WHITE PAPER

VeloRAIN: Next-Gen Security for AI-Driven Networks

Table of contents

Impact of AI workloads on WAN infrastructure

AI-driven applications are disrupting long-established upstream/downstream edge traffic ratios, their traffic is encrypted and therefore hard to identify for analysis and prioritization, and agentic AI workloads are sensitive to latency and dropped packets. In VeloCloud's State of The Enterprise Edge report, many organizations stated they are initiating "AI networking" infrastructure refresh projects to better support these unique AI workload requirements outside the datacenter.

The nature of AI applications and their workloads are fundamentally altering edge networking. Agentic AI traffic patterns have the characteristics of being highly sensitive to latency, bursty, and often operate on a peer-to-peer basis. These characteristics affect the evolution of AI workloads.

Latency sensitivity

While most network administrators are used to latency challenges (jitter, packet loss, etc.) causing video streaming issues, latencies can derail AI workload processing. AI workloads are distributed in nature as data is processed in parallel across clusters of workloads. Latency delays across these architectures can break distributed AI application processing, leading to inaccurate training and unreliable inferences.

Bursty traffic

AI applications demand large scale peer-to-peer connectivity and have bursty traffic patterns, demanding a dynamic WAN solution. The unpredictable nature of AI traffic requires networks to adapt and respond to these bursty patterns dynamically through intelligent traffic steering and shaping. Traffic steering should be pinned to individual application flows rather than the underlay to ensure that dynamic network adjustments are aligned with the unique requirements of AI traffic bursts.

Traffic symmetry

Most current enterprise WAN architecture is built to serve web application traffic. Web application traffic focuses on sending large volumes of data downstream to users and receiving very little traffic in return. AI training and inference traffic break this model, because upload and download traffic are comparable. AI workloads have low tolerance for delayed uploads. Many WAN architectures and Internet service providers need intelligent overlay to help manage traffic parity across both directions.

The challenge extends beyond developing superior algorithms for network management outside the data center. It also involves adapting to the unique requirements of these AI-driven applications through an intelligent software layer. This layer must not only understand the applications' needs but also dynamically adjust network resources to prioritize critical applications, ensuring optimal quality of service while deprioritizing less-critical traffic.

Security challenges with AI networking

While AI based systems promise significant advantages in efficiency, automation and decision-making, they also introduce new security challenges. The agentic nature of these AI systems, combined with their reliance on large datasets, dynamic learning processes, and real-time decision making, opens new vectors for potential attacks, threats and vulnerabilities. This white paper explores the security concerns associated with AI networking, presents key challenges in securing AI-driven networks, and discusses potential solutions and best practices to mitigate these risks.

Security for AI involves the ability to protect AI workloads along with the ability to secure unsanctioned use of AI apps which could lead to potential security threats. While the traditional security landscape helps with conventional WAN infrastructure workloads, they often need additional layers of defense to protect evolving GenAI applications. Security solutions for AI networking need to address new attack surface challenges including:

- **Expanding attack surfaces**: While the traditional zero-trust framework focuses on protecting users, devices, applications and cloud access, the agentic architecture of AI workloads adds a new dimension to the attack surface. Zero trust needs to evolve to support segmented access of users/devices/applications access to large language models (LLMs).

- **DOS attacks, adversarial attacks and data poisoning**: AI workloads running on edges can be manipulated through adversarial attacks, where malicious inputs can be introduced to mislead or deceive the model. These attacks can cause the AI system to make incorrect decisions, potentially jeopardizing the security and performance of the network. Additionally, AI agents running on edges can be denied access to resources through denial-of-service attacks that prevent them from operating.
- **AI system vulnerabilities**: GenAI models are complex and opaque, which makes them prone to errors and vulnerabilities. These vulnerabilities can be exploited by attackers to compromise AI-based network operations or manipulate the system in subtle ways.
- **Shadow AI:** Similar to shadow IT, shadow AI refers to the usage of AI tools unsanctioned by an organization's IT department. These tools might range from AI-powered chatbots, data analysis platforms, or machine learning models used to process business data. While the tools might help an employee be productive, inadvertent exposure of confidential information could result in severe damages to a company.
- **Privacy violations**: Agentic models may have sensitive data, such as personal information. If unsecured, this could lead to data leakage or misuse, becoming a privacy violation. Zero trust does not automatically ensure compliance with data protection regulations, so privacy-focused controls must be integrated into the architecture.

Challenges with traditional threat detection methods

Traditional threat detection methods rely on signature-based approaches, which focus on identifying known threats based on their patterns or signatures. However, this approach has several limitations:

- **Inability to detect unknown threats**: Signature-based detection methods are ineffective against zero-day attacks, which exploit previously unknown vulnerabilities.
- **Limited scalability**: Traditional threat detection methods can become overwhelmed by the sheer volume of network traffic, leading to false positives and decreased accuracy.
- **Lack of contextual awareness**: These methods often lack context about the network environment, making it challenging to accurately identify legitimate traffic.

VeloRAIN: A robust, secure SD-WAN AI solution

VeloRAIN is an AI networking architecture that enhances the performance, security, and scalability of distributed AI workloads beyond traditional data centers. Short for VeloCloud Robust Artificial Intelligence Networking, VeloRAIN builds on the leading VeloCloud SD-WAN platform to optimize AI workloads across distributed enterprise networks. It enhances VeloCloud Dynamic Multipath Optimization™ (DMPO) with AI and introduces Dynamic Application-Based Slicing (DABS) to assure quality of experience (QoE) across multiple networks.

