velocloud™
by Broadcom

WHITE PAPER

VeloCloud SD-WAN Dynamic Multipath Optimization (DMPO)

Table of contents

Introduction to Dynamic Multipath Optimization

VeloCloud SD-WAN™ enables enterprises and service providers to use multiple WAN transports simultaneously, optimizing bandwidth and ensuring reliable application performance. Its cloud-delivered architecture enhances both on-premises and cloud applications, including Software-as-a-Service (SaaS) and Infrastructure-as-a-Service (IaaS). The solution builds an overlay network with multiple tunnels that continuously monitor and adapt to real-time changes in WAN transport performance. To maintain a resilient and high-performing network, VeloCloud developed Dynamic Multipath Optimization™ (DMPO). This document explains the key features and benefits of DMPO.

As illustrated in Figure 1 below, DMPO operates between all VeloCloud SD-WAN components that process and forward data traffic, including VeloCloud Edges (VCE) and Gateways (VCG).

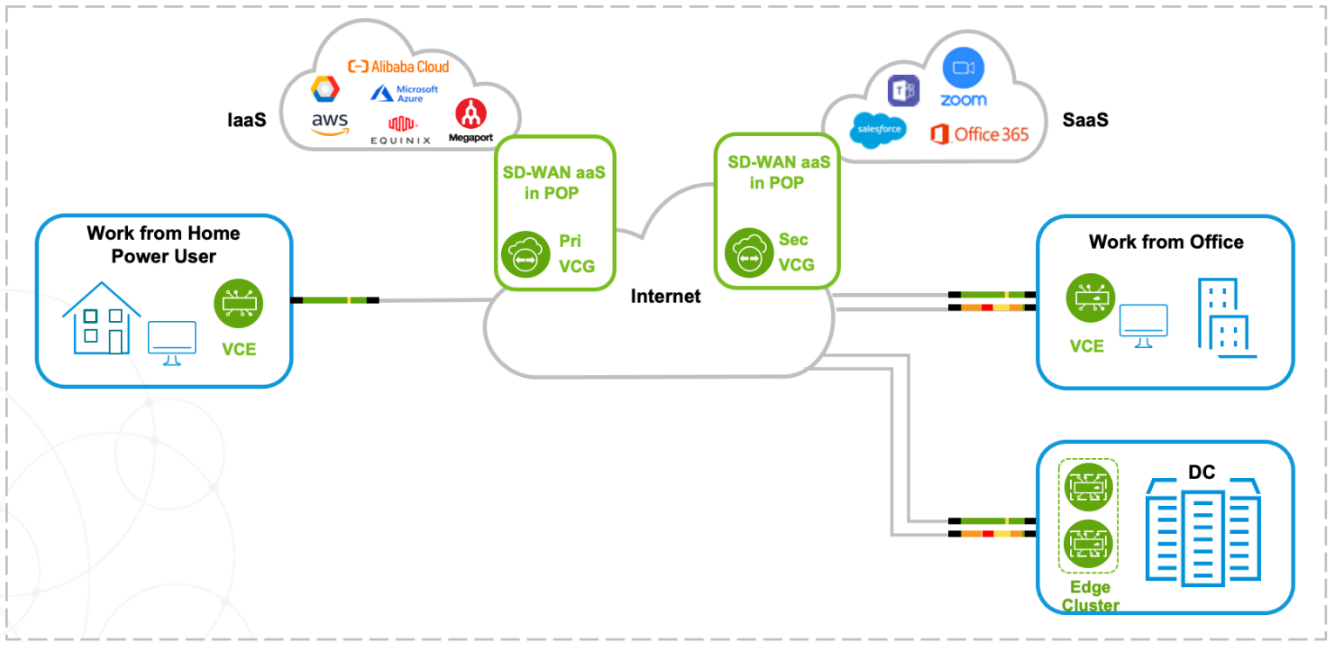

Figure 1: VeloCloud SD-WAN components capable of establishing full mesh connectivity using DMPO tunnels

- For enterprise site-to-site connectivity (e.g., branch-to-branch or branch-to-hub), VeloCloud SD-WAN Edges establish DMPO tunnels directly with each other.
- For cloud application access, each VeloCloud SD-WAN Edge forms DMPO tunnels with one or more VeloCloud SD-WAN Gateways.

Next, we'll outline the key features of DMPO.

Continuous monitoring

Automated bandwidth discovery

VeloCloud Edges are deployed using a zero-touch deployment process. Once an internet WAN link is detected, the VeloCloud SD-WAN Edge establishes DMPO tunnels with one or more VeloCloud SD-WAN Gateways and conducts a bandwidth test with the nearest Gateway. This test involves sending a short burst of bidirectional traffic and measuring the received rate at each end.

Because VeloCloud SD-WAN Gateways are deployed at Internet Points of Presence (PoPs), they can also determine the WAN link's real public IP address, even if the VeloCloud SD-WAN Edge is behind a Network Address Translation (NAT) or Port Address Translation (PAT) device.

A similar process applies to private links. For hub or headend VeloCloud SD-WAN Edges, WAN bandwidth is statically defined. However, when a branch VeloCloud SD-WAN Edge establishes a DMPO tunnel to a hub, the bandwidth test follows the same procedure as it does between the Edge and Gateway on a public link.

Continuous path monitoring for latency, jitter, and packet loss

DMPO continuously measures key performance metrics—loss, latency, and jitter—for every packet on every tunnel between DMPO endpoints, which include VeloCloud SD-WAN Edges and Gateways. These measurements are unidirectional, ensuring precise monitoring of both uplink and downlink performance.

VeloCloud SD-WAN's per-packet steering enables independent traffic decisions in both directions without causing asymmetric routing. DMPO employs both passive and active monitoring:

- **Passive monitoring**: When user traffic is present, DMPO embeds additional performance metrics in the tunnel header, including sequence numbers and timestamps. This allows DMPO endpoints to detect lost or out-of-order packets and calculate jitter and latency in real-time. Performance metrics are exchanged between endpoints every 100 milliseconds (ms).
- **Active monitoring**: If no user traffic is detected, DMPO sends an active probe every 100 ms. If no high-priority user traffic is present for five minutes, the probe frequency is reduced to 500 ms.

This continuous and adaptive measurement enables DMPO to quickly respond to changing WAN conditions, providing sub-second protection against link degradation or complete circuit outage.

Continuous path monitoring for available bandwidth

DMPO employs Dynamic Bandwidth Adjust (DBA) as a real-time network optimization technique that automatically modifies bandwidth allocation based on traffic demand, network congestion, application requirements, and QoS policies, ensuring efficient bandwidth usage by dynamically adapting to changing network conditions. DBA is employed to accurately measure variable bandwidth often found with wireless links, such as 4G/LTE, 5G, and satellite. DBA is only used when there is user traffic present and there is contention for the available bandwidth:

- **Passive monitoring**: When user traffic is present, DBA leverages consecutive data packets (with incremental sequence numbers) that are transmitted "back-to-back." The receiving end can then estimate the short-term bandwidth based on the data volume transmitted and the net time used to receive the consecutive packets by tagging them with special flags in the header.

Gateway and hub path Quality of Experience (QoE) monitoring

Each VeloCloud SD-WAN Edge establishes DMPO tunnels with both its primary and secondary Gateways, ensuring control plane resiliency. If the primary Gateway fails, traffic seamlessly redirects to the secondary Gateway. Metadata for each flow is shared with the secondary Gateway guiding how to processes data packets. Since overlay tunnels on the secondary Gateway are pre-established, they can carry traffic immediately, enabling seamless failover.

