The CIO's Guide to:

CISCO SDA FABRIC

The world of Cisco Software Defined Networking Explained

IPTel Solutions
Network Engineering Excellence

sales@iptel.com.au
www.iptel.com.au
SYDNEY - MELBOURNE - BRISBANE - CANBERRA

For great technical insights and information, visit our blog at:
https://blog.iptel.com.au

CISCO SDA FABRIC

Thank you for choosing to download this ebook, brought to you by the engineers at IPTel Solutions.

We aim to work with our customers to empower and educate in the use of the latest in networking technologies.

This ebook has been designed to walk through some of the key concepts in SDA Fabric, as well as define some key terms and use cases. We hope this is useful for you as you embark on your SDA Fabric journey and embrace the ease of automation and security of Zero Trust.

If you need any help with your Cisco Catalyst Center or Cisco SDA deployment, we are here to help: email us at sales@iptel.com.au and we can help you start your SDA journey!

Please enjoy this ebook,
The IPTel Team

Any comments or feedback on this publication, please email:
Mark McSherry
mark.mcsherry@iptel.com.au

BACKGROUND TO SDA FABRIC

The Background to Cisco SDA Fabric

Let's start at the beginning – what ***actually*** is Fabric? You hear this term spoken fairly frequently in networking, but let's kick off with what Fabric actually is:

- Fabric is a generic industry term.
- Fabric is an overlay – it exists in datacentres, service provider or campus.

Historically we could use MPLS as a Fabric, but with SDA Fabric, we use VXLAN. The concept has been around for a while, but its now getting a refresh.

We have a three-tier traditional network, which has been around for a long time. Typically this is built from using a mix of Layer 3 protocols (routing on the core switches for example) and Layer 2 protocols on the distribution and edge / access (spanning tree and a multitude of other protocols – stacking and securing the switching environment).

Wireless was never an original part of network architecture, and it's developed with a lot of tunnelling (CAPWAP tunnels to connect the APs to the WLAN controller), which is really building an overlay network across the traditional collapsed core or three layer network.

If you were architecting a network architecture now though, and needed to implement wireless, would you deploy the network architecture differently – the answer of course, if yes – and this is where Fabric comes in.

There is another shortcoming that traditional network deployments raised – which is security. Often networks are deployed and are insecure – with open ports, and free access between different clients and security domains. To address this, more protocols and overlays have been developed – 802.1X and a multitude of supplicants to sit on the end devices.

These are not always successful, and are not suited to 'black box' devices (think IoT devices), so a work around was needed for these - this is called MAB (MAC Authentication Bypass).

MAB essentially allows a workaround when a device cannot properly authenticate, but is itself susceptible to spoofing. In order to get around this, we then have profiling introduced, where we look at the traffic from that device, to really ensure it is that device.

It's exhausting following all this and each time a loophole emerges, we add another layer of complexity to the network authentication solution to close it.

This really sounds like network security is an afterthought, overlaid onto the traditional network.

All these protocols and layers have to be deployed and work together seamlessly – any misconfiguration and your network is wide open, and of course, the ever present danger as you add more features, you can expect more bugs.

IPTel Solutions
Network Engineering Excellence

SDA FABRIC KEY POINTS

What are the key benefits of SDA Fabric?

Let's take a look at a quick high level list:

- **Wireless is Designed as part of the Fabric:** The Fabric and not an overlay the way it currently is deployed
- **Security is inherent with Fabric:** Zero Trust is the major selling point
- **SDA Fabric can Support any Network Topology Equally:** Three tier, collapsed core, leaf and spine, daisy chain
- **Graphical Driven:** The deployment of policy and network is easy to do, with no CLI requirements
- Scale for large numbers of endpoints, with automatic trust ratings and quarantine ability
- Endpoint visibility to dynamically change end point trust levels and automatically quarantine suspicious clients

Key Aspects of Cisco SDA Fabric

Before we get into the technical detail of how Fabric works, let's define some key aspects that are used to build Cisco SDA Fabric:

- **Underlay**: The connection of the underlying network equipment. For this, a Layer 3 routing protocol is used (OSPF, IS-IS EIGRP or even static routing, for example). The underlay is just a transport – its all about a VXLAN tunnel operating between switch Loopbacks:
  - Maybe we should define why a Loopback? A loopback is a virtual interface – so it is not subject to a physical interface taking it up or down – its permanently up, which is a handy thing to separate the logical from the physical
- **Overlay**: On top of that Layer 3 network a set of point-point tunnels are built. The end user traffic flows over these