With VeloRAIN, organizations can use AI apps more easily with visibility, app prioritization, and automation that continually improves quality of experience (QoE). The VeloRAIN architecture enhances VeloCloud's DMPO network intelligence to accommodate the evolving AI workload demands efficiently. This ensures that enterprises can use AI capabilities effectively, without compromising on network performance or user experience.

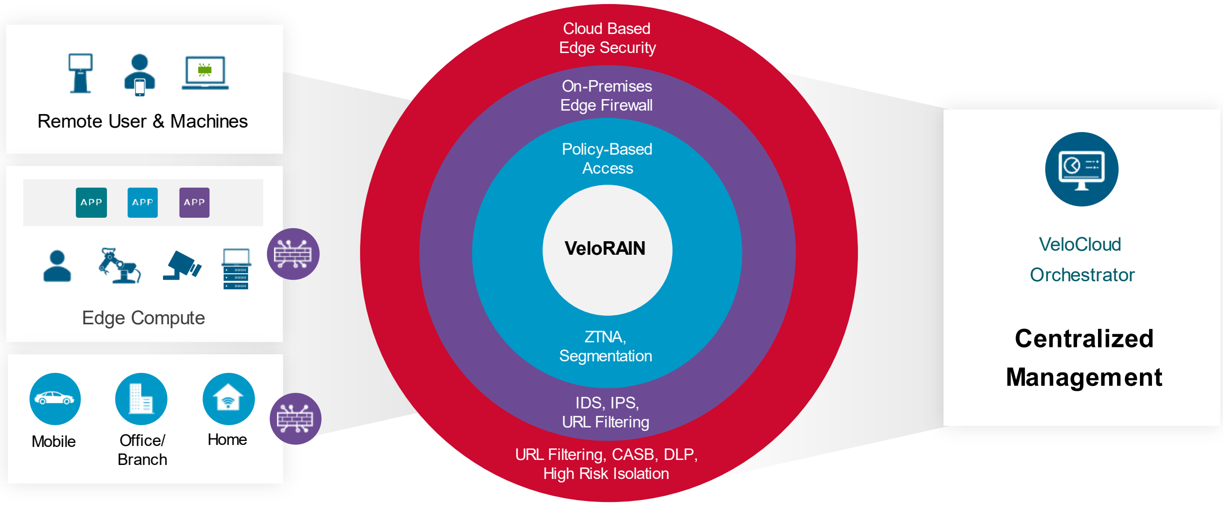

Figure 1: VeloRAIN solution framework reduces the attack surface for GenAI workloads, with centralized policy management and distributed enforcement at the edge and in the cloud

Core pillars of VeloRAIN architecture

Four critical aspects of the VeloRAIN architecture set it apart in optimizing AI applications.

AI-driven application profiling

VeloRAIN identifies and prioritizes applications with new intelligent capabilities. By identifying applications accurately through an enhanced machine learning (ML) system, VeloRAIN ensures that each AI and non-AI app receives appropriate network resources. VeloRAIN will also be able to identify encrypted application traffic that was previously unreadable for network optimization solutions. Organizations will be able to tailor network performance to meet the specific needs of or critical applications.

AI-based network optimization

VeloRAIN introduces Dynamic Application-Based Slicing (DABS), an innovative approach focusing on the application layer instead of traditional network-based slicing. This assures quality of experience (QoE) per application across multiple disparate underlying networks, whether they support network layer slicing or not. AI workloads receive the necessary bandwidth and low latency required for optimal performance, while less critical applications can be de-prioritized. DABS also includes user profiles to prioritize traffic based on a user's identity and attributes, helping ensure that key users receive the QoE they need.

AI also augments VeloCloud's DMPO (AI DMPO). Feeding packet capture data into an LLM, and then training that LLM to perform packet capture data, will allow end users to ask natural-language questions such as "Why is the network running slowly?" and receive an immediate response.

AIOps with real-time data

With data from the vast VeloCloud deployment base, VeloRAIN employs AI to automate network operations (AIOps), using anonymized real-time data to dynamically adjust policies. For example, in retail stores, vision applications analyzing in-store behavior could benefit from real-time policy adjustments, ensuring seamless experiences and high-value customer retention. This dynamic approach allows enterprises to adjust network configurations on-the-fly, ensuring the best possible experience for users and applications.

AI/ML based security

VeloRAIN's architecture is built on a robust foundation of AI/ML algorithms that analyze network traffic patterns to identify potential security threats. The platform uses machine learning models to learn normal behavior patterns and detect anomalies in real-time. This approach enables VeloRAIN to adapt to emerging threats, ensuring the security of the network.

Machine learning in threat detection

Traditional branch security systems rely on predefined signatures and rules to identify known threats. Threats introduced by adversarial AI require an ML-powered, real-time protection-based approach. The Intrusion Detection and Prevention System (IDPS) solution on VeloCloud SD-WAN Edge uses machine-learning based methodologies to generate signatures for unknown threats by continuously learning from network data. New intelligent capabilities in VeloRAIN enable the detection of encrypted application traffic that was previously unparsable, providing visibility into previously hidden threats. This helps identify and block AI-generated attacks.