DMPO also provides Quality of Experience (QoE) measurements for each Gateway path, tailored to different application types. The QoE score reflects application performance and the overall user experience delivered by the network over a given period. It's calculated by comparing all static tunnels (Edge-to-Gateway and Edge-to-Hub) and highlights the best-performing paths.

The QoE graph visually represents the quality score of a selected path before and after DMPO optimization, demonstrating its impact on application performance.

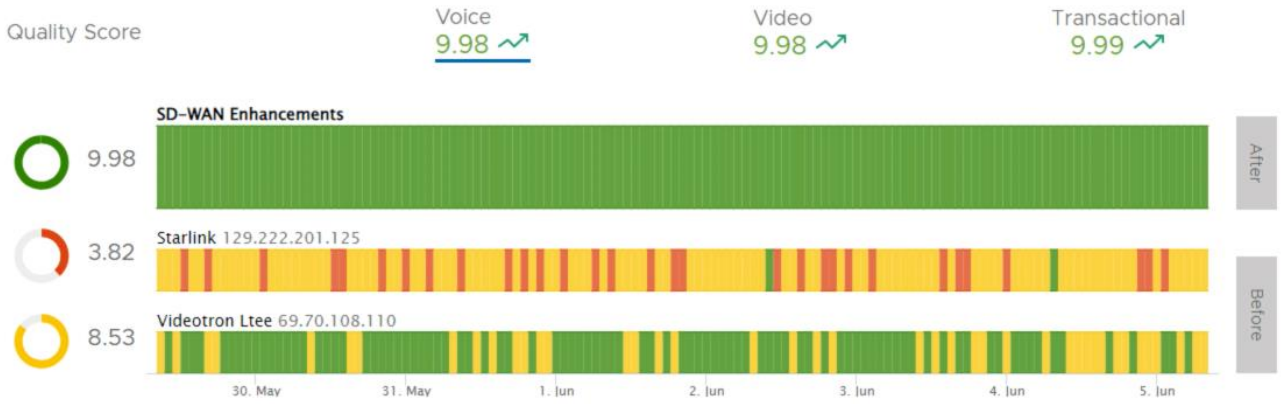

Figure 2: Quality of Experience (QoE) scoring for voice, video and transactional application types

Dynamic application steering

Application-aware per-packet steering

DMPO intelligently identifies traffic using Layer 2 to Layer 7 attributes, such as VLAN, IP address, protocol, and applications. VeloCloud SD-WAN appliances application-aware, per-packet steering based on business policies and real-time link conditions.

Policy defaults

VeloCloud SD-WAN includes pre-configured business policies with smart defaults for thousands of applications, specifying their priority and default steering behavior. This allows customers to immediately benefit from dynamic packet steering and application-aware prioritization without manually defining policies.

DMPO continuously monitors all WAN links and dynamically steers traffic based on current network conditions. A single traffic flow can seamlessly switch between DMPO tunnels mid-session without disruption and can adapt dynamically to different types of network conditions.

- **Outage condition**: A link that is completely down.
- **Degraded link condition**: A link that cannot meet the SLA requirements for a given application.

VeloCloud SD-WAN provides sub-second outage and network degradation protection. DMPO detects real-time link conditions within 300–500 milliseconds and instantly redirects traffic to maintain application performance with no observed impact on end users. If the affected link recovers, DMPO waits one minute before steering traffic back, as specified in business policies.

First-packet application steering

DMPO learns and caches application classifications to enable real-time traffic steering from the first packet. This supports application-based redirection, such as:

- Routing Netflix or YouTube traffic for Direct Internet Access (DIA) at the branch to bypass backhauling through DMPO tunnel to a Gateway or Hub.
- Backhauling GitHub traffic to a regional enterprise hub or data center for centralized security and compliance.

Example #1: Zoom traffic optimization

As an example of intelligent application steering, smart defaults classify Zoom as a high-priority, real-time application. In this example, consider an Edge with the following two WAN links:

- Link A: 50ms latency
- Link B: 60ms latency

If all other SLAs are met, DMPO chooses Link A due to lower latency. However, if latency on Link A increases to 200ms, DMPO immediately steers Zoom traffic to Link B (60ms latency) within milliseconds to maintain a high call quality.

Example #2: Data security

Intelligent application steering can also improve data security.

- DMPO ensures that sensitive data in PCI-compliant environments is only transmitted over secure links. If no secure link is available, VeloCloud SD-WAN blocks transmission to prevent exposure.

Example #3: Cost optimization

Intelligent application steering can also improve cost savings for remote workers.

- DMPO can prioritize low-cost connections (e.g., broadband) over expensive options (e.g., 4G/LTE, 5G, Satellite), reducing connectivity expenses for remote employees.
- Additionally, DMPO features such as 'Limit Control Traffic Frequency' can be enabled on metered wireless transports (e.g., 4G/LTE, 5G, satellite) to reduce the amount of control traffic generated by the Edge.

Multiprotocol Label Switching Class of Service (CoS)

For private links with a Class of Service (CoS) agreement, DMPO considers CoS levels when monitoring performance and making application steering decisions. In an MPLS network, service providers offer different SLAs for each CoS. DMPO can treat each CoS as a separate link, enabling granular, application-aware traffic steering based on performance and business policies.

For example, if a service provider offers two CoS levels—CoS1 and CoS2, each with its own SLA—DMPO can intelligently decide whether to route traffic over CoS1 or the Internet, or CoS2 or the Internet, depending on real-time network conditions and application requirements.

5G network slicing

5G carriers offer network slicing, which partitions physical 5G infrastructure into independent, isolated, and programmable virtual networks. Each network slice is optimized for specific applications and use case

- **Enhanced Mobile Broadband (eMBB):** High-speed data and bandwidth for applications like HD streaming and AR/VR.
- **Ultra-Reliable Low Latency Communication (uRLLC):** Enables mission-critical applications like autonomous vehicles and remote surgery.
- **Massive Machine Type Communication (mMTC):** Supports IoT devices with low data rates and massive connections.

VeloCloud's Dynamic Application-Based Slicing (DABS) is a policy framework that classifies and prioritizes application traffic for 5G Network Slices. The SD-WAN Edge connects to a carrier's network slice through a dedicated slice interface, creating a netlink. These interfaces leverage Qualcomm Multiplexing and Aggregation Protocol (QMAP), which efficiently multiplexes and aggregates multiple IP data flows over a single physical interface, optimizing performance and resource utilization.

Bandwidth aggregation

For applications that require high bandwidth, such as large file transfers, DMPO utilizes a per-packet processing engine to intelligently distribute traffic across multiple WAN links, maximizing bandwidth while ensuring seamless performance. This approach differs significantly from traditional flow-based load sharing solutions.

How per-packet load balancing works in DMPO

Unlike flow-based load sharing, which assigns entire traffic flows to a single WAN link, DMPO dynamically distributes individual packets of a single flow across multiple links based on real-time WAN performance metrics such as latency, jitter, and packet loss.

- Real-time path selection: DMPO continuously monitors WAN conditions and determines the optimal path for each packet, ensuring efficient use of available bandwidth.
- Packet-level steering: Instead of directing an entire flow to one link (as in flow-based load sharing), DMPO splits traffic at the packet level, allowing simultaneous use of multiple links.
- Resequencing at the destination: Since packets may arrive out of order when transmitted over different links, with different latencies, DMPO performs intelligent packet resequencing at the receiving end, ensuring that the application experiences a smooth and uninterrupted flow.