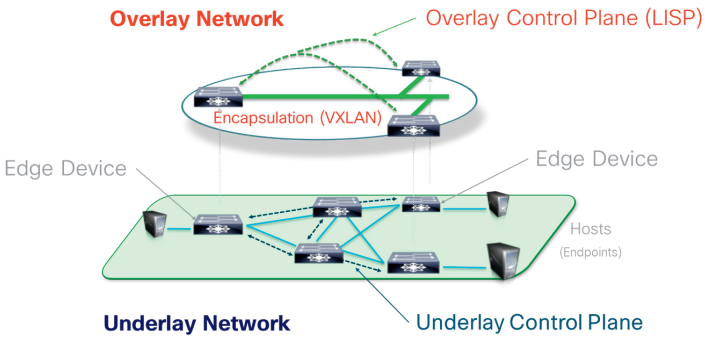

- **Management Plane:** This is the control of the whole network- the management station essentially translates GUI driven input ("intent") into CLI that is subsequently pushed to switches. It is out of band from fabric.
- **Tunnel:** A tunnel between two switches is just an encapsulation. This means traffic flowing from point-point is effectively switched across a tunnel, rather than being routed - but we get to avoid the shortcomings of our legacy Layer 2 networks (STP, etherchannel, etc)

WHY DO I NEED SDA FABRIC?

Why do I need SDA Fabric?
There are a number of reasons you might want SDA Fabric and it's important to have these firmly fixed in the mind at the start of any SDA Fabric discussion. These are the business outcomes by which any deployment success will be judged.

Here's the quick summary:

- **Network Automation**
- **Understanding of what is connecting**
- **Implementation of Zero Trust (Micro Segmentation)**

First up – ***automating your network operations***. This is an easy sell and a major win – the network operations are set to be undertaken more efficiently and without human error..

Next up, we have the outcome of ***understanding what end points are connecting to the network***. With the proliferation of IoT devices, as well as users taking it upon themselves to connect equipment, having a solid understanding of what is actually connected to your network is of vital importance.

The third – and really the key business outcome from SDA Fabric is ***micro-segmentation***. This the often-quoted concept of Zero Trust. What Zero Trust means is that connecting devices or users get no access unless it is specially allocated. Consider a normal network, where a device connects to an access port.

Typically, a device can see all the other devices on the same VLAN, as well as route across VLANs, with little or no restriction. Micro-Segmentation provides very specific access that you allow – for example, a device can connect to the internet, a printer and a specific database, but nothing else.

Zero Trust really turns each access switch into a firewall – how much more secure would your network be, if this were the case?

Even if you have a firewall in the core of your network, to separate your VLANs, these are built from access lists and rules – these often grow to be long and complex and subject to human error when changed. Applying a policy to a connecting device means this all goes away.

There's also a fourth – and worth talking about benefit – ***AI Endpoint Analytics***. Cisco SDA compatible switches have deep packet inspection capabilities on the switch, and can identify the endpoints by the DHCP / LLDP / CDP packet (profile the device), but once the workstation starts sending its actual data, the switch can determine if the device really is what it's meant to be.

This is sent as telemetry to Cisco Catalyst Center / ISE and this allows you to inspect end-point behaviour to confirm it is only doing what it is mean to be doing. This network self-awareness allows the network to change the trust level on a device, based upon its ongoing behaviour, is incredibly powerful. You can configure ISE policies to automatically contain any client which has a low trust level.

IPTel Solutions
Network Engineering Excellence

SDA FABRIC IS LAYER 3

SDA Fabric is Layer 3

Ethernet is ubiquitous and has been around forever (well, not exactly forever... we used to have SNA, Appletalk, Banyan Vines, and others – looking at you, Token Ring). Token Ring was often said to be superior to Ethernet – but Ethernet had a major advantage: it became a cheap technology to deliver mass connection of clients.

Once volumes started to increase and investment centred around one technology, the speeds of ethernet spiralled. However – we still have the drawbacks of Layer 2.

With SDA Fabric, we still have Layer 2- but only in a small, limited way. Layer 2 exists on the edge switch only – our laptop communicates with the switch via Layer 2, and from this point forward is encapsulated in a tunnel (this is called VXLAN).

This is the same process in campus, service provider and datacentre. The network is fully routed in the underlay (where the network converges), but switched in the overlay (where the user traffic is moved around the network).

The underlay network is now simple: lots of point-point links over which routing operates and essentially connects all the switches together as Layer 3 switches.

As mentioned above, the Layer 2 traffic is encapsulated within a VXLAN tunnel – but what does that actually mean?

The underlay routes between all switches – the overlay (unsurprisingly) sit on top of this. The Fabric builds tunnels between each of the switch loopbacks, over which the Layer 2 traffic can be switched directly to the destination. These tunnels are built using VXLAN.