Anomaly detection and predictive analytics

VeloRAIN's architectural framework will include machine learning algorithms used to analyze network traffic patterns, identifying normal behavior and flagging deviations. By analyzing historical data and real-time network activity, VeloRAIN can predict potential security risks and alert administrators to take proactive measures. Additionally, it includes behavioral analytics models, powered by machine learning, that can distinguish between normal user or device behavior and malicious activity. For example, the solution can identify unusual network traffic, or abnormal traffic destinations, providing early warning signals of potential breaches.

Symantec SSE AI

On the cloud front, Symantec SSE for VeloCloud includes an advanced set of capabilities, powered by AI/ML technologies. Advanced Machine Learning is part of the Symantec engine that is used to determine if a file is malicious or not with high accuracy and low false positives. The machine learning engine includes a model that is trained weekly on Symantec's Threat Labs against a large collection of known good/bad files, which in turn allows any new file presented to be classified as good or malicious based on its characteristics. For data protection, Symantec's DLP engine uses machine learning to protect unseen data without explicit instructions/rules. This is achieved by training the system for positive and negative data.

VeloCloud solutions

The growth in AI applications underscores the importance of responsive and adaptive network technologies. The VeloRAIN architecture is at the heart of VeloCloud's strategy, designed to not only recognize changes within network traffic but also to respond dynamically to these changes. VeloRAIN underpins the VeloCloud family of security solutions including VeloCloud Secure SD-WAN, VeloCloud SASE, secured by Symantec and VeloCloud SD-Access, for the best performance, effective security and enhanced user experiences across our customers' networks. Each of these solutions help address the AI-related security challenges outlined earlier in this document.

Security considerations for AI applications

When selecting an enterprise AI security solution, there are three primary considerations for enterprises:

- Inspection of user/agentic traffic
- Security of the enforcement points
- Security of the overall solution

VeloCloud is integrating new AI capabilities across its portfolio of products to secure enterprises' AI driven applications.

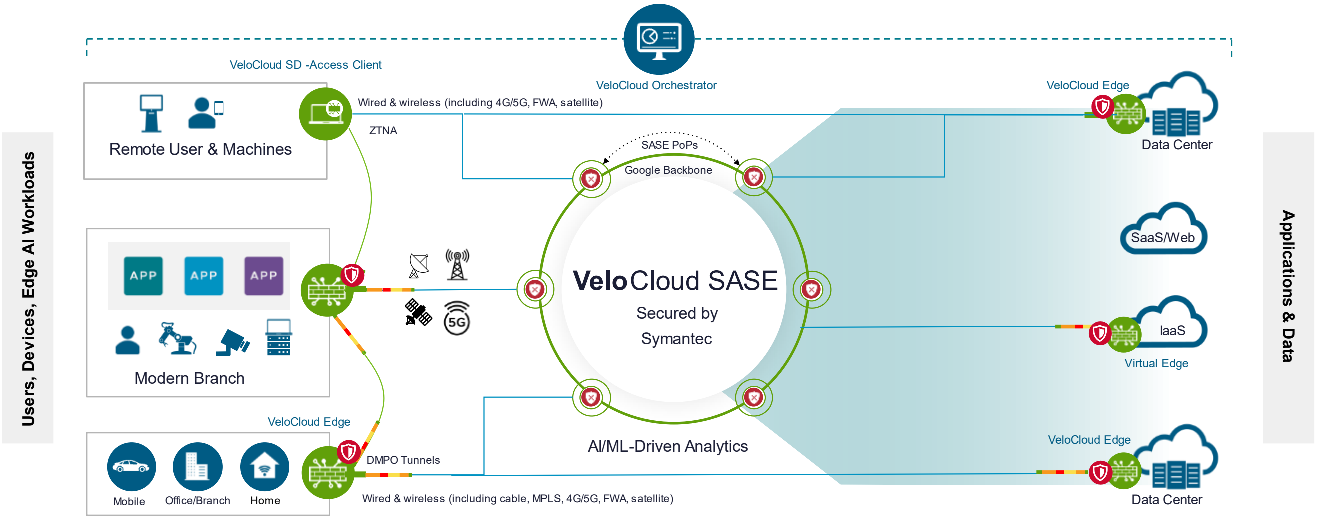

Figure 2: VeloCloud ecosystem with security enforcement points

Starting with the VeloCloud Edge, security is built in at the network via strong encryption methods for the link, network segmentation to isolate various types of network traffic and prevent leaks, advanced inspection of user traffic through built-in enhanced firewall services, and cloud enforcement using Symantec web services solutions. The architectural details of the security approach are discussed in detail in the following sections.

VeloCloud secure SD-WAN

Security is fundamental to VeloCloud SD-WAN, which is built on an architecture that ensures secure communication between the management, control, and data planes. In Figure 1, the management plane consists of the VeloCloud Orchestrator and the control plane consists of the VeloCloud Gateway. In the VeloCloud SD-WAN hosted SD-WAN deployment, the Gateway has a dual function: it can optionally participate in the data plane when the Gateway is set to process Internet-bound traffic to take advantage of DMPO capabilities.