Table 1: Key differences between DMPO per-packet bandwidth aggregation vs. flow-based load sharing

Feature	DMPO (per-packet bandwidth aggregation)	Flow-based load sharing
Traffic distribution	Splits packets of a single flow across multiple links	Assigns entire flows to a single link
Real-time adaptation	Adjusts per packet based on link conditions	Only reassigns flows when failures occur
Bandwidth utilization	Aggregates bandwidth across all links for a single flow	Limited by the capacity of a single link per flow
Resequencing	Performs packet resequencing at the destination to prevent out-of-order delivery	No need for resequencing, but limited to per-flow capacity
Performance impact	Optimized for high-bandwidth applications like file transfers	Can lead to underutilization if a single flow maxes out a link

Example: 100 Mbps file transfer over two 50 Mbps links

With a flow-based load sharing a 100 Mbps file transfer would be limited to a single 50Mbps link, leaving the second link underutilized, even though it's an active and available WAN link.

With VeloCloud DMPO per-packet load balancing, the transfer utilizes both links simultaneously, achieving the full 100 Mbps bandwidth. QoS policies are applied at both the aggregate level and the individual link level, ensuring smooth and efficient data transfer.

By using per-packet load balancing, DMPO delivers higher bandwidth, better link utilization, and improved application performance compared to traditional flow-based solutions.

On-demand remediation

In situations where traffic cannot be steered to a better-performing link—such as when only a single link is available or when multiple links experience issues simultaneously—DMPO applies error correction to maintain application performance during the disruption.

The specific error correction method used depends on:

- **Application type**: Different applications have varying sensitivity to packet loss, latency, and jitter.
- **Error type**: DMPO dynamically selects the most effective correction technique based on the nature of the network degradation.

By applying real-time error correction, DMPO ensures consistent application performance, even in challenging network conditions. Let's look at some different application types and some example DMPO techniques that are used to remediate the degraded network detected.

On-demand remediation on real-time application

Real-time applications such as voice and video benefit from Forward Error Correction (FEC) during packet loss. DMPO automatically applies FEC on both single and multi-link deployments to maintain call and video quality.

- **For multiple links**: DMPO selects up to two of the best-performing links at any given time for FEC. Duplicate packets are discarded, and out-of-order packets are resequenced at the receiving end before delivery.
- **For a single link**: DMPO transmits duplicate packets over the same path, increasing the likelihood that at least one copy reaches its destination.
- **To mitigate jitter**: DMPO enables a jitter buffer for real-time applications when WAN links experience fluctuations, ensuring smoother playback and reduced disruptions.

On-demand remediation on TCP-based application

TCP applications, such as file transfers, benefit from Negative Acknowledgment (NACK) in DMPO.

When a packet is lost, the receiving DMPO endpoint detects the missing packet and notifies the sending DMPO endpoint to retransmit it. This prevents the end application from experiencing packet loss, helping to maintain a stable TCP window and ensuring high TCP throughput, even in lossy network conditions.

Application-aware overlay QoS

In a VeloCloud SD-WAN network, DMPO tunnels are established in two scenarios:

1. Between a VeloCloud SD-WAN Edge and a VeloCloud SD-WAN Gateway
2. Between two VeloCloud SD-WAN Edges

Before leaving the VeloCloud SD-WAN Edge, a VeloCloud SD-WAN Management Protocol header is added to the packet, introducing an overhead of 59 bytes.

When the packet reaches the destination—either a VeloCloud SD-WAN Gateway or another VeloCloud SD-WAN Edge—all tunnel headers (including the VeloCloud SD-WAN Management Protocol and IPsec) are removed. The original user data is then forwarded to the next-hop router, which could be:

- A Provider Edge (PE) router in a service provider network
- An L3 switch or router in an enterprise network

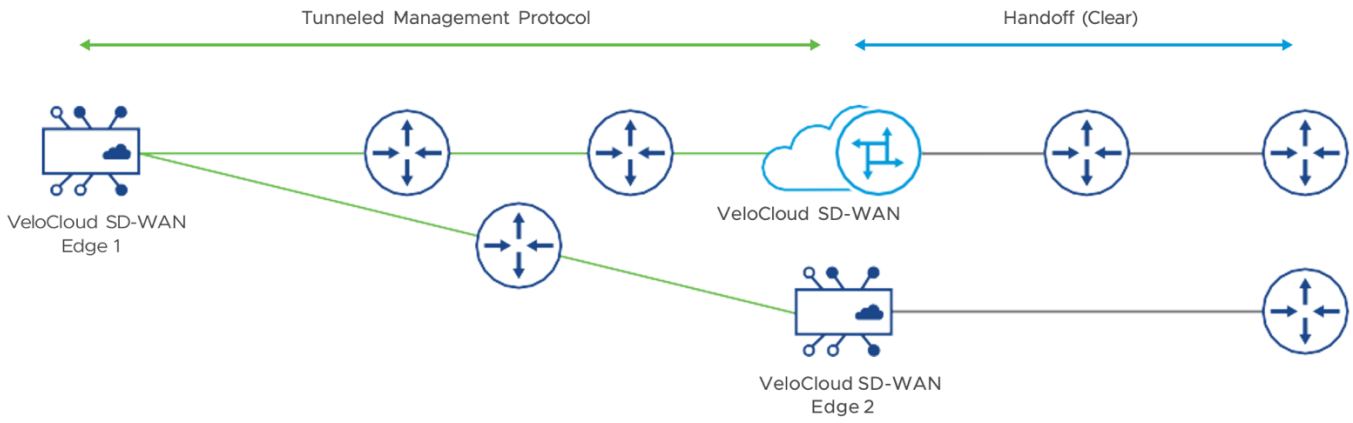

Figure 3: VeloCloud SD-WAN Tunnel Management Protocol: VCMP

QoS scheduling

A traffic class in VeloCloud SD-WAN is defined by a combination of priority (High, Normal, or Low) and service class (Real-Time, Transactional, or Bulk), creating a 3x3 matrix with nine traffic classes. Applications and categories, along with scheduler weights, are mapped to these traffic classes to ensure efficient traffic management.

Traffic class behavior

- **Aggregate QoS treatment**: All applications within the same traffic class share the same Quality of Service (QoS) policies, including scheduling and policing.
- **Guaranteed bandwidth during congestion**: Each traffic class is assigned a minimum guaranteed bandwidth, determined by the scheduler weight (percentage of bandwidth allocation).
- **Bursting beyond minimum allocation**: When there is no congestion, applications can use available bandwidth beyond their guaranteed allocation, up to the maximum aggregated bandwidth.
- **Bandwidth capping**: Policies can be applied to limit bandwidth usage for all applications within a specific traffic class.

Smart defaults and custom policies

- The business policy includes pre-configured smart defaults that automatically maps thousands of applications to the appropriate traffic classes. This allows customers to immediately benefit from application-aware QoS without manual configuration.
- Customers can also define custom policies for their own applications.
- Each traffic class has a default scheduler weight, which can be adjusted within the business policy settings to fine-tune bandwidth allocation.

