VERSA NETWORKS

WHITE PAPER

SD-WAN Brownfield Deployment

Introduction

There are multiple considerations to be made by a network architect before transitioning from a legacy network to SD-WAN. For example, is the target architecture a 'Do It Yourself' SD-WAN or is it purchased through a Service Provider or Systems Integrator? Furthermore, within each potential target architecture, there are further considerations to be made. For example, and continuing the scenario above, if its DIY, are the 'head end' components part of Infrastructure as a Service (IaaS) in a Cloud platform like Azure or AWS or deployed within an on-prem Data Centre? Ultimately, those decisions determine the target architecture for an Organisation on the journey to an SD-WAN network.

Most organisations find themselves starting from a 'Brownfield' position. Brownfield typically refers to networks that build on previously deployed infrastructure and features. Very few organisations transition from a position of Greenfield. Greenfield is considered the dichotomy of Brownfield. In Greenfield, these are typically considered ground up projects requiring no integration into existing infrastructure or features. Of the two approaches, Brownfield is not only the most common. From an engineering perspective, it is also the most interesting and challenging. This is because both old and new components need to become well integrated to be considered successful. As such, this Whitepaper shall focus on Brownfield SD-WAN deployments.

When transitioning to any new network including SD-WAN, there are multiple areas that need consideration. For example, training on the new platform and integrating with existing operational tools and processes to name a few. In this Whitepaper, focus is given to Branch related transition considerations. With this in mind, Branch related considerations are broken down into 'Transition Preparation' and 'Day of Transition' activities.

Transition Preparation

Prior to the actual act of physically transitioning a Branch to an SD-WAN network, there's transition preparation that needs to be undertaken. It cannot be underestimated that the time taken for the preparation phase will always be longer than the physical transition itself. If the transition is to go well, whilst delivering a good end user experience, preparation is key.

The following sections identify and discuss some of those Branch level considerations.

The Use of Templates

Any new SD-WAN platform will likely introduce the concept of 'templates'. The idea, as an example, is a 100-site network could be built using a single template. Rollout is therefore simplified as all sites use the same template. This reduces configuration time, errors and omissions as all sites have the same configuration regardless of activation engineer or expertise. This helps reduce transition time and improve the overall experience for the Organisation and specifically for the end users on the day of migration. Post transition, templates also simplify future changes to the network. Using the same 100-site network as an example, if DNS settings for end users need amending, this can be updated in the template. This saves time compared to amending 100 sites individually.

To take advantage of templates, first determine which SD-WAN branches share a common topology (e.g., one WAN or two WAN; one CPE or two CPE etc) and features (e.g., DHCP for clients; 802.1q trunking etc). For example, a fictitious organisation has determined it has four categories of branch sites across its 250-site network. There are data centre sites; large sites (with dual CPE and dual WAN links); medium sites (with single CPE and dual WAN links) and small sites (with single CPE and a single WAN link). Based on this information, the SD-WAN network for this Organization could be rolled out based on four templates – one for each category of branch site. As an example, here's the small site template with a single CPE and WAN:

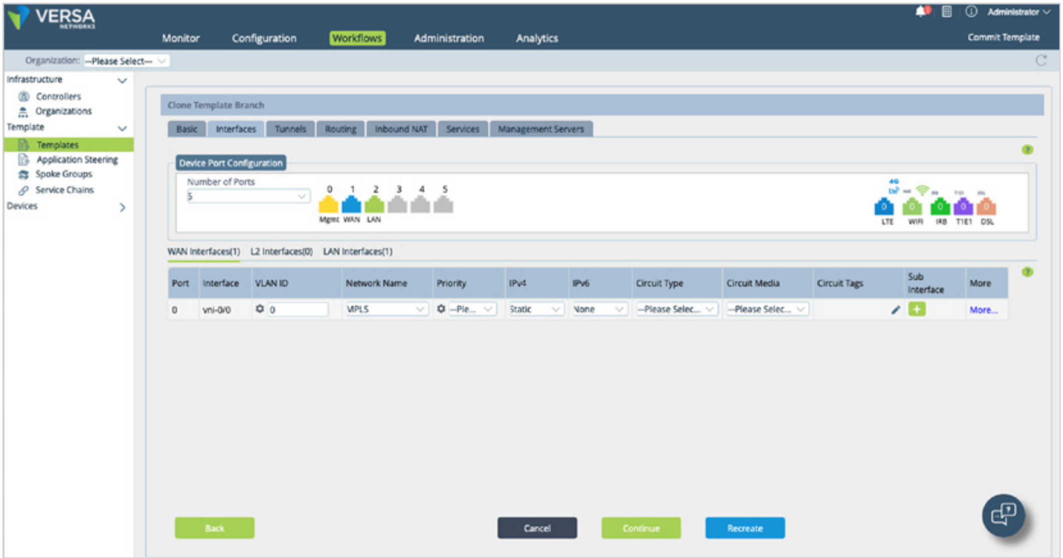

Organization: --Please Select--

Number of Ports: 5

Port	Interface	VLAN ID	Network Name	Priority	IPv4	IPv6	Circuit Type	Circuit Media	Circuit Tags	Sub Interface	More
0	vni-0/0	0	MPLS	--Ple...	Static	None	--Please Selec...	--Please Selec...		+	More...

Even if an organisation can't leverage templates in a physical topology way, templates may still be used for specific 'features' and associated with groups of sites. For example, Firewall policy may be common across the entire estate. In such instances, these can be built into a template and the template associated with all sites:

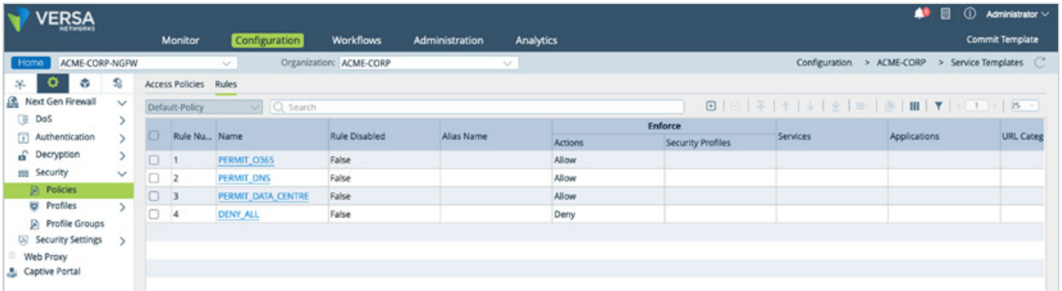

ACME-CORP-NGFW

Organization: ACME-CORP

Default-Policy

□	Rule Nu...	Name	Rule Disabled	Alias Name	Enforce		Services	Applications	URL Categ
					Actions	Security Profiles			
□	1	PERMIT_O365	False		Allow				
□	2	PERMIT_DNS	False		Allow				
□	3	PERMIT_DATA_CENTRE	False		Allow				
□	4	DENY_ALL	False		Deny				

Therefore, templates can be leveraged at a physical, feature or both physical and feature level. Other examples of features that may be common across all or groups of sites include Quality of Service, SD-WAN Traffic Steering and CPE Hardening.