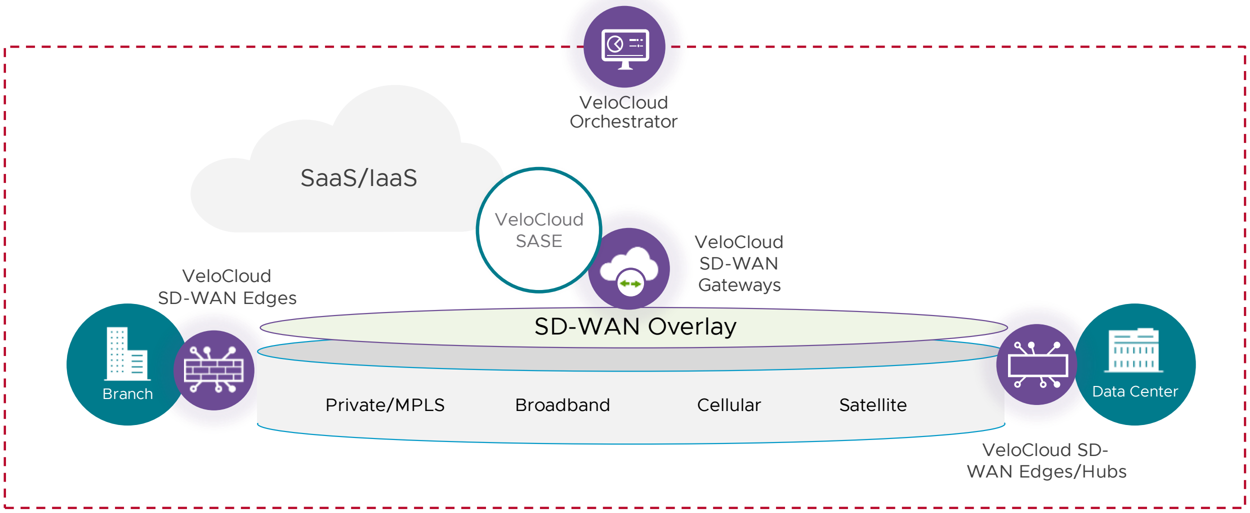

Figure 3: Architecture view of VeloCloud SD-WAN components

Inspecting user traffic

Multiple options exist within VeloCloud SD-WAN inspecting user traffic as it passes through the Edge appliance. Inspection for the purpose of threat detection and prevention (IDS/IPS, anti-malware, URL filtering etc.) is available either locally in the Edge or on cloud by steering user traffic to a cloud-hosted security service by policy.

Encryption between VeloCloud SD-WAN components

One of the critical components in securing AI-driven systems, particularly those that rely on distributed architectures or cloud-based infrastructures, is network link encryption. Encryption of network traffic provides a fundamental layer of protection against unauthorized access, data breaches, or cyberattacks. VeloCloud solutions use industry-standard encryption protocols, such as SSL/TLS and IPSec, to secure data transmissions over WANs.

Public key infrastructure-based authentication secures communication between VeloCloud SD-WAN data plane endpoints, namely the Edge and Gateway. VeloCloud Orchestrator acts as a central certificate authority (CA) server and manages the public key infrastructure (PKI) lifecycle of the VeloCloud SD-WAN deployment including certificate signing requests (CSRs) from Edges and Gateways, verifying certificate authenticity, and maintaining the certificate revocation list (CRL). VeloCloud Orchestrator has a built-in CA server. However, it also interoperates with external certificate authority for agencies that deploy an on-premises Orchestrator and must use their own CA rather than the default self-signed Orchestrator CA.

Security between SD-WAN components on the data plane is accomplished through Diffie-Hellman (DH) key exchange to generate symmetric keys between the VeloCloud SD-WAN data plane endpoints (e.g., VeloCloud SD-WAN Edges, Gateways). This key is then used to start encrypting the data using supported algorithms, ensuring the traffic going through the tunnel between the data plane endpoints is secured. Traffic between Edges, and from the Edge and Gateway, is secured with an VCMP-over-IPsec tunnel. IPsec uses AES-128 or AES-256-bit keys for confidentiality and up to SHA-256 for data integrity protection. Encrypted traffic is encapsulated in VCMP and UDP using source and destination port UDP 2426.

The need for segmentation

Network segmentation allows an organization to compartmentalize application traffic in the network infrastructure, separating users and assets belonging to different security zones. Some of the common use cases for segmentation include separating guest traffic from corporate traffic, and isolating credit card transactions from all other traffic types. With AI applications, segmentation can ensure LLM models do not talk to each other unless configured for security reasons.

In traditional networks, segmentation requires complex command line interface (CLI)-based configuration or dedicated hardware. More importantly, administration of the VLANs and virtual routing and forwarding (VRFs) increases the workload of the network engineer significantly. The biggest drawback of traditional network segmentation techniques is they are locally significant to the device, and the network engineer must provision every network device in the path, hop-by-hop, to extend the segmentation end-to-end. This results in network engineers having to spend significant effort to ensure segmentation is implemented correctly and policy compliance is met.

Below are a few common use cases for network segmentation:

- Line-of-business separation by department for security and audit.
- User data separation: guest traffic, payment card industry (PCI) data, employee traffic.
- Mergers and acquisitions: network segmentation allows overlapping IP addresses and secure access to shared assets.
- An external party such as a partner needing access to a subset of corporate data.

VeloCloud SD-WAN provides edge-to-edge and edge-to-gateway segmentation that can be centrally provisioned and pushed to some or all SD-WAN branch sites, making it scalable, easy to manage, and cost-effective. Each segment is treated as a separate configuration entity having its own set of cloud VPN, business policy, firewall, and QoS configuration elements.