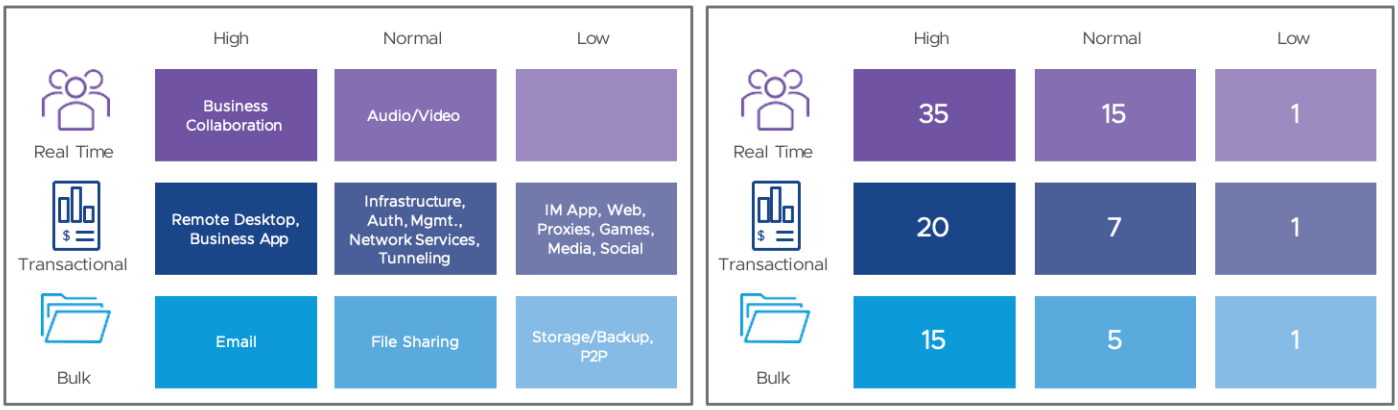

Figure 4: Default value for the 3x3 matrix with 9 traffic classes

For example, if a customer has a 90 Mbps Internet link and a 10 Mbps MPLS link, the total aggregate bandwidth is 100 Mbps.

Based on the default scheduler weights and traffic class mapping (as shown in Figure 4):

- Business collaboration applications (e.g., Microsoft Teams, Zoom) are guaranteed 35 Mbps of bandwidth.
- Email applications (e.g., Outlook, Gmail) are guaranteed 15 Mbps of bandwidth.

Business policies can be applied at different levels:

- **Category level**: e.g., all business collaboration applications.
- **Application level**: e.g., Microsoft Teams.
- **Sub-application level**: e.g., specific Teams features such as file transfer, audio, or video.

This flexibility allows customers to optimize QoS policies for different application types based on business needs.

Class of Service (CoS) marking

When traffic arrives at a VeloCloud SD-WAN Edge, the Differentiated Services Code Point (DSCP) values marked by the customer can either:

- Remain unchanged before being encapsulated in a tunnel.
- Be modified before transmission.

Additionally, the outer DSCP value in the tunnel header can:

- Be copied from the original inner packet.
- Be assigned a new value based on business policies.

Here is an example scenario (see Figure 5). Consider two traffic flows:

- Voice traffic (high priority)
- Data traffic (lower priority)

Customer's DSCP policy:

- Inner packet DSCP: The customer chooses to preserve DSCP values for both voice and data.
- Outer packet DSCP:

– Voice traffic: The DSCP value is copied to the outer packet, maintaining prioritization across the network.
– Data traffic: The outer DSCP value is reset to DSCP=0, deprioritizing it in the tunnel.

This approach ensures consistent QoS treatment while allowing flexibility in marking and prioritizing traffic within the SD-WAN overlay.

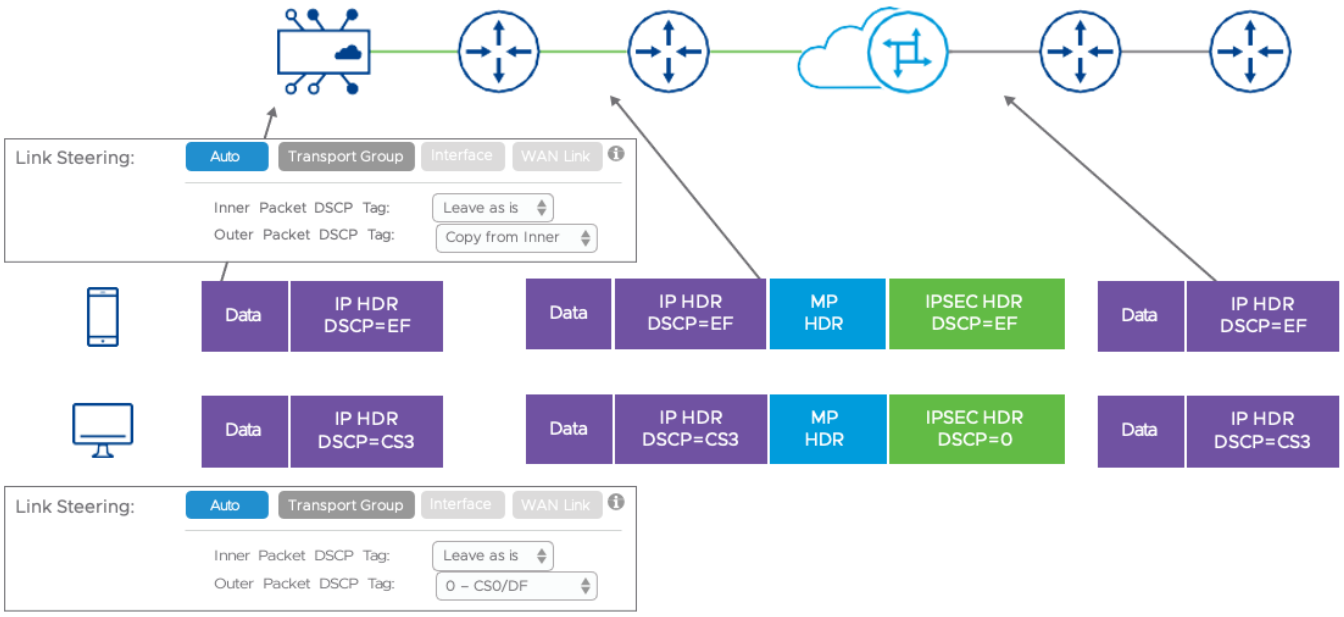

Figure 5: Traffic flow CoS marking example

Policing traffic class

In legacy WAN networks, service providers and enterprises allocate bandwidth and enforce traffic policies based on Class of Service (CoS) levels defined by the provider. With VeloCloud SD-WAN, a similar approach is needed for the WAN overlay, which may use multiple transport types from different service providers.

Overlay traffic policing in VeloCloud SD-WAN

IT administrators can apply traffic policing on the aggregated overlay tunnel to:

- Ensure high-priority business collaboration traffic aligns with the service provider's SLA.
- Limit non-critical applications for security or QoS compliance.

Traffic policing is defined at the traffic class level, which combines both service class (e.g., Real-Time, Transactional, Bulk) and priority (High, Normal, Low).

Example scenario: A customer has a 90 Mbps Internet link and a 10 Mbps MPLS link, providing an aggregated bandwidth of 100 Mbps. Based on the default traffic class mapping (Figure 4), business collaboration applications are guaranteed 35 Mbps of bandwidth.