VPN Topology

In simple terms, there are two types of VPN topology: fully meshed or hub and spoke. A fully meshed VPN allows all sites to communicate with each other directly. Consequentially, paths between sites are optimal avoiding unnecessary latency and jitter:

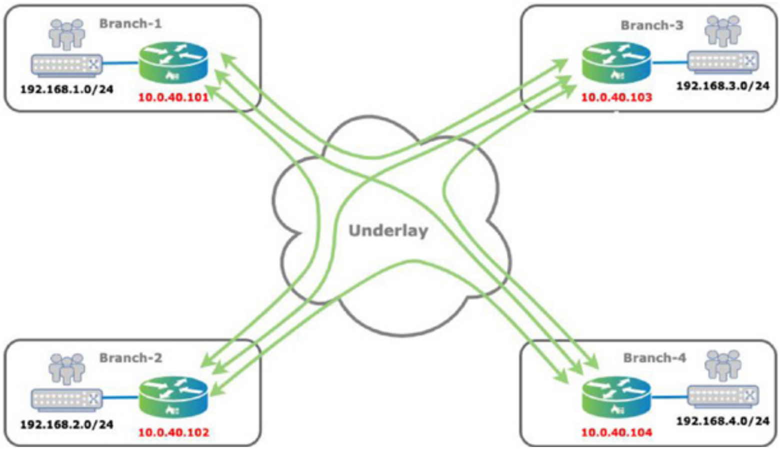

This differs to a hub and spoke topology. Traffic between sites passes through a hub device. As an example, traffic between Spoke 4 and Spoke 6 passes through the Hub:

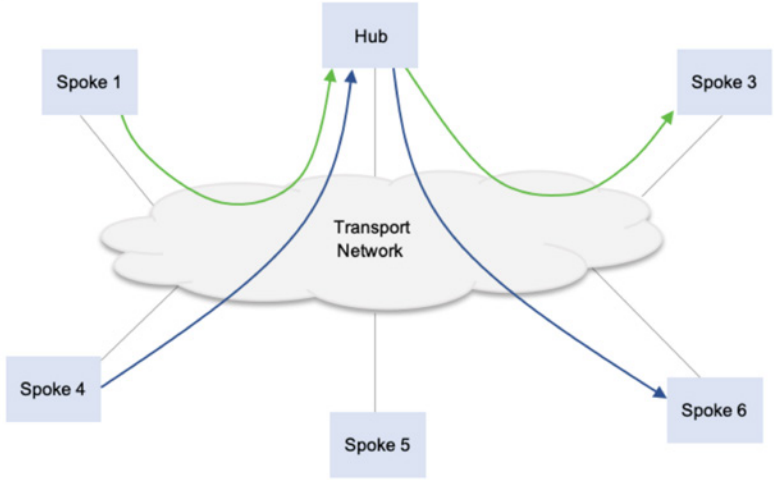

Although this may be sub-optimal from a routing perspective, a hub and spoke topology provides greater scalability and agility. For example, SDWAN networks use probes to determine the performance of the underlay network. These are periodic packets sent between CPE in the VPN. A large, fully meshed VPN composed of low bandwidth circuits can become overwhelmed with probes. A large fully meshed VPN composed of usage-based circuits can incur additional charges due to probe consumption. In contrast, a hub and spoke topology sends probes between hub and spoke – not between the spokes. This reduces bandwidth consumption, thereby giving more bandwidth for customer usage and avoiding unnecessary charges.

A VPN doesn't necessarily need to be a hub and spoke or fully meshed. Instead, it could be a blend of the two. In such cases, 'hybrid' or partial mesh topologies should be considered. Such topologies use a blend of fully meshed sites and hub and spoke sites to create an overall topology that provides advantages over using either of these topologies alone. This is achieved by creating 'Spoke Groups' and associating each Spoke Group with a VPN topology type.

For example, the figure below is composed of a Spoke Group built as a 'Spoke-To-Spoke Direct' VPN (DC Hub 1/2 and Spoke 1, 2 and 3). This allows communication between sites directly. For example, Spoke 1 can communicate directly with Spoke 2. Traffic does not need to traverse the Hub. This topology provides optimal routing between sites. Typically, such a topology is used between major sites in the VPN. In addition, the figure is also composed of a second Spoke Group. This has been built as a 'Spoke to Spoke via Hub' VPN (DC Hub 1/2 and Spoke 4 and 5). This allows communication between Spokes via the Hub. Direct connectivity between Spokes is not permitted. For example, Spoke 4 can communicate with Spoke 5 via DC Hub 1:

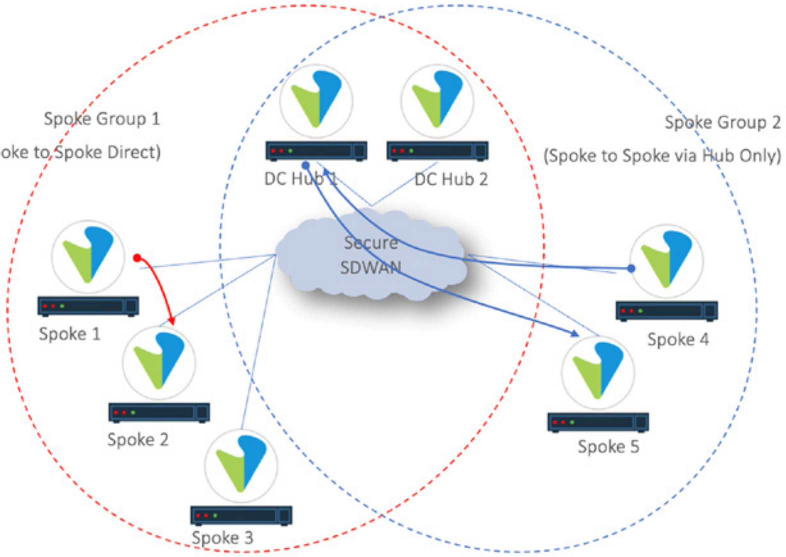

Such a hybrid topology gives the advantages of a fully meshed VPN for those sites connected to Spoke Group 1. (i.e., sites can communicate directly with all other sites, thus optimising paths between sites and limiting additional latency and jitter in the network). At the same time, for sites with small WAN uplinks (such as ADSL) or for usage-based WAN uplinks (such as 5G), these can be placed in Spoke Group 2 using a hub and spoke topology. This limits probe scope to the hub and spoke. This prevents low bandwidth circuits becoming overwhelmed with probes or incurring additional charges due to probe consumption.

Internet Access Policy

One of the advantages of an SD-WAN network is the ability for end users to access the Internet directly from the Branch rather than centrally at a Data Centre or via a Security Services Edge (SSE) service hosted in the Cloud. Such an architecture improves the end user experience. This is because latency between Entity (end user) and Resource (application) is reduced as there are fewer hops to the external Resource. Reductions in latency improve throughput and application 'responsiveness' to the end user.

If the new platform introduces the ability to break out internet destined traffic locally, there are several considerations that need to be made:

- 'which' applications shall be broken out locally at the Branch?
  - Of course, an Organisation could decide to break out all applications locally at the branch. This may be particularly pertinent to Branches enabled with NGFW

features such as URL Filtering; IP Filtering; AV and IPS. Conversely, an Organisation may choose to selectively break out 'trusted' applications locally whilst all other applications are broken out centrally or via an SSE Cloud hosted service.

- 'when' to break out locally?
  - During the transition, does internet access change during the Branch transition or at a later phase? Enabling local breakout later may have some advantages. Certainly, there's fewer moving parts on the day of the transition which may improve migration success. Additionally, with rich analytics and reporting from the new platform, Network Administrators can analyse the network to determine which specific applications to break out.
- 'who' to break out locally?
  - At a coarse level, the Branch may be composed of multiple VRFs. So, the 'which' and 'when' decisions can vary on a per VRF basis. This also extends to the 'who'. As an example, 'guest' users at a Branch may be placed into a guest VRF. Anyone within this VRF can access the Internet directly from the local branch.
  - At a more granular level and with a ZTNA architecture in mind, Entities such as end users may be authenticated onto the network. Moreover, access to Resources – including the internet can then be controlled via firewall policy.