VeloCloud segmentation provides network isolation using a VRF-like concept to protect sensitive traffic from potential threats. Segmentation is achieved by associating a segment ID with every packet entering the SD-WAN overlay fabric. DMPO isolates different segments as they traverse the overlay. The SD-WAN Edges use the segment IDs to route traffic only to their associated segment. Segmentation is assigned by interfaces and sub-interfaces (802.1q Virtual Local Area Network [VLAN]). Each segment has a unique route table entirely separated from the rest of the network. Edge-to-edge segmentation is centrally provisioned and pushed to some or all SD-WAN Edges, making it scalable, easy to manage, and cost-effective. VeloCloud SD-WAN Cloud VPN configuration enables traffic segmentation within geographic zones/regions. Inter-region traffic can be allowed with tier 1 (hub/data center) interconnects. All of this is centrally managed through VeloCloud Orchestrator.

Local inspection using VeloCloud Enhanced Firewall Service

Organizations with a security requirement to deploy a firewall locally at the branch can at the same time consolidate the branch hardware footprint by leveraging the Standard Firewall or Enhanced Firewall Service (EFS) on the VeloCloud Edge.

The Standard Firewall is a stateful firewall that monitors and tracks the operating state and characteristics of every network connection coming through the firewall and uses this information to determine which packets are permitted or denied. The Standard Firewall feature provides security benefits such as denial of service (DoS) and spoofing prevention, robust logging, and improved network security.

EFS builds on the Standard Firewall's security capabilities and provides additional security functionality such as URL and Category Filtering, URL Reputation Filtering, Malicious IP Filtering, and Intrusion Detection & Prevention System (IDS/IPS). Both the Standard Firewall and EFS protect Edge traffic locally and across the overlay (e.g., across branch-to-branch, branch-to-hub, or branch-to-Internet traffic patterns.

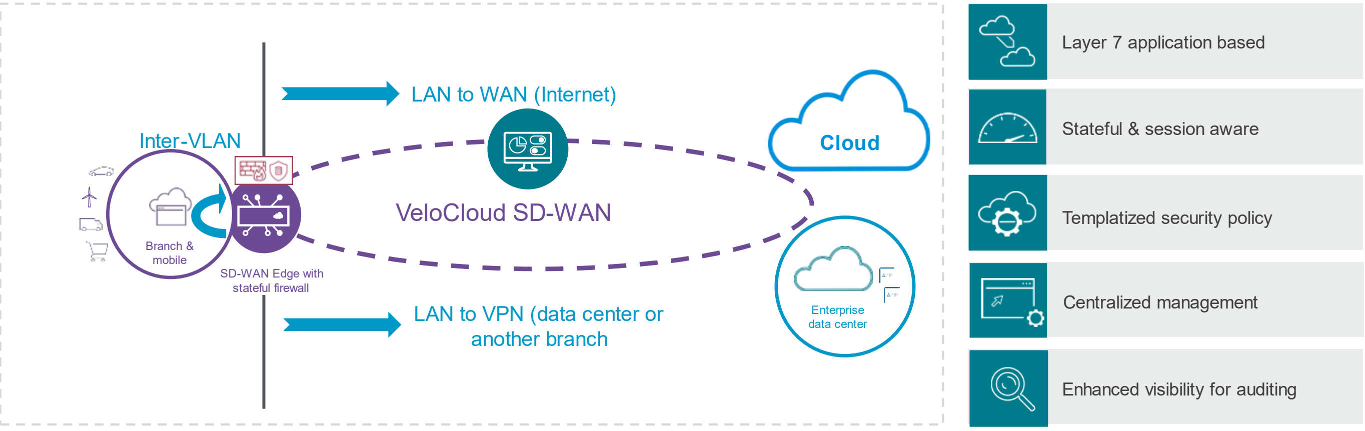

Figure 3: Overview of Edge Firewall security features

Integrating AI/ML capabilities within Enhanced Firewall Service

The built-in IDPS engine in EFS leverages a set of signatures to scan network payload for threats. Signature matching happens directly on the network components and covers both inbound and outbound activity. Signatures are sourced from third parties and generated in-house. The in-house signature uses a combination of manual analysis and machine learning-based candidate generation. The reputation and threat intelligence service uses a database of known bad IP addresses, DNS domains, and URLs to check for malicious traffic destinations. The database is constantly updated as newer bad destinations are unfolded with the advent of AI generated threats.

Symantec SSE for VeloCloud

Cloud-based security is a critical enforcement point of the overall AI networking solution. Just as we would expect to protect our data and apps on premises, we must ensure that our apps and services in the cloud are secure, as well as the users who are accessing and connecting to them. Symantec SSE for VeloCloud is a cloud-hosted service that protects users and infrastructure accessing SaaS and Internet applications from a changing landscape of internal and external threats, offers visibility and control, and ensures compliance.