Simply put: Under the surface, the network is full Layer 3. From a client perspective, the overlay is point-point, with 1 hop from any point to any other.

So – what is the point? Here's the value: the switches are now built all the same – they are just Layer 3 switches. It's the overlay that applies the magic.

IPTel Solutions
Network Engineering Excellence

SDA FABRIC ADVANTAGES

SDA Fabric Advantages

Let's talk more on what the magic is, that the overlay delivers. For the clients, their traffic is sent via the overlay, so the network no longer needs to use all the complicated mechanisms under the surface to support Layer 2, but from a client perspective, it's business as usual.

The key advantage for SDA: **Micro segmentation** - this delivers a very granular set of security controls. Traffic is tagged with Security Group Tags (SGTs) and this traffic passes over the overlay.

For the underlay, ***each switch is Layer 3:*** they're routers with lots of switch ports basically.

The majority of the protocols and issues introduced by all the workarounds at Layer 2 are dispensed with – we only use Layer 2 on one switch (YES – this is the end of Spanning Tree – if you've ever had a Spanning Tree loop, you will know what a celebration this is!)

Layer 2 is essentially stretched between switches, as traffic is switched over the overlay VXLAN tunnels.

Cisco Catalyst Center

For any fabric, you also need a management station, and in the case of Cisco SDA Fabric, this is Cisco Catalyst Center (formerly known as Cisco DNA Center).

For any automated networking deployment, there needs to be an automation and orchestration server - and this is performed for the Cisco Catalyst Center for Cisco.

BGP: (Border Gateway Protocol): Prefix Advertisement

There are 20 or so protocols in use in todays networks – with SDA Fabric we reduce this massively to around 3:

- LISP, a Layer 3 routing protocol,
- VXLAN, a Layer 2 tunnelling protocol, and of course,– BGP.
- BGP, a protocol used to connect the fabric to the outside world

Why BGP?

BGP is used in the same way it is on the internet – as a protocol to advertise prefixes – a handoff protocol between two independent networks, in layman's terms.

Add to that – VXLAN, which does the point-point tunnelling. The tunnels are dynamic and setting these up is a function of SDA Fabric – you don't do this manually.

CAMPUS NETWORKS

Campus Networks: Unique use case

Campus LAN networks have three unique challenges:

- **Three layer network is not always possible:** Sometimes networks are expanded or extended, or just cannot be designed as three layer for some reason
- **Wireless users are highly prevalent in any campus LAN**: They roam between access points and switches
- **Unmanaged Devices:** Many IoT sensors, BMS devices, cameras and so on now connect to corporate networks

For Datacentre networks, we typically have a fairly flat, interconnect-style network, where major services are connected and peered with each other.

Campus networks are different: the key aim is to service the needs of end users, interconnecting the datacentre to the users. The typical campus network is a 3 layer design, featuring Core, Distribution and Access switching. This is not always the case though, as some networks require a divergence from this standard design.

However, if you consider wireless users, they typically will connect to an access point (AP) on the access layer, but via a CAPWAP tunnel, traffic is sent to a WLAN Controller on the distribution layer, de-encapsulated and the traffic then sent out via the distribution layer (unless it was bound for another client).

The other increasing aspect in campus networks are the amount of IOT / BMS / control type devices – these devices do not have an active user connecting to them as they are used to monitor the environment or deliver services to end users (building way-finding, or measuring the temperature and so on).

Many customers, even those that are undertaking a 'wireless first' approach to their network, are discovering that during a campus network refresh they are actually increasing the number of switch ports – all these extra wired IoT devices need to be connected.

The above set of requirements for campus LAN leads us to some questions:

- With ***some*** campus LANs not being built as a three layer architecture, it no longer makes sense to deploy a firewall at the core, as traffic doesn't necessarily flow up and down in this manner any longer (case in point is that some high rise buildings are now putting in place a ring topology on each floor, connected to a collapsed core in the building basement)
- Wireless users could appear on any switch which has an AP connected, alongside any other traffic. ***All this has to be tunnelled back*** to the WLAN controller (there are exceptions to this, such as FlexConnect, but this doesn't scale well in the enterprise space)
- How can we best support this range of devices?
  - Many devices cannot have a authentication supplicant installed as they are black-box sealed units.
  - There are a unique set of security risks as they cannot be secured on the device end - this has to be done by the network.

WIRELESS AND SDA FABRIC

Wireless and SDA Fabric

There are generally three options for the way in which wireless is deployed.

When campus LAN design standards were originally conceived, Wi-Fi wasn't available – so the way in which wireless works is (generally) an overlay, not something that was originally built in to the designs.