Consideration should also be given to 'whitelisted' IP addresses. Access to specific Resources, such as SaaS applications may be controlled by source IP address. If the architecture model for internet access changes and traffic is permitted to break out locally, the source IP address will change. Therefore, for Resources accessed directly from the branch, ensure new source IP addresses are whitelisted with the respective platform owners. An example of such a use-case can be found by referring to 1.

Firewall Policy

There are two schools of thought on migrating firewall rules during a network transition. One says, "there's plenty of change in the network already, so don't amend firewall rules until after the migration. Post migration, use the rich logging and reporting capabilities of the new platform to clean up policy as appropriate". Another school says, "the transition is the perfect time to clean up policy so only policy that is required is migrated". There's no right or wrong answer to this decision and is likely influenced by multiple factors. For example, if there are transition time pressures whilst at the same time, firewall rules are large and complex, it may make more sense to adopt the 'migrate' approach and then review/refine post transition. Whereas, if the existing policy is simple and up to date, it may make more sense to review and transition to a new policy prior to migration. In either case, policy should always be reviewed and amended throughout the lifecycle of the network – particularly with a ZTNA architecture in mind. It's also worth noting, ZTNA may only be possible post transition. This may be due to technical barriers with the legacy platform. So, it may make more sense to migrate existing rules for the transition and as a Branch is uplifted to a ZTNA posture, policy is amended accordingly.

In addition to the rules themselves, does the decision on local internet breakout as described earlier change any of the existing policy rules? For example, are URL Filtering rules now required? Are other Threat management tools like AV; IPS and IP Filtering also required? And if yes, what is the policy associated with them. For example, is all traffic scanned or just traffic to the internet? This also extends to TLS Decryption. TLS Decryption is essential to the detection and prevention of malicious traffic. But what is the policy associated with TLS Decryption? For example, is policy for TLS to decrypt or inspect for URLs associated with personal information (such as banking and healthcare)?

A minor point, but worth considering is to check if there's an implicit 'deny' rule at the end of the policy ruleset. Typically, this is 'yes', but consideration should be given to configuring an explicit deny rule and enabling logging. This will help to quickly detect end user false positives during the day of transition. This is because logging will now be automatically surfaced into the SD-WAN platform reporting engine.

Quality of Service Policy

Even with an internet underlay, QOS can still be leveraged to control egress traffic from the branch. This allows the transitioning Organisation to prioritise traffic as it leaves the branch. Additionally, other QOS 'levers' such as shaping vs. policing or congestion avoidance vs. tail drop can be implemented to improve the end user experience.

When transitioning, check if there's already a QOS policy in use. Perhaps this can be used as a starting point:

1. How is traffic classified on ingress to the CPE?
    a. Are rules still being 'hit' or do they need updating?
    b. Does the transition allow policy to be updated? For example, the new platform may be able to match on Layer 7 information whereas the outgoing platform was only L3/L4 aware. This may simplify the configured policy and deliver operational efficiencies.
2. How is traffic being scheduled on egress from the CPE?
    a. Is shaping or policing being used and does this need updating? This is especially true where sites are transitioning to faster access technologies such as ADSL migrating to VDSL.
    b. Do bandwidth guarantees during periods of congestion require updating?
    c. Are more/fewer forwarding classes required to meet business outcomes?
    d. If the existing policy uses a tail drop congestion management approach, consider changing this to a congestion avoidance approach such as random early discard (RED). Technologies such as RED attempts to avoid congestion occurring in the first place rather than manage it after its occurred.

If there's no policy, perhaps a simple policy involves three forwarding classes:

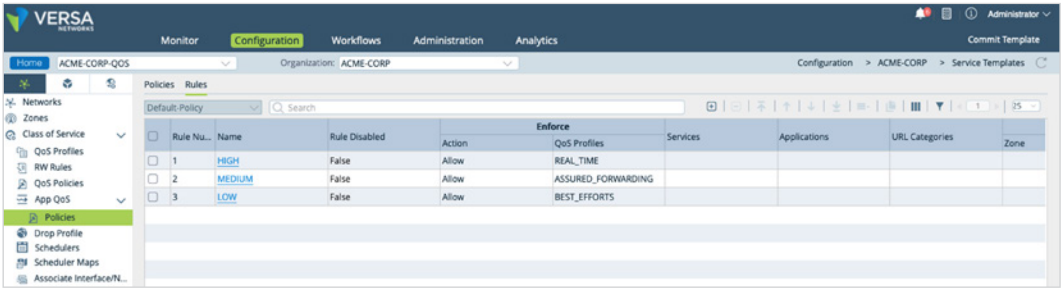

□	Rule Nu...	Name	Rule Disabled	Enforce		Services	Applications	URL Categories	Zone
				Action	QoS Profiles				
□	1	HIGH	False	Allow	REAL_TIME				
□	2	MEDIUM	False	Allow	ASSURED_FORWARDING				
□	3	LOW	False	Allow	BEST_EFFORTS				

The highest priority class is associated with critical applications, such as voice. The lowest priority class is associated with best efforts traffic – such as general internet access. The middle priority class is associated with all other traffic not already captured in the other two classes. Once deployed, the rich analytics and reporting tools of the new platform can then be used to refine which streams of data reside in which of the three forwarding classes. Reporting can also feed into the tuning of the schedulers themselves. For example, perhaps a forwarding class is being starved of bandwidth during periods of congestion. By adjusting the bandwidth guarantee of the class, the end user experience can be improved when accessing business critical resources.

Traffic Steering Policy

Complementary to QOS, SD-WAN traffic steering rules allow traffic between Branches to be forwarded according to the required end user experience. For example, mission critical applications can be forwarded over paths between branches with the lowest latency and lowest packet loss. Furthermore, such traffic could be subjected to packet replication. In this way, if a packet is lost, the stream is still protected as the replicated packet is received by the end device. For the end user, this also means throughput is kept high (as the window size between client and server isn't reduced due to packet loss).

In the same way firewall or QOS policy may already exist in the outgoing network, an existing traffic steering policy may be used as a starting point.

If there's no existing policy, perhaps a simple policy can be built that mirrors the QOS policy:

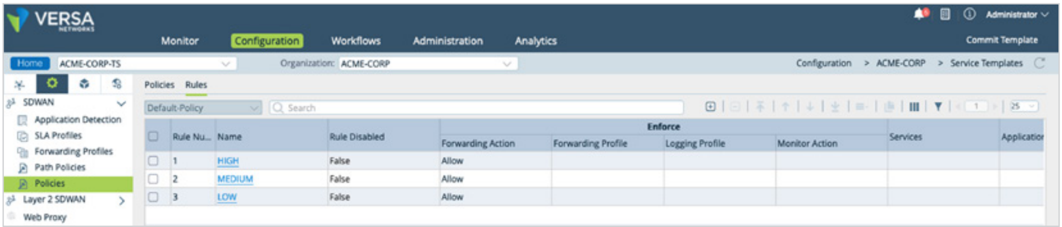

□	Rule Nu...	Name	Rule Disabled	Enforce				Services	Applicatior
				Forwarding Action	Forwarding Profile	Logging Profile	Monitor Action		
□	1	HIGH	False	Allow					
□	2	MEDIUM	False	Allow					
□	3	LOW	False	Allow					

