SOLUTION BRIEF

Versa Zero-Trust Architecture Overview

A New Strategy for "Network Security"

Introduction

As incidents of internal compromises have led to cyber security breaches, enterprises are increasingly implementing network admission control at the network edge, typically via 802.1x. Previously, a computer connected to the network was assumed to be trusted; and, thus, less network security was needed. With the increased cyber security breach threats, a connected device can no longer be trusted. This means that all access needs to be authenticated and authorized. The constant threat of cyber security breaches requires a new security paradigm, a Zero-Trust Architecture (Zero Trust Architecture).

What is Zero-Trust Architecture?

There are many definitions of a Zero-Trust Architecture, the best way to explain it is by using 3 principles – authenticated and authorized access only, least privilege (access), and continuous monitoring of access.

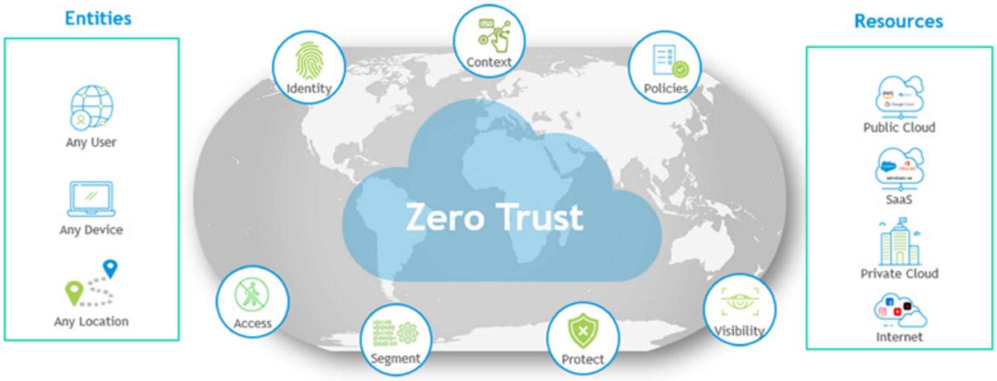

Zero Trust Architecture requires any entity requesting access to a network or resource to be duly authenticated and authorized. Entity is a term that refers to either a human or non-human initiator of the access request. A resource is a data store or application for which an entity requests access.

Furthermore, access and authorization decisions cannot be a single decision made in time and honored in perpetuity. Zero Trust Architecture requires that entities granted access to be monitored for either periodic re-authentication or re-authorization. This is needed due to potential changes in behavior (anomalous access request by the entity), changes in context (mobility, time of access, location of access request), changes in situation (IP reputation changes, URL reputation changes, location reputation changes), changes in the state of client device (ie: its OS or AV software getting out of date) or perhaps due to regulatory requirements.

'Authenticated and Authorized Access Only'

'Authenticated and authorized access only' means that every user, every device, and every application need to be authenticated and authorized before access is given to the enterprise network. A true Zero Trust Architecture does not allow any access to the network without an access control policy decision. While this concept seems like a given for security, many Enterprise networks are built upon the concept that access internal to the enterprise does not need to be authenticated and authorized with such level of details. So, any guest, contractor, or threat actor that has gained access to the enterprise premise or to Enterprise WLAN can attach to the Enterprise network and have implicit trust applied.

The identity and access management policy of a zero-trust architecture provides the mechanism to authenticate and authorize all entities. The identity and access management policy integrates with the enterprise's choice of identity provider. How the entity is authenticated depends on the type of entity being authenticated. For human entities, typically username and complex password is the norm for the credentials utilized in authentication. For application entities, typically a certificate is utilized. Device entities offer the widest variety of authentication methods as some allow for certificates or agents, and some offer multi-level authentication (human entity on a laptop), but IOT devices offer the most complex strategies as most do not allow for certificate nor agent installation. However, in a Zero Trust Architecture, every device, every application, and every user need to be authenticated to assure that they can be trusted. While most legacy security postures look only at the entity requesting access to a resource and not the details and security posture of the resource, Zero Trust Architecture looks at all entities and resources connected to the network. Both entities and resources need to be authenticated and authorized. Once the entity is authenticated then the identity and access management policy will authorize the entity's access. Authorization differs from authentication by not just looking at the validity of the credentials (certificates, username/password, or device profile) but considering context and situation to determine if a particular entity in a particular set of circumstances should be allowed to access a particular segment of the network and/or to a particular set of resources.