The service provider can enforce a policy on this traffic class, ensuring that all applications within this category are policed at 35 Mbps, even when there is no congestion in the network. This approach ensures consistent bandwidth allocation, prevents overuse of priority traffic, and helps enforce QoS policies across multiple transport links in the SD-WAN overlay.

Policing MPLS CoS

For private links with a CoS agreement from an MPLS provider, different Service Level Agreements (SLAs) apply to each CoS. DMPO can treat each CoS as a separate logical link, allowing for granular, application-aware traffic steering based on real-time conditions.

Policy enforcement for MPLS CoS underlay

To ensure the service provider's committed bandwidth SLAs are met, customers can define policies for MPLS CoS traffic within the SD-WAN overlay.

Example scenario: A branch SD-WAN Edge has a 10 Mbps MPLS link. The MPLS provider's SLA allocates:

- 40% (4 Mbps) for CoS1 (e.g., real-time traffic, tagged as DSCP=EF, CS5).
- 60% (6 Mbps) for all other traffic.

The service provider's Provider Edge (PE) enforces these limits:

- Total MPLS traffic is capped at 10 Mbps.
- CoS1 traffic exceeding 4 Mbps is dropped, potentially impacting QoS.

How DMPO ensures compliance: A traffic policy on the SD-WAN Edge enforces a 4 Mbps limit for CoS1 to prevent packet loss due to provider policing. Remaining traffic can use up to the full 10 Mbps if no congestion exists. During congestion, bandwidth is allocated based on configured minimum guarantees. By proactively managing CoS-based traffic policies, VeloCloud SD-WAN ensures compliance with MPLS SLAs, prevents unnecessary packet drops, and optimizes application performance across the private link.

Rate limiting an application or application category

Rate limiting in VeloCloud SD-WAN can be applied to both inbound and outbound traffic for specific applications. When a rate limit is enforced and congestion occurs, traffic is queued, and if the queue reaches capacity, excess packets are dropped to prevent network overload.

Example: Managing Hulu traffic. Consider a scenario where customers access Hulu, where:

- Outbound traffic (requests to Hulu) is minimal.
- Inbound traffic (streaming video) is significant.

Traditional WAN challenge: In a traditional WAN setup, congestion is only detected after traffic has reached the edge router, by which time the WAN link is already saturated, leading to buffering, packet drops, or degraded performance.

VeloCloud SD-WAN solution: With inbound QoS, VeloCloud SD-WAN can proactively signal the streaming application to reduce its bitrate, preventing the Hulu traffic from exceeding the configured inbound bandwidth limit (Figure 6). This ensures:

- Optimized bandwidth usage without oversaturating the WAN link.
- Better performance for other critical applications running on the network.
- A smoother streaming experience without unexpected buffering or quality degradation

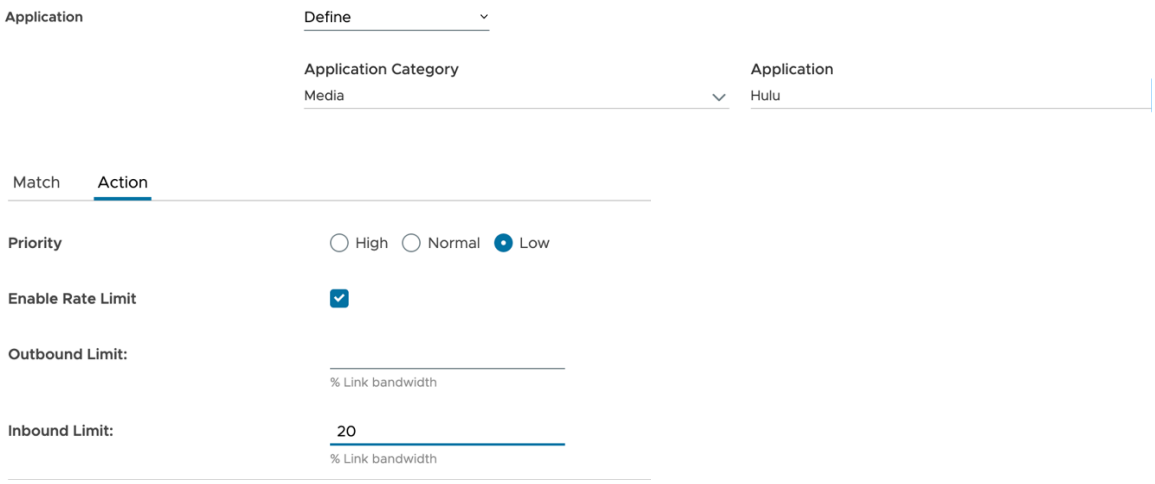

Figure 6: Rate limiting an application

DMPO tunnel shaper for service providers with a partner Gateway

Service providers may offer SD-WAN services with a lower capacity than the total aggregated bandwidth of a branch's WAN links. This can occur when:

- A customer purchases a broadband link from another vendor, outside the service provider's control.
- The service provider hosts a VeloCloud SD-WAN Gateway but has no direct control over the underlay broadband connection.

Ensuring SD-WAN service capacity compliance

To prevent congestion toward the partner gateway and ensure the SD-WAN service capacity is honored, the service provider can enable a DMPO tunnel shaper between the VeloCloud SD-WAN Edge and the Partner Gateway.

In an example scenario shown in Figure 7, the VeloCloud SD-WAN Edge has two WAN links:

- 20 Mbps Internet
- 20 Mbps MPLS

The service provider's SD-WAN service capacity is limited to 35 Mbps. Without traffic shaping, the combined traffic could exceed 35 Mbps, potentially causing congestion. To enforce the 35 Mbps limit, the service provider applies a tunnel shaper on the DMPO tunnel, ensuring that total traffic toward the partner gateway (X in Figure 7) does not exceed the subscribed SD-WAN capacity.

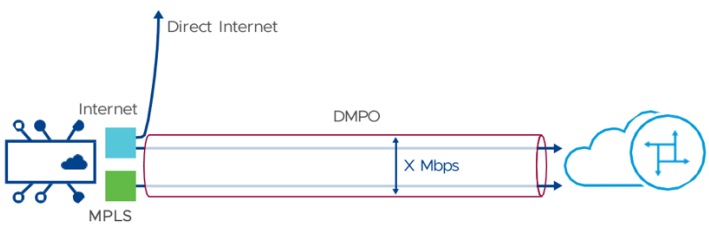

Figure 7: DMPO tunnel shaper

Business priority monitoring

Application traffic can be monitored in real time based on its assigned priority, with historical data available for analysis. Traffic metrics can be viewed in various formats, including:

- Bytes sent and received
- Packets sent and received
- Average throughput

This allows administrators to track network performance and analyze trends over time.