For example, assuming the simple QOS model described earlier is used, for the highest priority QOS class, this can be associated with a Traffic Steering forwarding profile that uses the best path between two branches from a latency, jitter, and packet loss perspective:

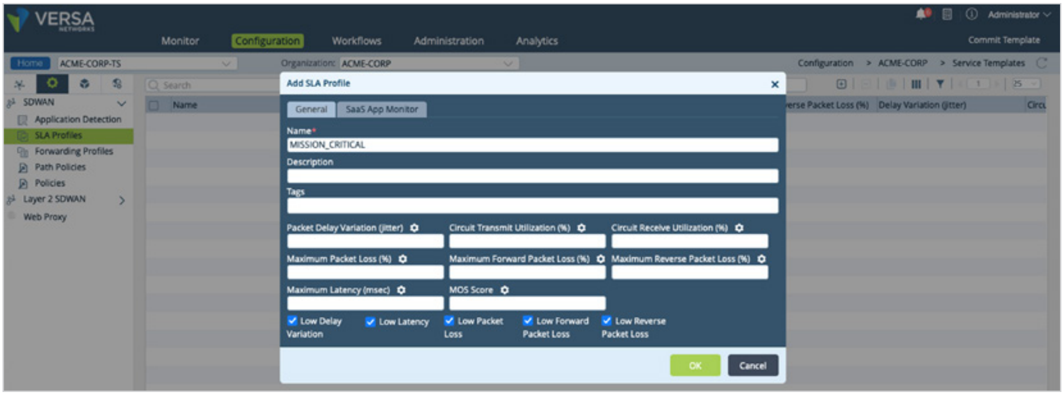

Add SLA Profile

Name* MISSION_CRITICAL

Description

Tags

Packet Delay Variation (Jitter)

Circuit Transmit Utilization (%)

Circuit Receive Utilization (%)

Maximum Packet Loss (%)

Maximum Forward Packet Loss (%)

Maximum Reverse Packet Loss (%)

Maximum Latency (msec)

MOS Score

☑ Low Delay Variation ☑ Low Latency ☑ Low Packet Loss ☑ Low Forward Packet Loss ☑ Low Reverse Packet Loss

The lowest priority QOS class is associated with a forwarding profile that uses any path between two branches so long as it is up. In the event there are multiple paths, traffic is furthermore load balanced across all available paths. The middle priority QOS class is associated with a Traffic Steering forwarding profile that uses the best path from a packet loss perspective (as traffic is less time sensitive so there is no requirement to also monitor latency and jitter). Again, in the event there are multiple paths, traffic is furthermore load balanced across all available paths.

If the new platform supports the concept of a Unified Policy Engine, when traffic is classified by QOS policy into a forwarding class, the Traffic Steering policy can leverage that classification as its own policy 'match' criteria. This makes management of policy simpler as only QOS policies need granular details of the match criteria. Once classified, the Traffic Steering policy simply matches on the forwarding class the traffic is classified as:

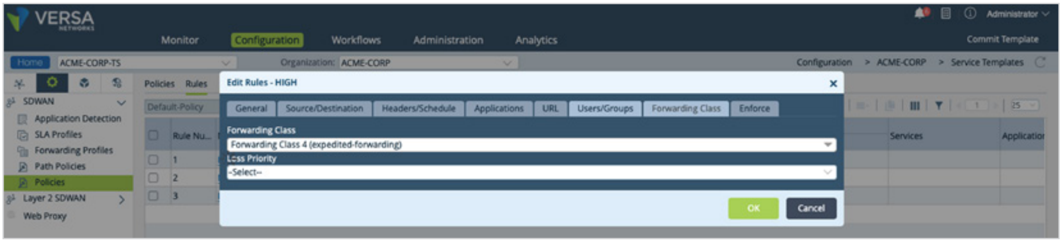

ACME-CORP-TS

Organization: ACME-CORP

Default-Policy

Edit Rules - HIGH

Forwarding Class: Forwarding Class 4 (expedited-forwarding)

Loss Priority: -Select-

As an example, an Organization may have 20 QOS policy rules. Each rule classifies traffic into one of three forwarding classes. Each rule is configured with match criteria including such details as L7 or tuple information. By contrast, this would only require three Traffic Steering policy rules. This is because the Traffic Steering rules match on the forwarding class. There's no requirement to mirror the same 20 QOS policy rules as Traffic Steering policy rules.

Day of Transition

Interconnect

To transition from the legacy network to the new, an 'interconnect' is required. The interconnect connects the networks together. As sites migrate to the new network, they are still able to access internal resources hosted on the legacy network via the interconnect and vice-versa. As an example, the diagram below shows a VOS Appliance. This is connected to both the new SDWAN and legacy networks. This appliance provides the interconnect between User and Applications:

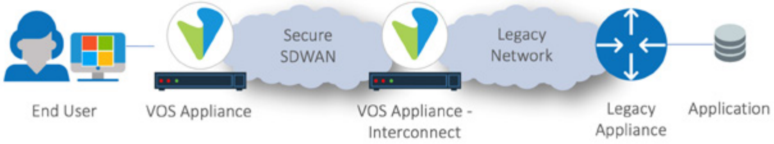

It is recommended a dynamic routing protocol like BGP or OSPF is used over the Interconnect. This simplifies the day of transition activities. For example, when the branch is disconnected from the legacy network, the LAN range(s) associated with the site are no longer advertised into the new network over the Interconnect. Instead, the LAN ranges are advertised from the new network over the interconnect to the legacy network. It also avoids issues when the transitioned site can't access resources on the legacy network because the pre-requisite routing steps were missed! A routing protocol also allows migrated sites to be rolled back to the legacy network without amending the Interconnect.

In summary, using a dynamic routing protocol allows the network to determine for itself where the LAN ranges reside.

Once the interconnect is in place, Branches can be transitioned to the new network whilst still providing connectivity to the legacy network.

Due to the importance of the Interconnect, ensure any SLA wrap is appropriate for the criticality of the data passing over it. For example, a best endeavours SLA is probably inappropriate for an Interconnect when point of sale (PoS) devices on the new network still need to reach into the legacy network.

It is also worth considering that the transition plan is in lock step with the available bandwidth of the Interconnect. For example, a Data Centre site is unlikely to be at the beginning or end of a transition plan. It's more likely to be in the middle. From an interconnect bandwidth perspective, this ensures only half the required bandwidth of the Data Centre is required over the interconnect. Likewise, migrating several large sites may place extra demands on the available bandwidth of the interconnect which may then drive the Data Centre transition before further large sites are migrated.

Minimize the Downtime

Where possible, try to onboard the new branch CPE prior to cutting over the LAN from the legacy CPE to the new CPE. For example, if the Branch has dual WAN links, (as shown in the left-hand diagram below) migrate one of the circuits to the new CPE (as shown in the right-hand diagram below). This allows the new CPE to be onboarded onto the SD-WAN platform.

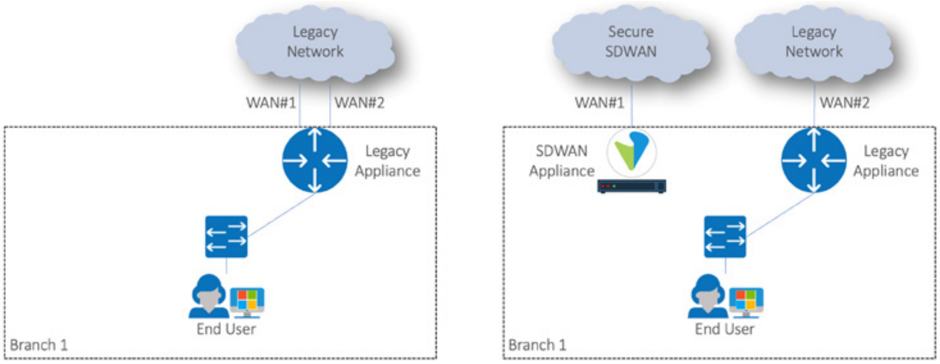

Once the WAN is migrated, check the CPE status. Are all services running correctly? If the new platform uses 'probes' to characterise the performance of the underlay network, check metrics such as latency; jitter; packet loss. This ensures when the LAN is cut over, the WAN is already performing within specification. It's one less thing to check after the LAN is migrated and avoids troubleshooting live performance issues and affecting end users.

