAN.

Summary

Customers who are required to secure their networks using IVDs can do so in a manner that is both cost-effective and operationally effective by using Cisco SD-Access. Not only does Cisco SD-Access overlay the IVD encrypted network more efficiently than other methods, but it also provides many other security enhancements that are necessary to implement a zero-trust architecture. Using the methods outlined in this paper, we can close the gaps in visibility, consistently enforce policy across the entire network, and help ensure least-privilege access to users.

For additional information, contact Aaron Warner, Systems Architect, at aawarner@cisco.com or Todd Vogt, Network Consulting Engineer, at tovogt@cisco.com.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA C11-743733-00 07/20