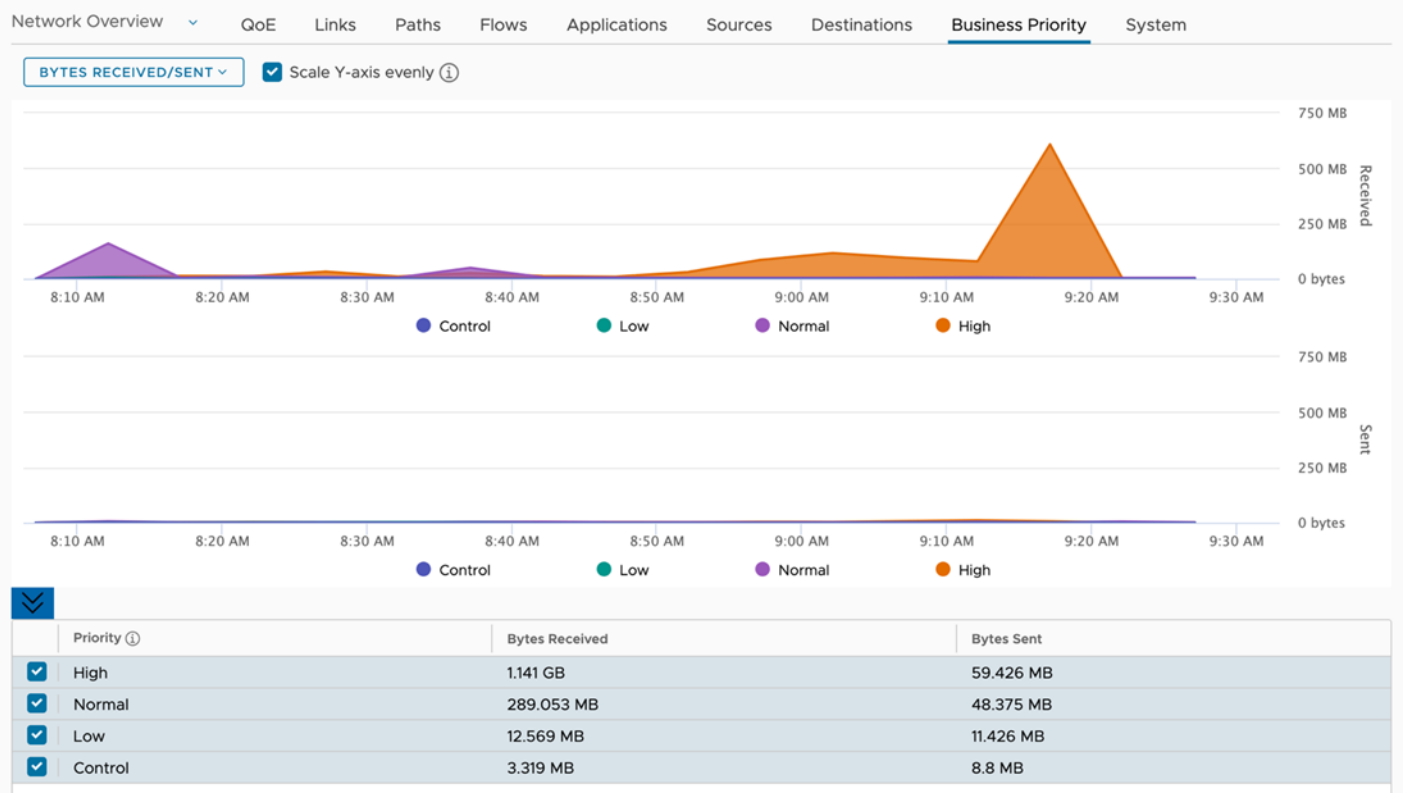

Figure 8: Business priority monitoring dashboard

Business policy framework and smart defaults

The IT administrator manages VeloCloud SD-WAN Quality of Service (QoS), traffic steering, and network services for application traffic through the Business Policy framework.

Business policy and smart defaults: Smart defaults provide pre-configured business policies for thousands of applications, enabling automatic optimization without manual configuration.

DMPO dynamically makes traffic steering decisions based on:

- Application type
- Real-time link conditions (congestion, latency, jitter, packet loss)
- Business Policy rules

Application categories and traffic management: Each application belongs to a category, which has a predefined default action. A default action consists of:

- Traffic class (a combination of priority and service class)
- Network services
- Link steering behavior

In addition to the predefined application list, customers can manually define custom applications and assign specific policies. The following section provides an example of a business policy configuration.

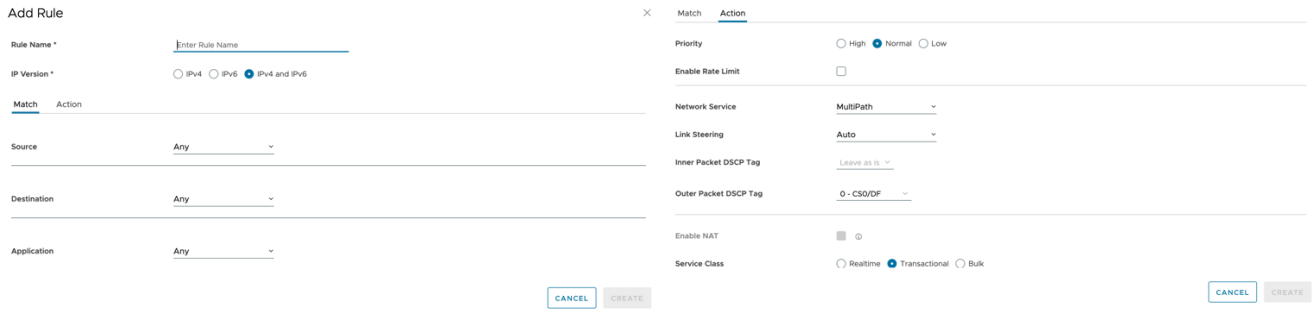

Figure 9: Business policy (match & action)

Traffic class (priority and service class)

Applications and categories are assigned to a traffic class based on a combination of priority and service class. All applications within the same traffic class receive aggregated QoS treatment, which includes scheduling and policing, as outlined in the "Application-aware overlay QoS" section.

Network services

By default, each application is assigned one of three network services, which can be modified by the user based on business needs:

- **Direct**: Used for non-critical, trusted Internet applications that do not require SD-WAN optimization. Traffic is sent directly to Internet (DIA), bypassing the DMPO tunnel. Load balancing is applied at the flow level. By default, all low-priority applications are assigned to the Direct network service.

  Example: Netflix, a high-bandwidth, non-business application, should not consume SD-WAN resources.

- **MultiPath**: Used for important business applications that benefit from SD-WAN optimization. Sends Internet-based traffic to a VeloCloud SD-WAN Gateway or Hub for enhanced steering and remediation. By default, high and normal-priority applications are assigned the MultiPath network service.

  Table 3 illustrates the default link steering and remediation techniques based on service class.

  MultiPath additionally can redirect application traffic to Cloud Security Services (also known as a non SD-WAN destination) for additional inspection and filtering.

  Example: Websense (now Forcepoint)

- **Internet backhaul**: Redirects Internet applications to a specified enterprise location for security processing. Used when security devices such as firewalls, intrusion prevention systems (IPS), and content filtering must inspect the traffic before it exits to the Internet. The enterprise location may or may not have a VeloCloud SD-WAN Edge.

Important note: Table 2 displays the default network service assignments for different applications.

VPN traffic is always sent through SD-WAN tunnels, even if an application is assigned the Direct network service.

Table 2: Default values for a network service action

Priority	Destination: Internet (e.g., SaaS, web traffic	Destination: Within the enterprise VPN
High	MultiPath (through DMPO tunnels)	MultiPath
Normal		
Low	Direct	

Link steering

In the Business Policy, there are four link steering modes for traffic routing:

- **Auto**: DMPO dynamically selects the best path based on real-time network conditions.
- **By transport group**: Traffic is steered based on predefined groups of transport types (e.g., broadband, MPLS, LTE).
- **By WAN link**: Traffic is directed over a specific WAN link (e.g., a designated Internet or MPLS connection).
- **By interfaces**: Traffic is steered based on a specific network interface (e.g., Ethernet or LTE).