Onboarding the CPE prior to LAN cut over also allows the network engineer time to ensure software levels are correct and remediate any patching – especially if reboots are required to complete the software upgrade process. This check should be extended to security packages to ensure the latest software is installed. This ensures more advanced features, like AV scanning, function as expected. Such features are reliant on up-to-date definitions to be effective. If URL or IP Filtering are enabled, are advanced features such as Cloud Lookups functioning as expected?

Check management access to the CPE and any related management functions (e.g., syslog and SNMP) operate as expected. For example, can you remotely access the CPE? Is the CPE synchronised to an NTP clock source – especially important for SAML authentication use-cases. If access to the CPE is remotely authenticated via protocols such as LDAP and SAML, are these successful?

Completing all these steps before the LAN is cut over improves migration success and smooths the end user experience.

If new WAN links are being used, there's even less reason not to onboard the CPE prior to the LAN migration.

Cut over the LAN

You've planned the transition. You've built and deployed all the configuration perhaps using templates to a greater or lesser extent. You've onboarded the CPE and they are ready for use. It's now the final piece of the puzzle – as shown below, move the LAN cable from the legacy CPE on the left-hand side of the diagram below to the new CPE on the right-hand side!

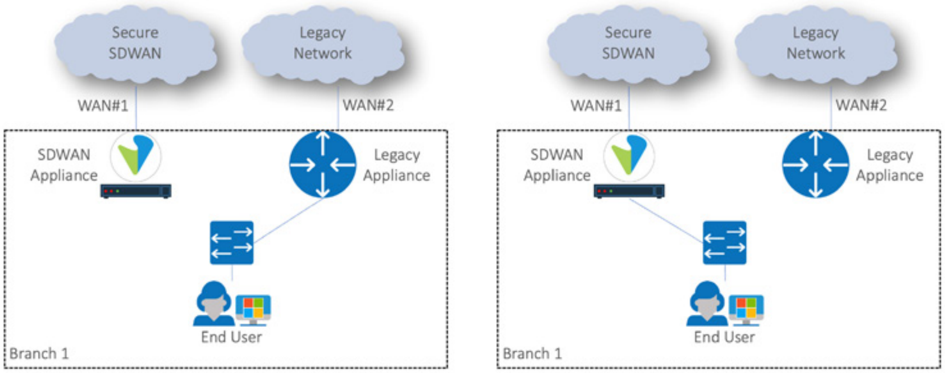

After you've let your site contact know the site is now being migrated and there will be an outage whilst the networks reconverge, it's now time to check the transitioned site is working as expected. When performing these checks, it's always worth listening for negative end user feedback. Such a feedback loop can help fast track problem resolution.

Once the LAN is cut over, the first check is to ensure the LAN interface has come up and just as importantly, at the correct speed and duplex. It seems obvious to check speed or duplex but poor performance feedback from the site may be simply down to speed or duplex mismatches and can be quickly resolved by checking the LAN interface post cut over.

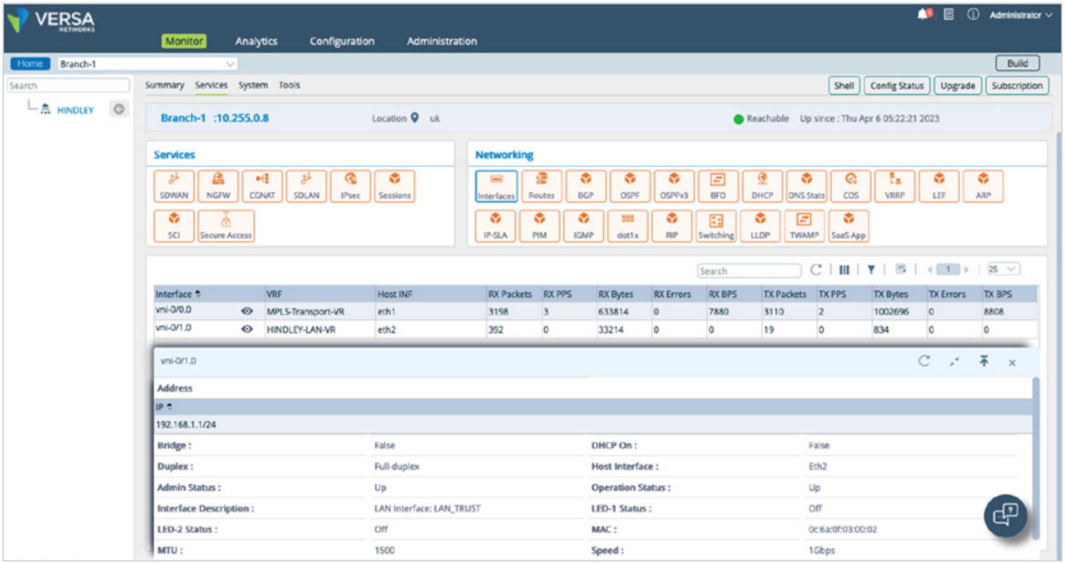

Branch-1 :10.255.0.8 Location uk ● Reachable Up since : Thu Apr 6 05:22:21 2023

Interface	VRF	Host INF	RX Packets	RX PPS	RX Bytes	RX Errors	RX BPS	TX Packets	TX PPS	TX Bytes	TX Errors	TX BPS
vni-0/0.0	MPLS-Transport-VR	eth1	3198	3	633814	0	7880	3110	2	1002696	0	8808
vni-0/1.0	HINDLEY-LAN-VR	eth2	392	0	33214	0	0	19	0	834	0	0

vni-0/1.0

Address

IP
192.168.1.1/24

Bridge :	False	DHCP On :	False
Duplex :	Full-duplex	Host Interface :	Eth2
Admin Status :	Up	Operation Status :	Up
Interface Description :	LAN Interface: LAN_TRUST	LED-1 Status :	Off
LED-2 Status :	Off	MAC :	0c:6a:0f:03:00:02
MTU :	1500	Speed :	1Gbps

Post Transition Checks

After the LAN interface has been checked, the following high-level checks should be undertaken.

It's worth noting, these checks are 'service' orientated. In other words, they are focussed first and foremost on understanding the transitioned end user experience. These checks are not focussed on whether, for example, a route exists in the routing table. Of course, this is still an important tool in the network engineers tool bag. It may also be the reason the end user cannot connect to a Resource in the network. However, as SD-WAN CPE are session and L7 aware, it allows the engineer to move up the OSI stack and from an application perspective view the network from an end user point of view:

1. If the LAN is enabled for DHCP, are end users receiving IP addresses via the DHCP service or relay service running on the CPE.
   a. Don't forget DHCP issues may not manifest themselves straight away. They may only arise for new clients or as existing clients renew their leases. So, keep checking back on DHCP in case issues are masked post LAN cut over.