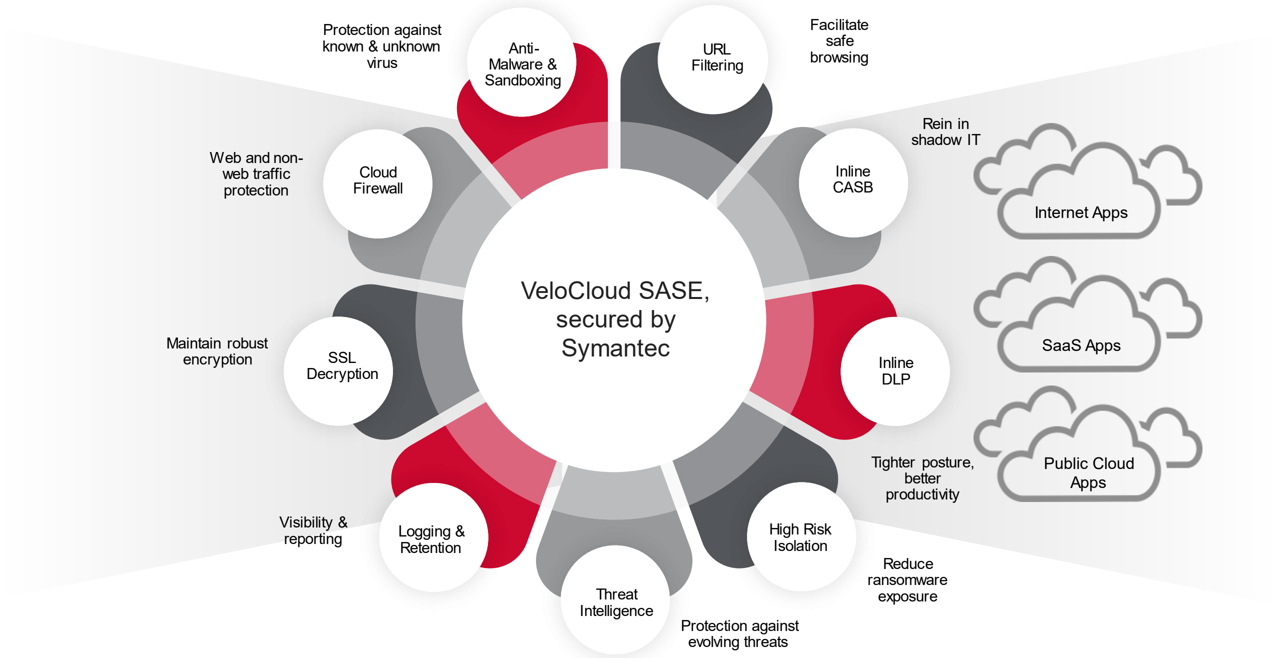
Protection against known & unknown virus
Anti-Malware & Sandboxing
URL Filtering
Facilitate safe browsing
Rein in shadow IT
Inline CASB
Internet Apps
Web and non-web traffic protection
Cloud Firewall
VeloCloud SASE, secured by Symantec
SaaS Apps
Maintain robust encryption
SSL Decryption
Inline DLP
Public Cloud Apps
Tighter posture, better productivity
Visibility & reporting
Logging & Retention
High Risk Isolation
Reduce ransomware exposure
Threat Intelligence
Protection against evolving threats

Figure 4: Symantec SSE for VeloCloud provides comprehensive threat and data protection

Encrypted traffic protection and management

Symantec SSE for VeloCloud provides high-speed decryption and re-encryption of SSL/TLS traffic to intercept and decrypt SSL and TLS traffic to uncover threats and potentially malicious content hidden in encrypted traffic. Further, the solution streamlines customer PKI management with self-managed certificates.

Proactive threat mitigation and data protection

Symantec SSE for VeloCloud delivers improved threat efficacy and response to manage continuous exposure to threats. The solution reduces the risk from a new generation of attacks or mutating threats using the largest civilian threat intelligence platform, Global Intelligence Network (GIN). GIN collects over a billion security telemetry data per day from 157 countries. This data is analyzed by a team of over 300 threat researchers using AI and ML for proactive threat detection and prevention. Symantec SSE helps reduce exposure to the most sophisticated and completely new generation of threats. For data protection, Symantec's DLP engine helps monitor and protect sensitive data using advanced DLP matching and recognition engines. The DLP engine is augmented with machine learning techniques to protect unseen data without explicit instructions/rules.

Machine learning-based malware detection

Symantec SSE for VeloCloud employs machine learning-based malware detection through its advanced threat analysis capabilities. The solution uses a combination of static and dynamic analysis techniques to identify potential threats. Static analysis involves examining the file's metadata, such as its name, size, and hash values, while dynamic analysis involves executing the file in a sandboxed environment to observe its behavior. Additionally, Symantec SSE for VeloCloud uses layered security with sandboxing to detect zero-day threats while blocking known threats using anti-malware and deep file inspection. The solution provides better detection and response along with prevention taking advantage of the analysis done on the telemetry data.

Tighter security with high-risk isolation

Every day over 250,000 web sites are introduced worldwide. The risk and reputation of these sites are not known to security vendors at the time of launch. About 50% of intrusion incidents involve credentials stolen when users interact with web applications, according to the Verizon 2024 Data Breach Investigations Report. Instead of blocking users from accessing these sites, Symantec SSE for VeloCloud creates an air gap between the users and the application to prevent users from being tricked into sharing credentials. It adds another layer of defense, at no extra cost, when users access websites that are uncategorized or categorized to be a higher risk. This approach strikes a balance between employee productivity and infrastructure security.

Reduced shadow AI risk using inline CASB

Inline CASB solution from Symantec SSE helps identify and control usage of third-party AI applications and tools not sanctioned by the enterprise IT. It helps enterprises by providing visibility into the Shadow AI application footprint used while reducing the security risks associated with the usage through policy management of unsanctioned applications.