Link steering: Auto

By default, all applications use automatic link steering mode, allowing DMPO to dynamically select the best WAN link based on real-time network conditions and application type. When needed, on-demand remediation is automatically applied to maintain application performance.

For Internet applications, there are four possible combinations of link steering and on-demand remediation.

For enterprise (VPN) traffic, DMPO tunnels are always used, ensuring it automatically benefits from on-demand remediation at all times.

Table 3: Default link steering and on-demand remediation for a given service class

		Destination: Internet	
Service class		Network service: MultiPath Link steering: Auto	Network service: Direct Link steering: Auto
Real time	Link selection behavior	Per-packet steering	Flow-based load balancing
	On-demand remediation	FEC and jitter buffer	—
Transactional	Link selection behavior	Per-packet load balancing	Flow-based load balancing
	On-demand remediation	NACK	—
Bulk	Link selection behavior	Per-packet load balancing	Flow-based load balancing
	On-demand remediation	NACK	—

The following examples illustrate the default DMPO behavior for various real-time application types under different network conditions:

Table 4: Default DMPO behavior for different real-time application types and link conditions

Scenario	Expected DMPO behavior
At least one link that satisfies the SLA for the application	Pick the best available link.
Single link with packet loss exceeding the SLA for the application	Enable FEC for the real-time applications sent on this link.
Two links with loss on only one link	Enable FEC on both links.
Multiple links with loss on multiple links	Enable FEC on two best links.
Two links but one link appears unstable, i.e. missing three consecutive heartbeats	Mark link un-usable and steer the flow to the next best available link.
Both jitter and loss on both links	Enable FEC on both links and enable jitter buffer on the receiving side. Jitter buffer is enabled when jitter is greater than 7ms for voice and greater than 5ms for video. The sending DMPO endpoint notifies the receiving DMPO endpoint to enable jitter buffer. The receiving DMPO endpoint will buffer up to 10 packets or 200ms of traffic, whichever happens first. The receiving DMPO endpoint uses the original timestamp embedded in the DMPO header to calculate the flow rate to use in de-jitter buffer. If flow is not sent at a constant rate, the jitter buffering is disabled.

Link steering by transport group

Different locations may use different WAN transports, such as various carriers or interface types. To simplify network management, DMPO utilizes transport groups to abstract WAN carriers and interfaces from business policy configurations.

How transport groups work: Instead of specifying individual WAN links, business policies can define traffic steering based on transport groups, such as:

- Public wired (e.g., broadband, fiber)
- Public wireless (e.g., LTE, 5G)
- Private wired (e.g., MPLS)

This allows the same business policy to be applied across different locations and device types, even if they use different WAN carriers or interfaces.

When DMPO discovers WAN links, it automatically assigns them to the appropriate transport group. This approach eliminates the need for IT administrators to manually configure policies based on specific physical connectivity types or WAN providers.

Using transport groups in business policies ensures scalability and consistency, allowing IT teams to apply policies universally across locations without concern for underlying WAN differences.

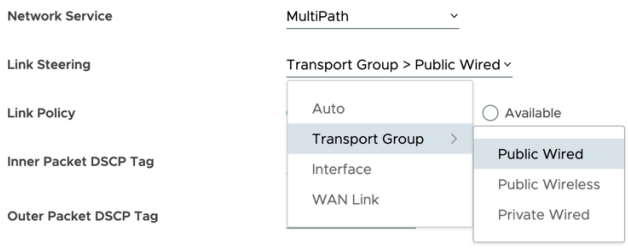

Figure 10: Link steering by transport group

Link steering by WAN link

Each WAN interface connects to a WAN carrier, which is specific to the location of the VeloCloud SD-WAN Edge.

DMPO automatically identifies the WAN carrier using GeoIP lookup. Alternatively, the IT administrator can manually specify the WAN carrier.

Link steering can be based on MPLS CoS settings defined in the WAN overlay. Figure 11 illustrates an MPLS CoS agreement with the following bandwidth allocations:

- CoS1 (CS5, EF): 60% guaranteed bandwidth (for high-priority traffic)
- CoS2 (AF41, CS4): 20% guaranteed bandwidth
- CoS5 (AF21, CS2): 20% guaranteed bandwidth

MPLS CoS1 ensures that no more than 60% of the total bandwidth is used for that class.

This approach allows DMPO to optimize traffic steering based on real-time network conditions while ensuring MPLS SLAs are met.

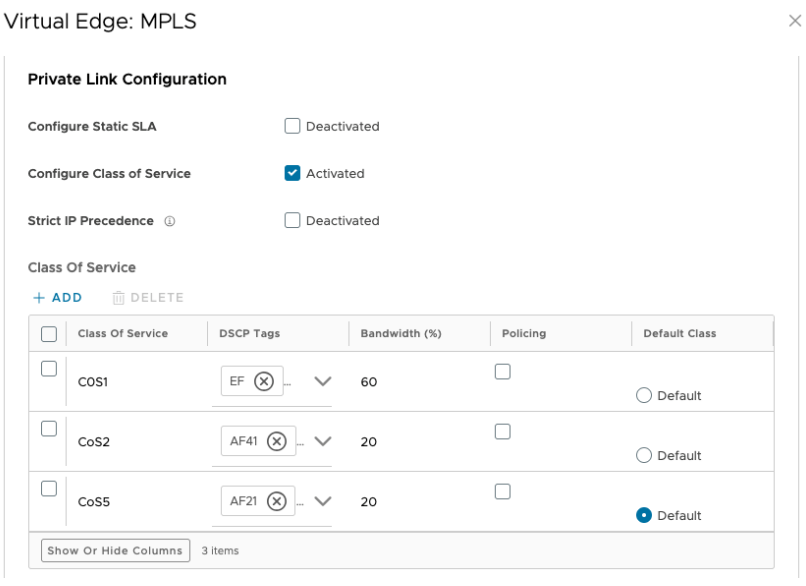

Figure 11: MPLS CoS agreement with three classes of service

In the Business Policy above, link steering options include:

- Internet
- MPLS - CoS1 (high-priority traffic)
- MPLS - CoS2 (medium-priority traffic)
- MPLS - CoS5 (low-priority traffic)

This allows traffic to be dynamically routed based on application requirements and network conditions, ensuring optimal performance and adherence to service provider SLAs.

Link steering by interface

The link steering policy can be applied to a specific interface (e.g., GE2, GE3), which varies depending on the VeloCloud SD-WAN Edge model and deployment location.

IT administrators must manually configure policies based on the physical connection layout of each Edge device. Since interface assignments differ across Edge models and locations, this approach lacks flexibility and scalability. For these reasons, interface-based steering is the least desirable option in the Business Policy, as it requires detailed knowledge of each Edge's connectivity setup.