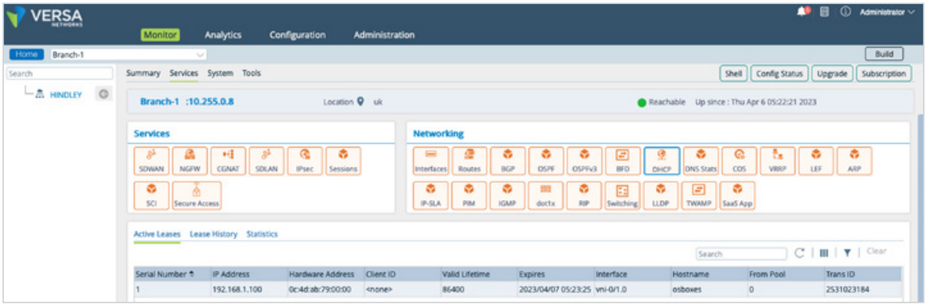

Branch-1

Branch-1 :10.255.0.8 Location uk Reachable Up since : Thu Apr 6 05:22:21 2023

Serial Number	IP Address	Hardware Address	Client ID	Valid Lifetime	Expires	Interface	Hostname	From Pool	Trans ID
1	192.168.1.100	0c:4d:ab:79:00:00	<none>	86400	2023/04/07 05:23:25	vni-0/1.0	osboxes	0	2531023184

2. If end users have an IP address, focus on DNS next. Without DNS, clients can't connect to Resources such as applications.
   a. Whilst troubleshooting DNS issues, if the new CPE is session aware (which is typical of SD-WAN CPE), you'll be able to monitor DNS either by application or protocol (UDP) and port number (53). Be aware it may take both DNS-request and DNS-response for the application to be detected by the new platform. And if DNS is broken, the application may never be detected. Therefore, it is recommended at this stage in the migration window to monitor sessions based on UDP port 53 and not the application itself.
   b. In the session table of the SD-WAN CPE, check for packets received on the LAN interface (i.e., DNS-requests received by the SD-WAN CPE from the end user) and packets forwarded to the LAN interface (i.e., DNS-responses sent to the end user). If the session table only shows packets received from the client, DNS is broken as there are no DNS responses. Check routing to ensure both the DNS server is known to the new network and the transitioned LAN range is known to the legacy network. Additionally, check firewall rules in case traffic is being dropped.

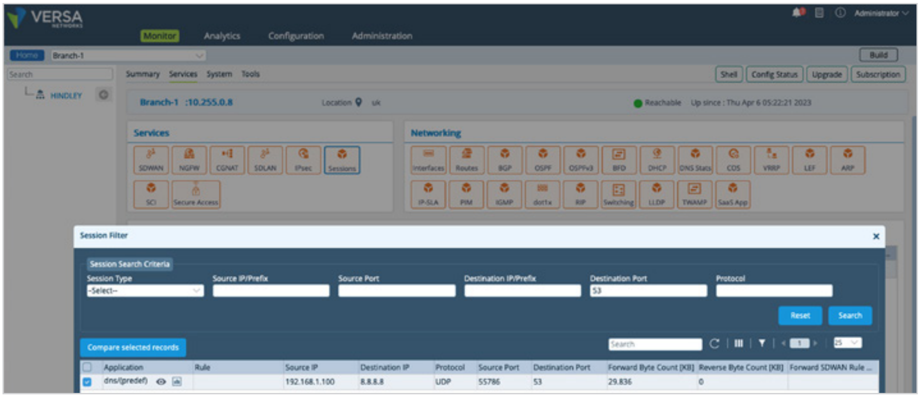

Branch-1 : 10.255.0.8 Location uk Reachable Up since : Thu Apr 6 05:22:21 2023

Session Type: --Select--
Source IP/Prefix:
Source Port:
Destination IP/Prefix:
Destination Port: 53
Protocol:

	Application	Rule	Source IP	Destination IP	Protocol	Source Port	Destination Port	Forward Byte Count [KB]	Reverse Byte Count [KB]	Forward SDWAN Rule ...
☑	dns/(predef)		192.168.1.100	8.8.8.8	UDP	55786	53	29.836	0	

c. If supported by the SD-WAN platform, perform TCPDumps on the LAN interface of the source and destination branch SD-WAN CPE. This will help isolate where packet loss is occurring in the network.

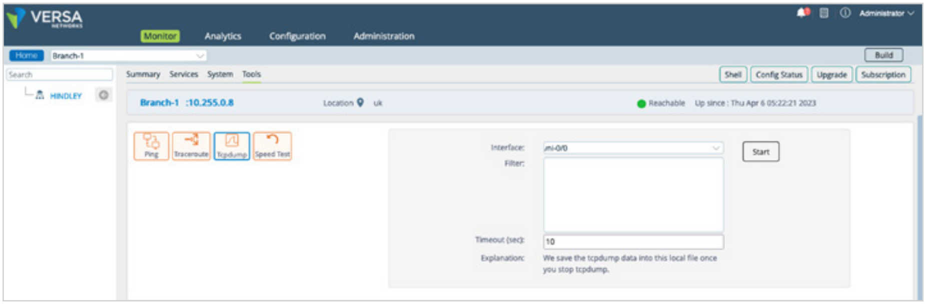

Branch-1 : 10.255.0.8 Location uk Reachable Up since : Thu Apr 6 05:22:21 2023

Interface: vni-0/0
Filter:
Timeout (sec): 10
Explanation: We save the tcpdump data into this local file once you stop tcpdump.

3. Assuming DNS is functioning normally, check internal applications first followed by external applications.

   a. Focus on internal mission critical applications first. Again, check in the session table. Filter on L3/L4 information rather than L7 in case the application isn't detected if its broken.

   b. Check packets are received and sent and troubleshoot accordingly. As per DNS, this could be routing or firewall related. As appropriate, don't forget to use TCPDump on the LAN interface of the source and destination branch SD-WAN CPE

   c. Once internal applications are confirmed working, check external applications. This can take longer to troubleshoot depending on features on the CPE – such as TLS Decryption; URLF; AV; IPS etc. Hence it is recommended to check these last

   d. For external applications breaking out locally at the branch, are the new source IP addresses whitelisted on the SaaS platform?

4. Although real time statistics are a quick way to temperature sense the transition and aid in troubleshooting, once you've got to this step and everything appears to be working normally, it's time to take a step back and look at non real time data. It's time to look for unusual patterns or behaviours in the new platform's dashboards and reports. For example:

   a. Assuming probes are used by the new SD-WAN network, check the performance of the underlay over different time periods (e.g., last 5 minutes; last hour; last x hours). Has it changed from when you last looked during the onboarding period? Any changes in performance may indicate underlay performance issues which manifest themselves in poor end user experience. As examples, xDSL and DOCSIS technologies may continue to train the line post migration. This may result in changes in throughput as well as line drops. LTE technologies may be influenced by environmental conditions (or the cabinet door being closed!!). By checking back on the performance metrics of the underlay, issues can be addressed proactively.