Least Privilege

Least privilege is defined as only allowing permissions or privileges which are needed to perform a specific task or role. The old paradigm of VPN allowing ubiquitous access to the network no longer fits. Zero Trust Architecture maintains that access to the network should be for legitimate authorized purposes only. The use of policies in Zero Trust Architecture is very powerful as it determines which access is granted, under what conditions it is granted, and for how long it is granted. This enables the enterprise to control how long an entity can be on the network, what actions an entity can perform while on the network, and under what conditions information can be uploaded, downloaded, modified, or created.

Continuous Monitoring

Continuous monitoring is the last critical component of Zero Trust Architecture. By continually monitoring all access, the Zero Trust Architecture can protect against the lateral or tangential access by an entity. With the ever-present threat of a cyber security breach, any device or application can become compromised. With 34 percent[1] of successful network breaches being started from inside the enterprise, no longer is the threat to the enterprise only coming from external factors. Being an employee, contractor, or vendor of an enterprise is not sufficient. From accidental disclosure, or intentional abuse by disgruntled employees, to unauthorized access and disclosure of sensitive enterprise resources, enterprises now needed to find ways to control lateral movement of data, even within the company.

Zero Trust Architecture requires that every access attempt to an enterprise resource must be authenticated, authorized with least privilege, and monitored. This concept applies to access attempts by employees, customers, contractors, partners, and especially any threat actor, access requests from inside the company or outside the company, and to access requests from anywhere in the world.

Many Zero Trust Architecture solutions incorporate an agent to help with the device and user authentication. An agent enables the solution to gather information that will be part of the Endpoint Information Profile (EIP). The agent, which typically comes with a vender trusted certificate, is used to establish the initial trust mechanism. Thus, anyone accessing the device could have an authentication challenge and access can be permitted or denied by policy on the device itself. The use of an agent has limited scope in the context of Zero Trust Architecture. In a closed system, where only employees access the network, this model is very effective. However, most enterprises have a myriad of people accessing the enterprise resources – employees, partners, contractors, and in many cases, guests (customers, potential customers, auditors, etc.). Contractors and partners have relationships with many different enterprises. There is a high probability that the agent installed for one enterprise might conflict with an agent installed for another enterprise. Partners, contractors and employees with personal devices may not adopt a given enterprise agent. Guests might be enticed to temporarily utilize an agent that can easily be downloaded and removed. But many guests do not want to be bothered installing a 3rd party software. In addition, many enterprises have IoT devices that will not support 3rd party software.

Therefore, any Zero Trust Architecture Solution needs to incorporate an agentless mechanism for human and non-human entity authentication. Since an agent-based solution provides a litany of information that is beneficial to the authorization decision, a hybrid approach, where an agent-based mechanism is utilized for enterprise employees and devices (where an agent can be installed) and an agentless mechanism is utilized for all other human and non-human entities, is needed. This model offers the most flexibility and the most control.

Versa Networks Zero-Trust Architecture

Versa Networks' Zero Trust Architecture solution is incorporated into the Versa Networks SASE Solution. This architecture utilizes a perfect marriage between networking and security. Following diagram depicts a high-level Versa SASE reference architecture view. The Versa Networks Zero-Trust Architecture aligns with both the NIST SP 800-207 Zero Trust Architecture[2] publication and the MEF 118 Zero Trust Framework for MEF Services[3] standard.