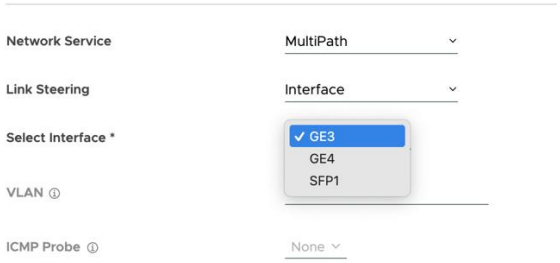

Figure 12: Link steering by interface

Traffic steering types

For link steering by transport group, interface, or WAN link, there are three steering sub-options:

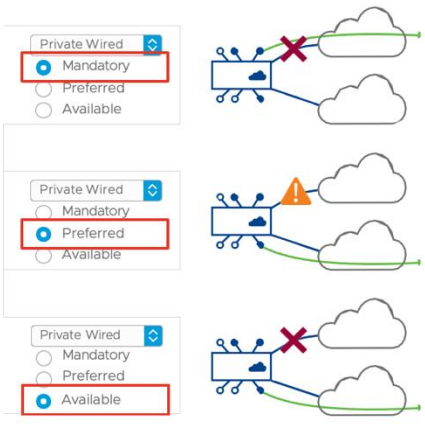

Mandatory: The application is pinned to a specific path, even if the link fails. (Example: PCI-compliant traffic, which must follow a designated secure path.)

Preferred: The application is prioritized on a specific path but will be redirected if the path cannot meet the required SLA. (Example: VoIP, which requires low latency and jitter.)

Available: The application prefers a specific path but will switch to another if the link fails. (Example: Web browsing, which can be redirected to any available connection.)

Figure 13: Sub-options for link steering by transport group, interface, or WAN link

Mandatory

Traffic is pinned to a specific link or transport group and will never be steered away, even if the link is experiencing degradation or an outage. On-demand remediation is enabled to mitigate brownout conditions such as packet loss and jitter.

Example: Netflix, a low-priority application, is required to always stay on public wired links, even if performance degrades.

Preferred

Traffic uses the preferred link as long as it meets the application's SLA. If the preferred link fails to meet the SLA, traffic is steered to another link. If no other link meets the SLA, on-demand remediation is applied instead of immediate steering. DMPO can continue using the degraded link with remediation until performance deteriorates beyond recoverable limits, at which point traffic is steered to a better link.

Example: Video collaboration applications prefer the Internet link but will switch to a private link if the Internet connection fails to meet video quality SLAs.

Available

Traffic uses the selected link as long as it is active, regardless of SLA performance. On-demand remediation is enabled if SLA conditions degrade, but traffic is only moved if the link completely fails.

Example: Web traffic is backhauled over the Internet link to a hub site, staying on the Internet link as long as it remains active, even if it does not meet SLA requirements.

Secure traffic transmission

For private or internal traffic, DMPO encrypts both the user traffic (payload) and the tunnel header using IPsec transport mode, ensuring end-to-end security.

Overhead varies from 59 bytes to 120 bytes based on the level of encryption.

Enterprise administrators have the control to set the encryption key. Through the ReST API this can be rotated on a schedule of the administrator's choosing. At the time the keys are rotated both the new key as well as the old key will be accepted to avoid any network disruption.

Encryption and security standards:

- **Encryption**: Supports AES-128 and AES-256 for data confidentiality.
- **Integrity**: Uses SHA-2 and SHA-1 algorithms to verify data integrity.
- **Key management**: Utilizes IKEv2 for secure key exchange.
- **Authentication**: Implements Public Key Infrastructure (PKI) for authentication.

This ensures secure and tamper-proof communication across the SD-WAN network.

Ports used

Both data and control traffic use UDP port 2426.

DMPO real-world results

Scenario 1: Branch-to-branch VoIP call on single link

The results in Figure 14 highlight the benefits of on-demand remediation, specifically Forward Error Correction (FEC) and jitter mitigation, on a single Internet link, comparing traditional WAN and VeloCloud SD-WAN performance.

A Mean Opinion Score (MOS) below 3.5 indicates unacceptable voice or video call quality. The figure illustrates how VeloCloud SD-WAN improves MOS scores, ensuring better call clarity and overall user experience.

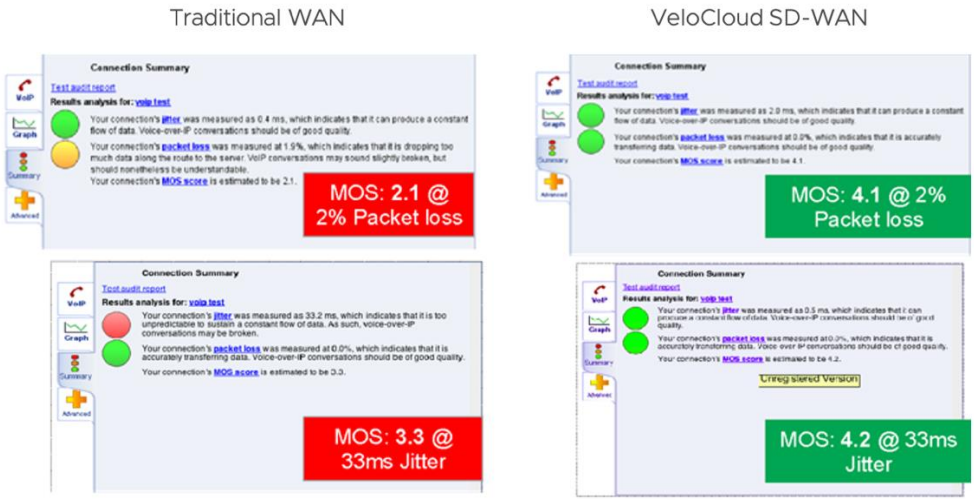

Figure 14: Results for a branch-to-branch VoIP Call over a single link optimized with DMPO

Scenario 2: File transfer from Box.com on dual links

The results in Figure 15 highlight the advantages of bandwidth aggregation and on-demand remediation when downloading a 50MB file from Box.com using two 20 Mbps links. The comparison between traditional WAN and VeloCloud SD-WAN demonstrates how SD-WAN optimizes throughput and improves download performance by intelligently utilizing both links.

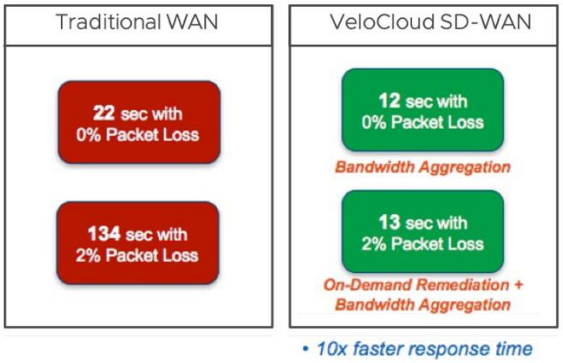

Figure 15: Results for a file transfer from Box.com on dual links with DMPO

Scenario 3: Branch-To-branch video call on dual links

The results in Figure 16 showcase the benefits of sub-second link outage protection in VeloCloud SD-WAN. When a link fails, application flows are seamlessly steered to an available Internet link, while on-demand remediation ensures optimal performance on the new path. This demonstrates how VeloCloud SD-WAN maintains application continuity and quality even during sudden network disruptions.

Figure 16: Results for branch-to-branch video call on dual links with DMPO

Summary

VeloCloud SD-WAN DMPO enhances network performance with three key capabilities:

- Application-aware, dynamic per-packet steering to optimize traffic flow in real time.
- On-demand remediation to mitigate packet loss, jitter, and latency issues.
- Overlay Quality of Service (QoS) to prioritize critical applications.

DMPO ensures high-performance SD-WAN connectivity for even the most demanding applications, across any transport (Internet or hybrid networks) and to any destination (on-premises or cloud).

For more information, visit the VeloCloud SD-WAN web page.

velocloud™
by Broadcom

Item No: VeloCloud-DMPO-wp-2025 Mar-25