Continuous web traffic protection

URL Filtering and Categorization service process over 6 billion web requests and block millions of web attacks and social engineering scams daily. Symantec SSE web protection uses dynamic, real-time risk ratings via the Global Threat Intelligence solution.

VeloCloud SD-WAN component security

VeloCloud SD-WAN Edge

The Edge is built on top of a customized Linux operating system distribution with an up-to-date, hardened kernel. All non-essential services, utilities, and accounts have been removed. If a service or utility is necessary, security configuration best practices have been implemented to reduce the attack surface.

The Edge implements control plane stateful firewalling and connection rate limiting to harden it against man-in-the-middle (MITM) and denial-of-service (DoS) attacks. Edges should be configured not to respond to connections on local ports.

An Edge configured for dynamic branch-to-branch virtual private network (VPN) will listen on UDP port 2426 for inbound tunnel connections from other branch Edges with configuration. While the use of this open UDP port is limited to only VeloCloud SD-WAN Edges that are authorized to connect with the appropriate IKEv2 credentials, some highly secure organizations may still choose to disable this feature, thereby leaving no ports open on the Edge.

VeloCloud Orchestrator and VeloCloud Gateway

The cloud-hosted VeloCloud SD-WAN service provides Orchestrators and Gateways running in secure SSAE Type II data centers and Tier 1 cloud data centers. Network ports are locked down to the minimum required. For the Gateway, only UDP 2426 traffic is permitted; and for the Orchestrator, only TLS (TCP/443) traffic is permitted. All other ports are blocked and all traffic to these ports is silently discarded.

As part of a comprehensive threat management approach, VeloCloud conducts various types of security penetration testing and vulnerability scanning, including internal and via third-party engagement. Such testing is carried out periodically and whenever the underlying product architecture undergoes significant changes. Vulnerabilities identified are assessed and patched based on the criticality of the scan result. There are three categories of scan results: Confirmed Vulnerability, Potential Vulnerability, and Information Gathered. Service-impacting vulnerabilities are fixed and patched within 24 hours of identifying the issue.

Figure 5 shows the methods used to secure communication between control, data, and management planes.

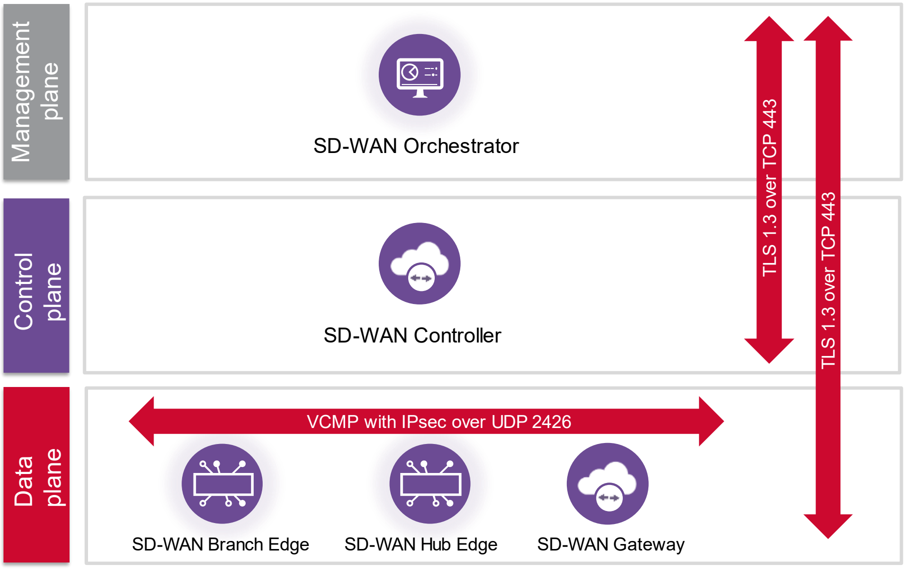

Figure 5: Secure communication between VeloCloud SD-WAN components

Security in the management plane

When a new VeloCloud Edge or Gateway is first provisioned, an activation key is generated by the Orchestrator. Edges and Gateways establish a Transport Layer Security (TLS) 1.2 or 1.3 encrypted session to the Orchestrator. The Orchestrator's public certification authority (CA) signed certificate is used to establish the TLS connection. Once the TLS connection is established, the activation key is used by the new device to authenticate and download configurations and policies from the Orchestrator. Along with the configuration and policy download, an authentication token and optionally a certificate is issued to the Edge and Gateway at time of activation. The authentication token (and a certificate if applicable) allows the Orchestrator to uniquely identify all the devices in its management domain. The process described here happens seamlessly during the zero-touch provisioning process.

Security in data plane and control plane

Security between the data plane and control plane components is managed through VeloCloud Management Protocol (VCMP), tunneling over User Datagram Protocol (UDP) port 2426. Internet key exchange version 2 (IKEv2) is used for IPsec negotiation and tunnel establishment between VeloCloud SD-WAN Edges and from VeloCloud SD-WAN Edges to Gateways. IKEv2 offers robust security and resiliency against attacks by bad actors.