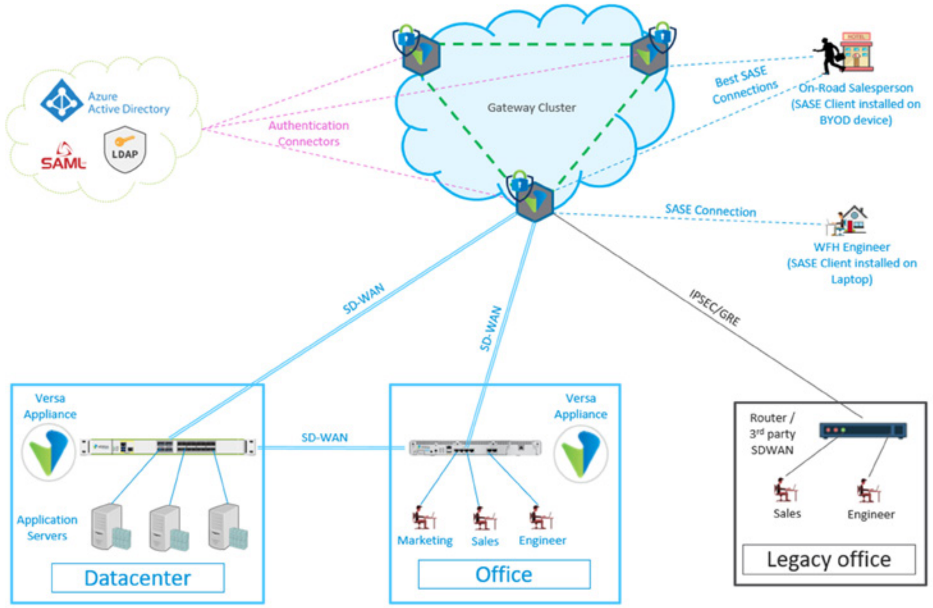

Versa SASE architecture components

- Secure SD-WAN – Connection between offices and workers (Edge) protection
- Zero-trust Network Access (ZTNA) – Secure access to company assets and cloud-based applications or data for all entities
- Secure Web Gateway (SWG) – Secure Internet access for all entities including those on-premises or off-premises.
- Versa Director/Concerto – Centralized Network and Security Policies
- SASE Agent – Agent-based software

Main Versa SASE benefits

- Zero Touch Provisioning for SD-WAN and SASE Clients
- Integrated SD-WAN and SD-Security in a single platform
- Single-Pass architecture – unparalleled performance at scale
- Unified security policy management – one software stack (VOS™) for on-premises, SASE cloud and the edge
- SD-WAN Lite for the SASE Clients
  - › Performance based Gateway selection and Traffic Steering
  - › Application/FQDN-based Traffic Steering
  - › Multi-Access Link support on the SASE Client
  - › Forward Error Correction (FEC) on the SASE Client
- SD-WAN Lite between SASE Gateway and 3rd-party branch routers using IPSec, eBGP and TWAMP
- Optimal Traffic Steering and SaaS acceleration using Versa Traffic Engineered SASE backbone
- Support for multi-cloud deployments
- Big Data Analytics with AI/ML functions
- Hierarchical multi-tenancy & granular RBAC – SPs can conserve costs with full separation of roles
- Dynamic tenant & virtual Gateway instantiation – elastic auto-scaling and network intelligence to meet real-time capacity demands
- Flexible deployment options – cloud delivered, on-prem model, DYI

One of the core principles of the SASE architecture is to secure communication between resources, branches and entities.

Main principles of the Versa ZTNA access

- Zero trust for entities and resources.
- External and Internal threats always exist on the network.
- Every device, user, application, and network flow must be authenticated and authorized.
- Granular Access: Access to the resources should be provided not only based on the source IP/entity but also based on the Layer4-Layer7 Application, URL category and reputation, and deep packet inspection of the application payload.
- The security posture and geo-location of the entity also must be taken into consideration.