There are three broad ways of delivering wireless on a campus LAN:

- **Option 1: Tunnel from the access point to the controller:** This is a CAPWAP tunnel for Cisco. This is by far the most common approach to wireless deployment:
  - This is somewhat suboptimal – all traffic is sent to the WLAN controller, which can be a bottleneck
  - The key point is that the traffic isn't flowing point-point, but tromboning via the WLAN controller
  - Consider how Wi-Fi 6E, with an exponential increase in throughout affects this – you will need larger scale WLAN controllers as well as distribution switches which connect to the controller. Customers are seeing rates in excess of 1.5Gbs - from a busy AP!
  - With the increasing throughout made possible on Wi-Fi 6E, the controllers will just not be able to cope with the amount of throughput
- **Option 2: FlexConnect at the Access layer:** This is the second most common approach. Cisco supports this, but with a limitation on the numbers of APs (300 APs) and Meraki supports this as its native deployment method:
  - The issue with this design type is that for clients to roam effectively, the same VLAN will need to appear across all the access switches to which a particular SSID connects.
  - There are limits to how large a layer 2 domain can grow, so large scalability is the drawback of this methodology
- **Option 3: FlexConnect at the Distribution layer:** In this model, you convert your distribution switches to Layer 3. That way the VLANs from the access layer only need to transit up one level to have traffic switched, not two through to the core:
  - This spreads the routing workload, but still requires the use of Layer 2 VLANs on the access layer

Next Generation Network Architecture has to Resolve these Issues

With SDA Fabric, we introduce a new option, which removes the need for Layer 2 to transit beyond the access layer, and moves the network to a full SDA Fabric deployment:

- **Option 4: SD Access with LISP:** When a receiving switch receives traffic, it notifies the LISP Control Plane Switch (this is a switch which is just a database, detailing what devices are connected where), which is used to allow traffic to be switched directly over VXLAN tunnels from the end point to the destination:
  - Roaming updates when a user moves, update the LISP database and VXLAN only needs to setup tunnels to the switches which are involved in sending or receiving the client traffic

KEY SDA FABRIC TERMS

LISP

LISP (Locator/ID Separation Protocol) is at the heart of SDA Fabric.

This is the mechanism by which all devices report in where they are in the network - this allows the Fabric to set up VXLAN tunnels from their location to their traffic destination. Traffic is then switched over that tunnel, a faster and more efficient process than routing.

LISP was originally a Cisco development, but is now an open standard (RFC9300). The SDA Fabric deployment is managed from Cisco Catalyst Center – it has to be because so much on the network has to be automatic, so a management station is needed.

The main function of LISP is to track where clients are, in order that the Fabric can be setup to allow the VXLAN tunnels to be set up and client traffic switched.

NAC

A Network Access Control (NAC) server is a fundamental aspect of the SDA Fabric story. In Cisco terms, this is Cisco ISE.

This is needed for a number of functions – the first and most obvious one, is to allow authentication of clients. Typically the most secure authentication is via a certificate deployed to the device – this though requires a supplicant to be installed which can interact with ISE.

Many devices fall into the IoT camp though – these devices do not have a supplicant and assume they will just be connected to a network. The NAC component with Cisco ISE can profile devices, so we have a secondary method of connecting devices which don't support a local supplicant.

This is of course less secure than requiring a certificate – and this is where SDA Fabric mitigates the risk by implementing micro-segmentation and only allowing the device specific and defined access.

The NAC function is fundamental to SDA Fabric - you need to positively identify the connecting device to allocate it the access it needs.

SGTs

Security Group Tags are a fundamental aspect of segmentation, used to specify the access for a traffic source and these are attached to the traffic. Essentially the SGTs specify the privilege any traffic has within the Fabric.

VXLAN

VXLAN (Virtual Extensible LAN) is used to create the overlay network for the SDA Fabric.

This essentially introduces tunnels on top of a Layer 3 routed network, to allow traffic to be switched point-point, similar to how a VLAN operates

Thank you for reading our ebook.

CISCO SDA FABRIC

The world of Cisco Software Defined Networking Explained

IPTel Solutions: Your Enterprise Networking and Wi-Fi Experts.

If you're looking for professional assistance with your networking project, look no further than IPTel Solutions.
Our team specialises in enterprise networking and Wi-Fi, and we're here to help.

Contact us today at sales@iptel.com.au or visit our website at www.iptel.com.au for more information.

I you're interested in more great reading, from IPTel, read the ***The Top 8 Wi-Fi Secrets:***

https://blog.iptel.com.au/top-8-secrets-to-great-wi-fi