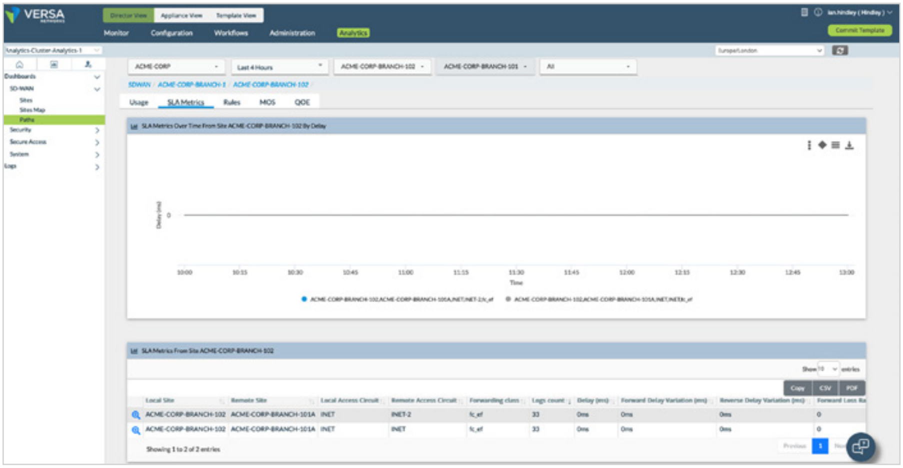

SLA Metrics From Site ACME-CORP-BRANCH-102

Show 10 entries

Copy CSV PDF

Local Site	Remote Site	Local Access Circuit	Remote Access Circuit	Forwarding class	Logs count	Delay (ms)	Forward Delay Variation (ms)	Reverse Delay Variation (ms)	Forward Loss Ra
ACME-CORP-BRANCH-102	ACME-CORP-BRANCH-101A	INET	INET-2	fc_ef	33	0ms	0ms	0ms	0
ACME-CORP-BRANCH-102	ACME-CORP-BRANCH-101A	INET	INET	fc_ef	33	0ms	0ms	0ms	0

Showing 1 to 2 of 2 entries

b. Check WAN utilisation. Is it where it should be in both transmit and receive directions? If it's a branch and there's very little in the receive direction, its likely there's an issue as clients at the branch can't receive data from the data centre.

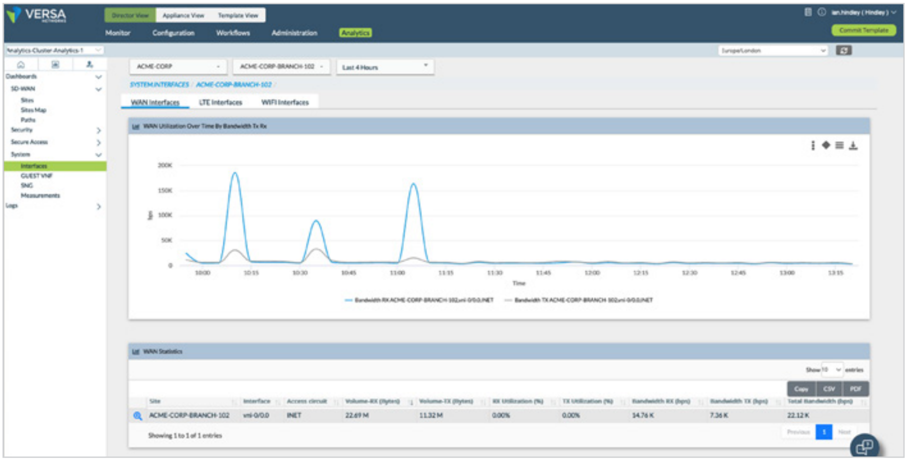

WAN Statistics

Show 10 entries

Copy CSV PDF

Site	Interface	Access circuit	Volume-RX (Bytes)	Volume-TX (Bytes)	RX Utilization (%)	TX Utilization (%)	Bandwidth RX (bps)	Bandwidth TX (bps)	Total Bandwidth (bps)
ACME-CORP-BRANCH-102	vni-0/0.0	INET	22.69 M	11.32 M	0.00%	0.00%	14.76 K	7.36 K	22.12 K

Showing 1 to 1 of 1 entries

c. Check firewall logs. Is traffic being blocked that should be permitted?

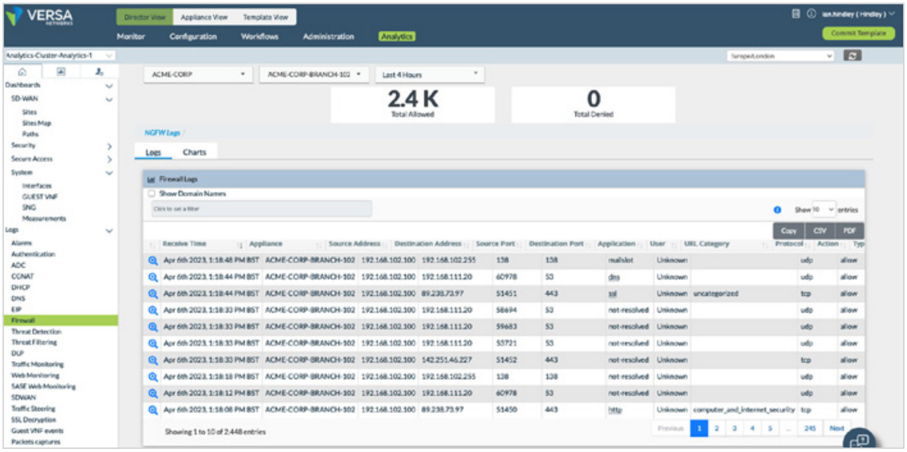

2.4 K Total Allowed

0 Total Denied

Firewall Logs

□ Show Domain Names

Click to set a filter

Show 10 entries

Copy CSV PDF

Receive Time	Appliance	Source Address	Destination Address	Source Port	Destination Port	Application	User	URL Category	Protocol	Action
Apr 6th 2023, 1:18:48 PM BST	ACME-CORP-BRANCH-102	192.168.102.100	192.168.102.255	138	138	mailslot	Unknown		udp	allow
Apr 6th 2023, 1:18:44 PM BST	ACME-CORP-BRANCH-102	192.168.102.100	192.168.111.20	60978	53	dns	Unknown		udp	allow
Apr 6th 2023, 1:18:44 PM BST	ACME-CORP-BRANCH-102	192.168.102.100	89.238.73.97	51451	443	ssl	Unknown	uncategorized	tcp	allow
Apr 6th 2023, 1:18:33 PM BST	ACME-CORP-BRANCH-102	192.168.102.100	192.168.111.20	58694	53	not-resolved	Unknown		udp	allow
Apr 6th 2023, 1:18:33 PM BST	ACME-CORP-BRANCH-102	192.168.102.100	192.168.111.20	59683	53	not-resolved	Unknown		udp	allow
Apr 6th 2023, 1:18:33 PM BST	ACME-CORP-BRANCH-102	192.168.102.100	192.168.111.20	53721	53	not-resolved	Unknown		udp	allow
Apr 6th 2023, 1:18:30 PM BST	ACME-CORP-BRANCH-102	192.168.102.100	142.251.46.227	51452	443	not-resolved	Unknown		tcp	allow
Apr 6th 2023, 1:18:18 PM BST	ACME-CORP-BRANCH-102	192.168.102.100	192.168.102.255	138	138	not-resolved	Unknown		udp	allow
Apr 6th 2023, 1:18:12 PM BST	ACME-CORP-BRANCH-102	192.168.102.100	192.168.111.20	60978	53	not-resolved	Unknown		udp	allow
Apr 6th 2023, 1:18:08 PM BST	ACME-CORP-BRANCH-102	192.168.102.100	89.238.73.97	51450	443	http	Unknown	computer_and_internet_security	tcp	allow

Showing 1 to 10 of 2,448 entries

Previous 1 2 3 4 5 … 245 Next

d. Check QOS statistics. Is traffic correctly hitting policy rules? Is traffic being dropped as the queue depth is too shallow?