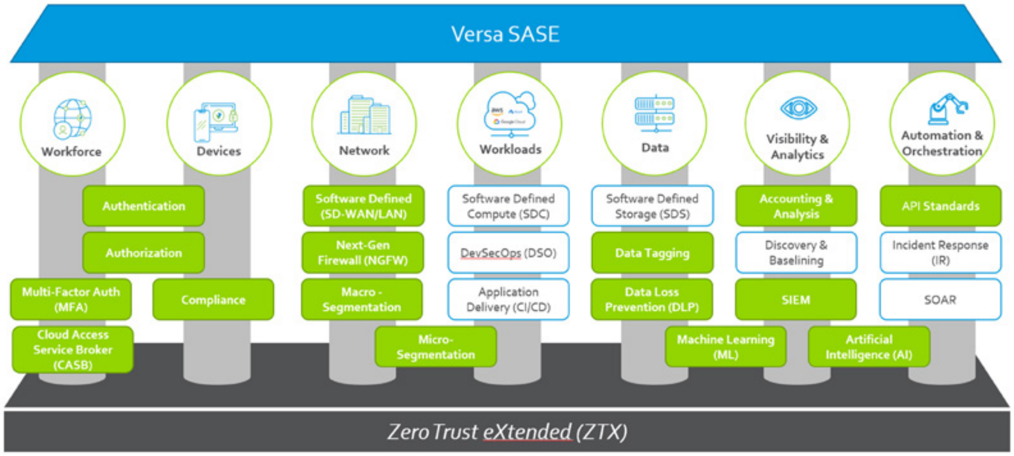

The Versa SASE architecture utilizes authentication connectors to exchange authentication information between the Cloud Gateways and the customer's Identity Provider (IdP). There are multiple types of authentication connectors supported by the Versa SASE architecture:

- LDAP
- Active Directory
- Azure Active Directory SAML
- SAML
- RADIUS
- Local Database managed by Versa SASE Service

Versa Cloud Gateways

Versa Cloud Gateways are the central component of the SASE architecture. Cloud Gateways enable entities to have a secure and robust access to the internet, corporate applications and an optimized experience using a global SASE fabric of Versa Cloud Gateways. Versa Cloud Gateways perform the ZTNA, SWG, CASB, DLP functions and optimally steer traffic to SaaS and internet destinations.