IP Header	UDP Header	ESP Header	ESP IV	VCMP Header	VCMP Data	User Data	ESP Trailer	ESP Auth	Padding

Figure 6: How user data is encapsulated and encrypted for transport

VeloCloud SD-Access

VeloCloud SD-Access provides a secure, zero trust remote access solution designed to support distributed users, applications, and AI workloads. By optimizing connections for speed and quality, SD-Access enhances productivity while ensuring robust security. Its software-based, distributed architecture eliminates the need for traditional network hardware and static topologies, seamlessly integrating networking and security into a unified, cloud-managed solution. With unparalleled versatility for AI workloads, SD-Access offers:

- **User-to-application and workload-to-workload connectivity**: SD-Access can be seamlessly deployed directly on users and workloads, enabling direct, secure, robust communication without intermediary complexities.
- **Optimized AI traffic handling**: Native peer-to-peer connectivity support and AI-driven route optimization between nodes ensure low latency, high bandwidth, and seamless scalability for distributed AI processing.
- **Dynamic scalability**: VeloCloud SD-Access is a fully software-based solution that uses lightweight, encrypted tunnels, enabling flexible and scalable deployment across nodes. By eliminating the constraints of traditional hardware capacity, it is ideally suited to meet the demands of AI-driven environments.

Leveraging VeloCloud SD-Access for AI workloads

VeloCloud SD-Access integrates advanced networking and security features, offering unparalleled support for AI workloads. Key benefits of its robust architecture include:

Encrypted peer-to-peer traffic support

AI workloads often rely on peer-to-peer communication for distributed processing. SD-Access facilitates encrypted, direct connections between nodes, bypassing traditional bottlenecks and ensuring secure data exchange.

End-to-end encryption and advanced protocols

VeloCloud SD-Access ensures that all traffic traversing its network is encrypted end-to-end using AES-256, the industry standard for secure communication. By using modern DTLS protocols, SD-Access provides faster and more secure connections compared to legacy protocols like IPsec, which are prone to vulnerabilities and slower performance.

Key benefits include:

- **Dynamic session security**: When a connection is initiated, a lightweight tunnel is established, and a unique 2048-bit RSA certificate is generated for each session. This ensures that every connection is uniquely secured.
- **Enhanced encryption standards**: Data passing through the SD-Access tunnel is secured with DTLS, preventing any unauthorized interception or manipulation.
- **Man-in-the-middle attack prevention**: During the DTLS handshake, certificate fingerprints are validated to ensure no unauthorized entity can interfere with the connection. Session certificates never leave the device, ensuring that data remains fully encrypted and inaccessible to any external party, including relay nodes or potential eavesdroppers.

This advanced encryption framework ensures AI workloads can securely exchange data even in sensitive, distributed environments.

Zero trust networking

VeloCloud SD-Access employs a zero-trust architecture, where all user connections/attempts are strictly authenticated and authorized before being exposed to any network.

Key features include:

- **Identity-based policies**: All network and security policies are tied to strong user identities rather than IP addresses, ensuring precise and reliable access control.

- **Cloaking**: By ensuring no open TCP/IP ports on network nodes and disallowing inbound traffic, SD-Access removes its external attack surface, greatly reducing the risk of exploitation.
- **Contextual access rules**: Granular policies enable access control based on specific contexts such as time, geographic location, device posture, operating system, and encryption.

Authentication and authorization

Before any user or device can connect to the SD-Access network, the control plane authenticates and authorizes them. No unknown traffic is allowed, and every packet carries an identity for traceability and control. SD-Access integrates with third-party identity providers to enable:

- **Single sign-on (SSO):** Seamless user authentication.
- **Multi-factor authentication (MFA):** Enhanced security by requiring multiple layers of identity verification.
- **Corporate directory synchronization**: Ensures streamlined user and device management.

Native network segmentation for AI workloads

SD-Access provides native segmentation, ensuring each private network is visible only to its authorized members. In the event of a node compromise, the damage is contained within that specific network, limiting the "blast radius." This segmentation is especially critical for isolating AI workloads from other network traffic.

Best path selection and self-optimizing routes

SD-Access dynamically selects optimal routes using AI-driven analytics. By monitoring network conditions in real-time, it ensures that AI workloads receive the necessary bandwidth and low-latency pathways for uninterrupted operation.

Key benefits for AI workloads

- **Enhanced security**: End-to-end encryption, cloaking, and zero trust networking ensure robust protection of AI data and applications.
- **Dynamic scalability**: Encrypted, lightweight tunnels support high-performance peer-to-peer communication for distributed AI processing.
- **Resilience against threats**: Identity-based policies, cloaking, and contextual access rules significantly reduce attack surface.
- **Efficient performance**: AI-driven route optimization ensures low-latency, high-bandwidth paths tailored to the unique demands of AI workloads.
- **Improved segmentation**: Native segmentation prevents cross-contamination between different networks or compromised nodes.

By combining cutting-edge encryption, intelligent routing, zero trust principles, and seamless integration with identity management, VeloCloud SD-Access is designed to meet the stringent requirements of AI workloads while ensuring a secure, efficient, and scalable networking environment.

Conclusion

AI in networking is a transformative journey for every enterprise, offering many opportunities to enhance the efficiency, intelligence, and automation of modern infrastructures. Securing these systems is paramount to ensuring their safe and ethical deployment. As AI technologies continue to evolve, so too must the security strategies and frameworks that protect them.

By adopting a multi-layered security approach, leveraging robust defense techniques, and maintaining ethical standards, VeloCloud can help customers mitigate the risks and maximize the benefits of AI networking.

velocloud™
by Broadcom

BROADCOM®
connecting everything®

Item No: VC-SD-WAN-VeloRAIN-Security-WP100