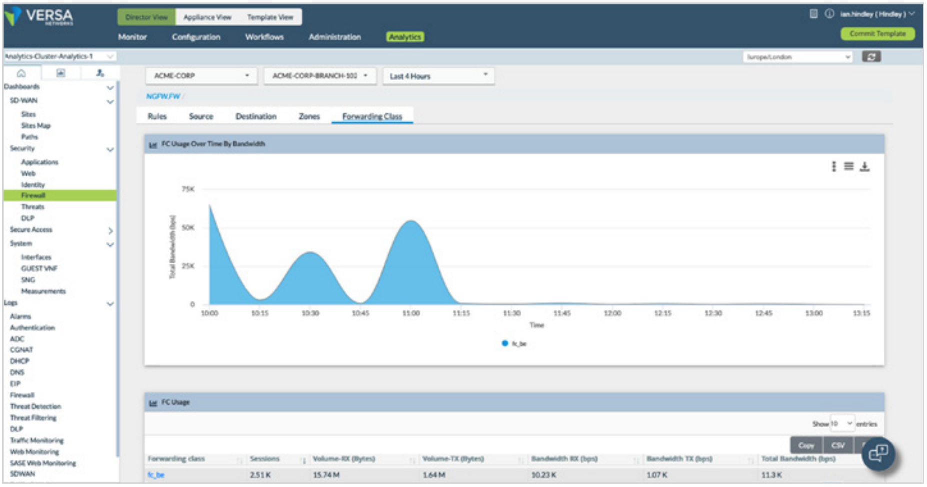

Analytics-Cluster-Analytics-1

Europe/London

ACME-CORP · ACME-CORP-BRANCH-102 · Last 4 Hours

NGFW.FW /

FC Usage

Show 10 entries

Forwarding class	Sessions	Volume-RX (Bytes)	Volume-TX (Bytes)	Bandwidth RX (bps)	Bandwidth TX (bps)	Total Bandwidth (bps)
fc_be	2.51 K	15.74 M	1.64 M	10.23 K	1.07 K	11.3 K

e. Check traffic steering statistics. Is traffic correctly hitting policy rules? Are forwarding profiles correctly using the appropriate paths for the application in question?

Rollback

Any transition should include a rollback plan in case of unforeseen issues. Before executing, ensure diagnostic information is gathered so it may be analysed offline. After you've let your site contact know the site is now being rolled back, the first step is to migrate the LAN back to the legacy router as shown in the diagram below. In this example, the LAN interface is moved from the SDWAN Appliance (on the left of the diagram below) back to the Legacy Appliances (on the right of the diagram):

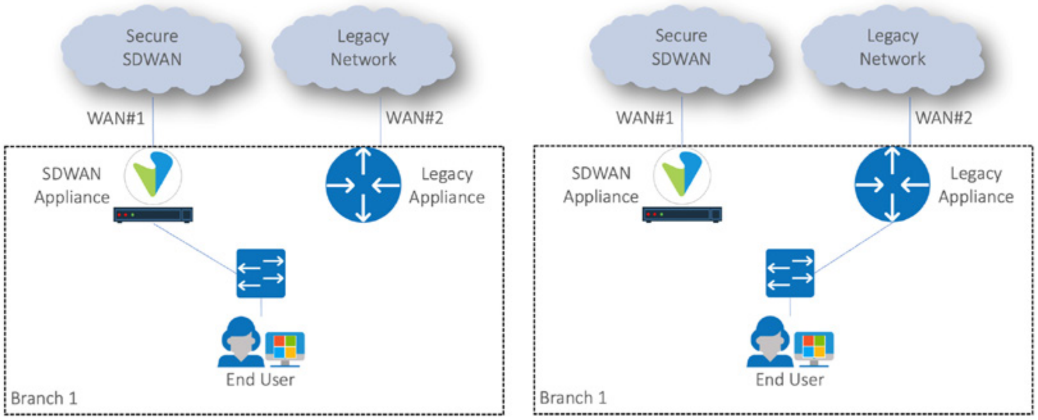

Once moved, there will be an outage whilst the networks reconverge. Once complete, check the site is working as expected. Additionally, if a WAN circuit was moved from the Legacy Appliance to the SDWAN Appliance, this needs to be moved back too (as shown in the figure below):

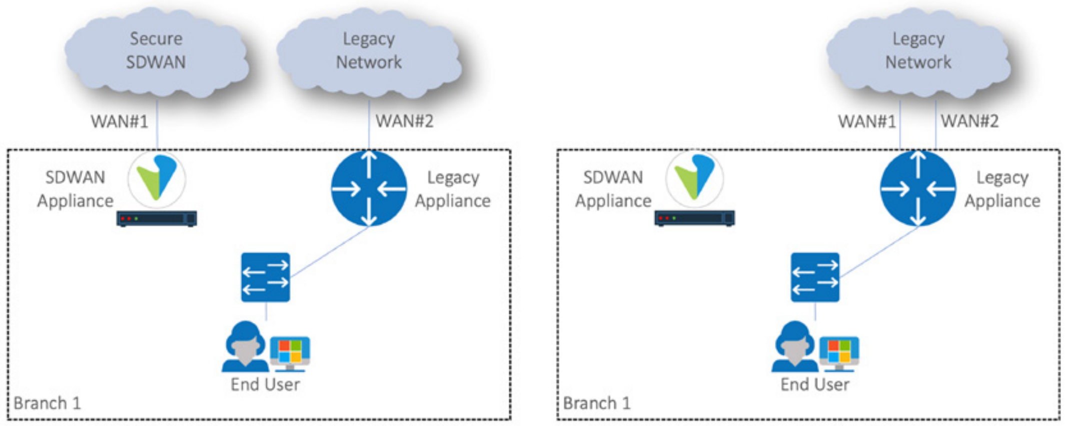

Conclusion

Transitioning from any network is never a trivial task. This is particularly true when it's a brownfield deployment and any unplanned or extended downtime is never welcome.

Unsurprisingly, a successful transition is all about planning and preparation before the day of transition. As examples, consideration needs to be made on the use of templates to simplify rollout of the new network or even to facilitate mass network updates. There are several policy related topics that also need addressing. For example, internet access; Firewall Policy; QOS Policy and Traffic Steering Policy. The challenge is whether to reuse existing policy or recreate from new. In most cases, they should be used as a starting place and refined from there.

Once the planning and preparation is complete, its then down to the day of transition. A prerequisite is for the interconnect to be in place. This ensures connectivity between legacy and new networks is in place and smooths the transition journey for the end user. Assuming this is in place the benefit of onboarding the SD-WAN CPE before the LAN is migrated to the new network cannot be underestimated. This gives time to ensure the new CPE is correctly prepared and functioning before the service is migrated. By resolving potential issues during the onboarding process avoids end users being exposed to issues that negatively impact the transition experience.

Once the LAN is cut over, there are several post migration checks that can be undertaken. In this whitepaper, focus was given to 'service' orientated checks. This document listed these in order of priority. And after the migration is complete, it's important to then step back and look at historical data to spot unusual behaviours or patterns after the site has been migrated.

In the event the transition is rolled back, this can be achieved quickly by moving the LAN cable back to the Legacy Appliance. It is recommended this is undertaken after diagnostic information is gathered.

For more information on Versa Networks, please visit https://versa-networks.com, contact us at https://versa-networks/contact or follow Versa Networks on Twitter @versanetworks

Reference and Resources

[1] Improving the End-User Experience when Using Azure Active Directory and MFA

Versa Networks, Inc, 2550 Great America Way, Suite 350, Santa Clara, CA 95054
+1 408.385.7660 | info@versa-networks.com | www.versa-networks.com