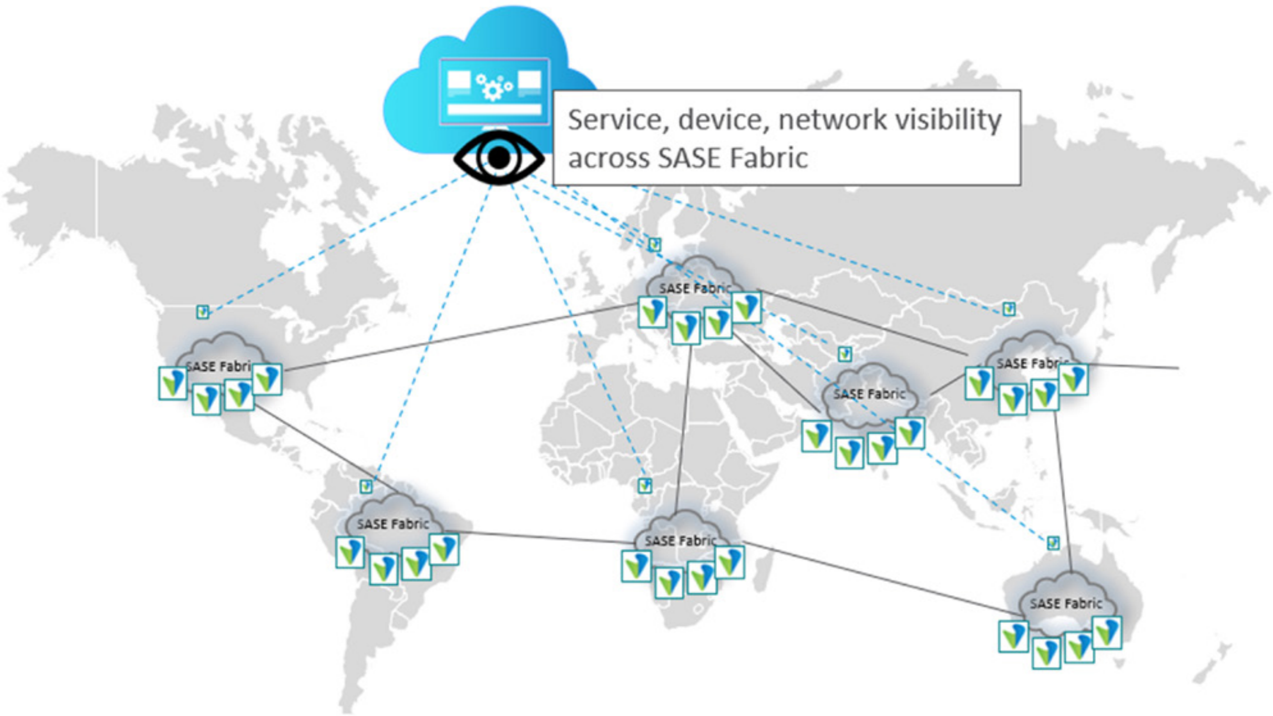

The main benefit for the enterprise is that the cloud-delivered model does not require an enterprise to host security components on-prem. Versa Cloud Gateways provide high-available environment with automatic scale-out capabilities.

Customers can connect to the Versa Cloud Gateways in any of the following ways:

- Legacy branches which need additional security
- Third-party SDWAN branches which need security features
- Versa SD-WAN devices to extend customer's internal network
- Home offices
- Roaming entities
- Bring Your Own Devices (BYOD)

Enterprise locations connect to the Versa Cloud Gateways to provide a comprehensive set of security functions. By connecting to the Versa Cloud Gateways, many security functions that would have been done at the edge can be offloaded to the cloud. Reducing the need for large and expensive appliances at the Enterprise network edge. This architecture allows the network to scale better and provides a more centralized security management and analytics.

Any Legacy branch or non-Versa SD-WAN can be connected to the Versa Cloud Gateways for protection and traffic optimization purposes. This connection can be established using IPSec or GRE tunnels.

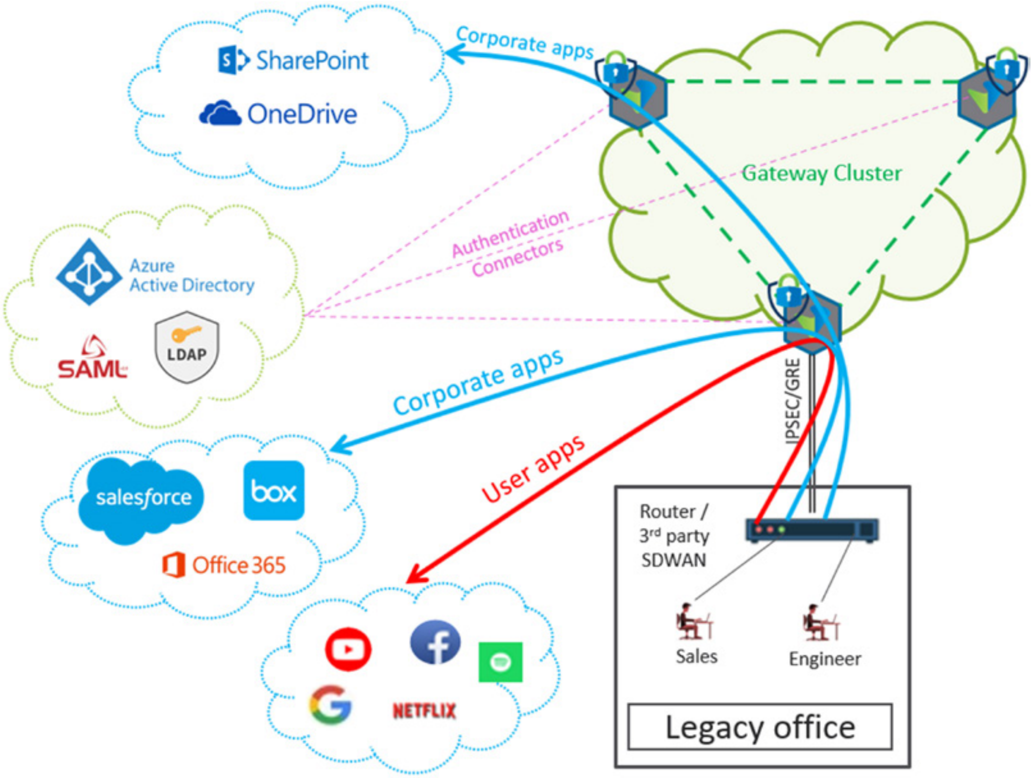

As demonstrated in the above diagram, the enterprise receives additional security benefits and application performance by steering all the traffic to SaaS applications and Internet-Destinations via the Versa SASE Fabric.

If an enterprise has already deployed SD-WAN service from Versa, then that enterprise can extend its existing SD-WAN infrastructure to the Cloud Gateways. In addition to all the benefits of the Versa Secure SD-WAN such as, Forward Error Correction (FEC) and Packet Replication, the enterprise benefits from centralized security policies across the enterprise environment. The Versa appliances at the edge of the enterprise connect via SD-WAN to the Versa Cloud Gateways. There is no need to manually provision IPsec tunnels. Also, the traffic from the enterprise edge is routed to the optimal Versa Cloud Gateway using an underlay which offers the best application SLA.

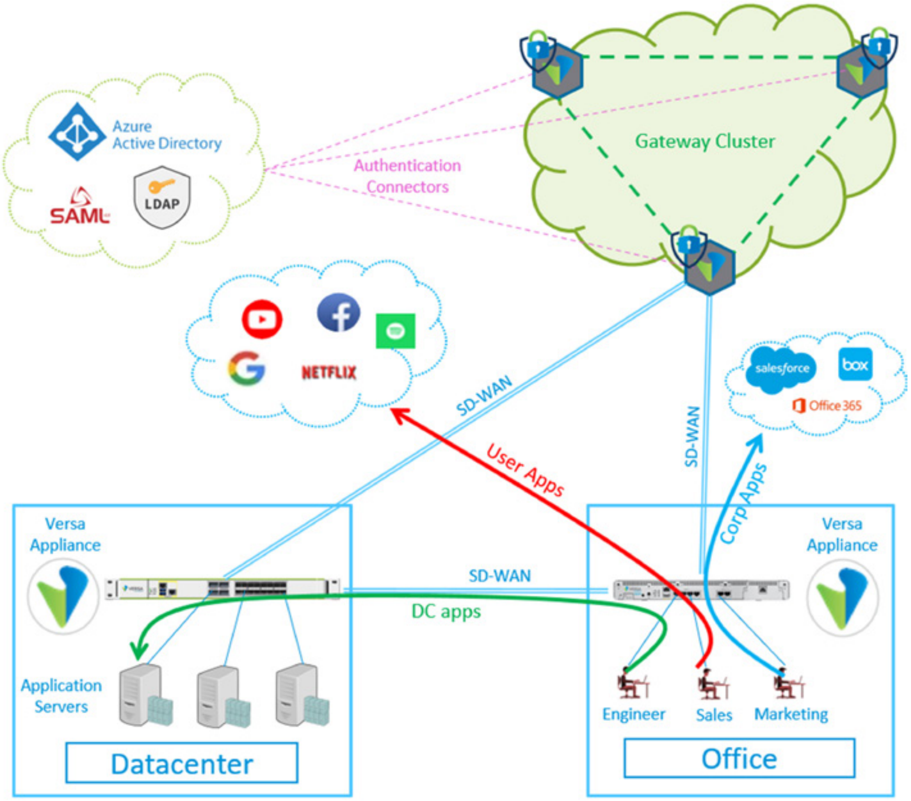

Versa SASE Client

The Versa SASE Solution extends the ZTNA concept by providing a client agent which can be installed on the enterprise device. The agent is known as the Versa SASE Client and allows remote entities to securely and optimally connect to their corporate resources (in Private DCs or Public Cloud Service Providers), SaaS applications.

The Versa SASE Client has advanced features, including SD-WAN Lite, best gateway selection, selective Direct Internet Access (DIA), visibility about the security posture of the endpoint and others.

Reference architecture for the remote client connections:

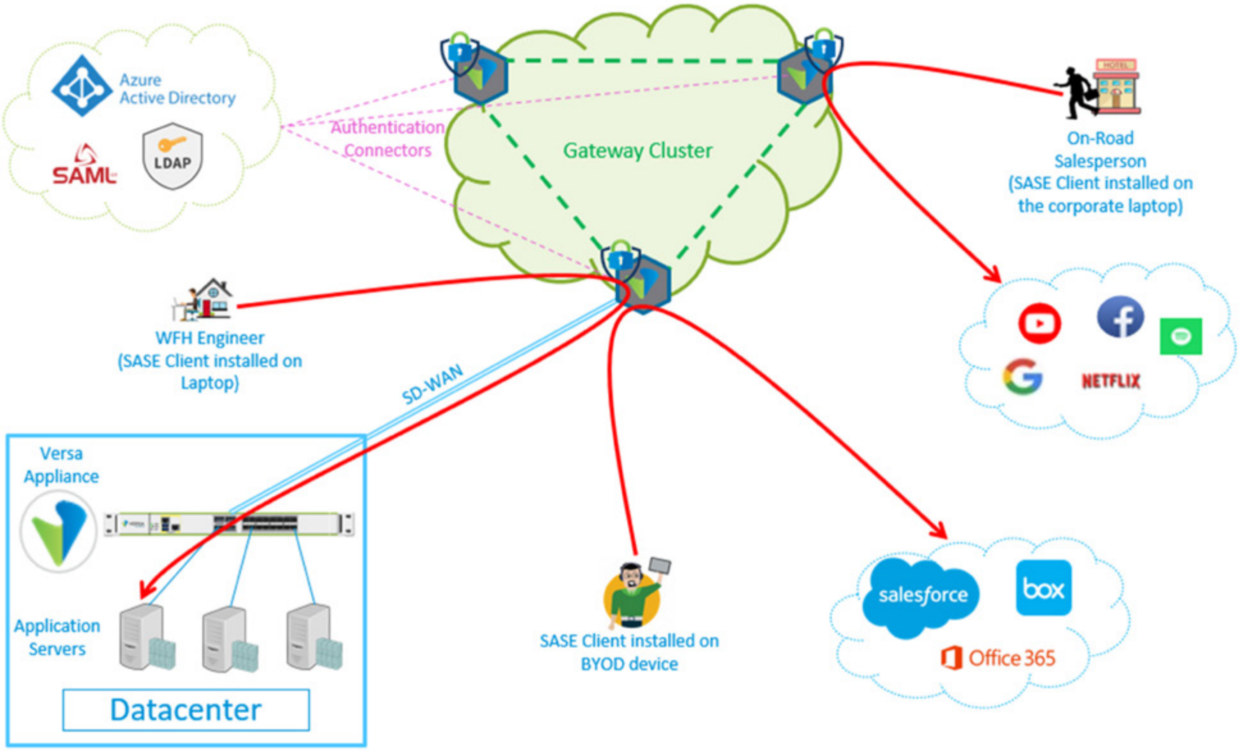
Azure Active Directory
SAML
LDAP
Authentication Connectors
Gateway Cluster
On-Road Salesperson (SASE Client installed on the corporate laptop)
NETFLIX
WFH Engineer (SASE Client installed on Laptop)
SD-WAN
Versa Appliance
Application Servers
Datacenter
SASE Client installed on BYOD device
salesforce
box
Office 365

Versa Secure Gateways can be grouped per-geographic region, per-continent or dedicated per-customer. Each entity can be assigned to one of the gateway groups at the time of registration:

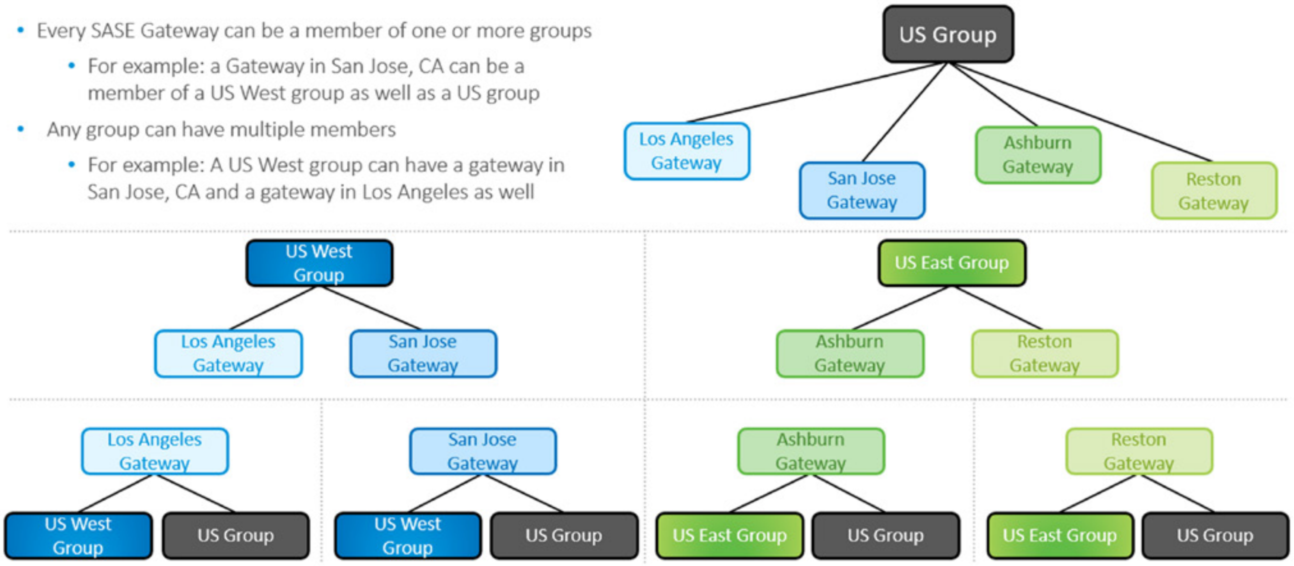
- Every SASE Gateway can be a member of one or more groups
  - For example: a Gateway in San Jose, CA can be a member of a US West group as well as a US group
- Any group can have multiple members
  - For example: A US West group can have a gateway in San Jose, CA and a gateway in Los Angeles as well

Depending on company policy, an Enterprise-Admin can enable entities to connect to a set of gateways or to any of the gateways based on the best performance metrics measured from the SASE client.

SASE Client enables roaming entities to automatically connect to the best gateway, which is determined based on several factors, some of which are given below:

- Physical Proximity – by identifying network location based on IP subnet.
- Network Proximity – automatic gateway selection based on the round-trip delay.
- Real-Time load metrics of the gateways – selection of the least loaded gateway to load-balance the load.

The SASE client develops an endpoint information profile that allows for the Zero Trust Architecture policy to permit or deny access based upon the following parameters:

- Anti-virus version
- Anti-virus signature version
- Operating System type
- Operating System version
- Operating System Patch version
- Specific software installation and version
- Storage encryption
- Other parameters

Versa Agentless SASE

However, for those entities who cannot install or do not want the Versa SASE Client, the Versa SASE Architecture allows for the following mechanisms:

- Captive Portal – utilizing 802.1x for the on-premises devices.
- Identity Provider redirection – Utilizes a login page for remote access.
- Enterprise provided front-end web-application redirection.
- Fully qualified domain name registration – special URL for remote access which is directed to Versa Cloud Gateways.
- Device identification and Fingerprinting based access control.

Just like the SASE Client model, the entities utilizing a non-agent-based access method need to register and get enterprise approval for access to the network.

Conclusion

In conclusion, adoption of Zero-Trust Architecture can significantly improve the enterprise security model. Versa Networks' SASE solution enables the enterprise to implement Zero-Trust Architecture, either with or without a client agent, by enforcing access authentication and authorization, granting only the access that is necessary, and by continuously monitoring that access to assure compliance corporate security policies. Versa Networks has market leading Zero Trust Architecture solutions that are cloud and on-premises delivered to address the security and networking needs of our customers comprehensively. For more information on Versa Networks' Zero Trust Architecture solutions, please visit https://versa-networks.com, contact us at https://versa-networks.com/contact, or follow Versa Networks on Twitter @versanetworks.

References and Resources

[1] Verizon, 2019 Data Breach Investigations Report, 2019, https://www.verizon.com/business/resources/reports/2019-data-breach-investigations-report.pdf

[2] NIST, SP 800-207, August 2020, https://csrc.nist.gov/publications/detail/sp/800-207/final

[3] MEF, MEF 118 Zero-Trust Framework for MEF Services, October 2022, https://www.mef.net/resources/mef-118-zero-trust-framework-for-mef-services/