

Versa SD-WAN over Starlink

Technical Validation of Dual-Terminal Load Balancing, Failover, and TCP Optimization

Table of Contents

1. Executive summary	2
<i>Who should read this</i>	2
<i>What this paper does not claim</i>	2
2. The Versa and Starlink solution	2
3. Topology and methodology	3
<i>Evidence sources</i>	4
<i>Terminology and operating rates</i>	5
4. Throughput & load balancing measurements with both terminals	5
5. Latency measurements with both Versa and Starlink	7
6. Using Starlink telemetry and Versa SLA probes	8
7. Measuring traffic behavior during failover and recovery	9
<i>How long the control plane took to notice</i>	11
8. Understanding when to use local breakout or SD-WAN tunnel for workloads	11
9. Forwarding profile behavior and TCP optimization	12
<i>When to use TCP optimization for a workload</i>	13
10. Configuration reference	14
<i>Path liveness</i>	14
<i>TCP optimization</i>	14
<i>Forward error correction, packet cloning and header compression</i>	15
11. Conclusions and configuration guidance	15

1. Executive summary

Starlink delivers broadband capacity and low-earth-orbit latency to sites where terrestrial circuits are unavailable, uneconomic, or slow to provision. In addition, workloads placed at those sites typically bring requirements that no single access link is designed to satisfy: use of more than one path at once, continuity when a path is lost, forwarding policy per application, encryption to a defined endpoint, and per-path measurement. Versa Secure SD-WAN, in conjunction with Starlink, offers a solution that delivers a new era of global digital connectivity.

In this technical report, two Starlink terminals were placed behind a single Versa branch appliance and measured across sixteen forwarding profiles, TCP and UDP, and in both directions. The analysis also included induced terminal outages, repeated over the SD-WAN tunnel and over local breakout. Four primary findings from this solution pairing are summarized below.

- Two Starlink terminals can be used to actively carry traffic at the same time, consistently enough to recommend provisioning a site using a pair when such a requirement is needed. Section 4.
- The Starlink terminal's telemetry when combined with the Versa SD-WAN tunnel's telemetry can help isolate which terminal and which layer a problem is on. Section 6.
- Loss of a Starlink terminal's availability is absorbed by the SD-WAN tunnel such that no interruption occurs for the tested SD-WAN tunnel flows. An application's impact still depends on protocol timers and session behavior. The same loss of terminal availability over local breakout flows ends established sessions permanently, because those sessions are anchored to the address of the terminal that went away. Section 7.
- Throughput using Starlink is dependent on the application requirements. While the SD-WAN tunnel delivers lower aggregate throughput than direct local breakout throughput, performance can be tuned to minimize impacts. Sections 8 and 9.

Who should read this

This paper is intended for enterprise architects, satellite connectivity teams, SD-WAN engineers, and customer facing technical teams designing resilient Starlink based branch connectivity.

What this paper does not claim

These results should not be treated as universal Starlink capacity numbers. They reflect controlled and measured results of a single site, terminals, time windows, forwarding profiles, hub placement, and clear sky conditions.

2. The Versa and Starlink solution

The measurements that follow compare forwarding profiles when a Versa SD-WAN appliance uses two Starlink terminals. Within Versa, a forwarding profile is a named combination of the features below, bound to traffic by policy. The paper assumes the reader has no prior familiarity with them, so each is described once with the behavior that matters on a satellite underlay.

Feature	What it does	Why it might matter on Starlink
SD-WAN tunnel forwarding	Encapsulates branch traffic in a tunnel to a defined endpoint, with optional encryption	Provides policy, encryption and per-path measurement that raw internet access cannot provide
Load balancing	Distributes sessions across more than one WAN path	Two terminals become one logical service instead of an active path and an inactive spare
TCP optimization	Terminates and proxies TCP at each end of the SD-WAN tunnel so the long path is handled by the appliance	Removes the round trip to the receiver from the sender's control loop, which is the dominant limit on a satellite path
Packet cloning	Duplicates packets across two paths and discards the later copy at the far end	Masks packet loss before a transport protocol interprets it as congestion, at roughly double the WAN bandwidth cost
Forward error correction	Adds redundancy so a receiver can reconstruct lost packets without retransmission	Targets recovery on a lossy path where a retransmission would cost a full round trip
Header compression	Reduces per-packet header overhead inside the tunnel	Improves the payload to overhead ratio, which matters most for small packets
Per-packet load balancing	Spreads individual packets of one flow across multiple paths	Allows a single high-rate UDP stream to exceed the capacity of one terminal. Not suitable for TCP
Local breakout	Sends traffic straight to the Internet from the branch instead of through the SD-WAN tunnel	Shortest path for internet-destined traffic, at the cost of the SD-WAN tunnel's policy and resilience

Table 1. The forwarding features combined into the profiles measured in this paper.

3. Topology and methodology

Two independent Starlink terminals were attached to a single Versa branch appliance as SD-WAN path candidates toward a Versa SD-WAN hub deployed in a public cloud region. A Linux test host behind the branch generated iperf3 traffic. A forwarding profile was selected by destination service port, so a given test port maps to exactly one policy and one feature combination. That is what makes a profile comparison possible without reconfiguring the appliance between runs.




Figure 1. Test topology.

Two live Starlink terminals were actively available for this validation, so all measured results in this paper use two paths by policy. Note: The Versa Networks solution is not limited to two Starlink terminals or connections. A single Versa VOS instance can support up to 15 WAN links for SD-WAN traffic. For direct internet access and local breakout designs requiring additional physical WAN links, VOS instances can be cascaded as aggregation nodes, allowing the architecture to scale beyond that limit. The same multi-link design principles are used in production multi orbit deployments.

Evidence sources

The validation used three primary measurement sources and one correlation workflow.

- iperf3 JSON provides receiver payload throughput at one second resolution for each run and direction. This is the authority for delivered application throughput.
- Starlink terminal telemetry was collected per terminal through gRPC and stored as time series data. It provided reported rate in each direction, terminal to point of presence latency, loss, and obstruction fraction. These counters describe the terminal as a whole and are not per policy or per forwarding class counters.
- Versa SLA measurements and path state events were exported through syslog and parsed into event records. They provided delay, jitter, loss, and path state separated by SD-WAN path and forwarding class. The forwarding class carrying test traffic used 1 second probes with a loss threshold of 10 probes. A second higher priority class with no test traffic used 2 second probes with a threshold of 3 probes.

The automation orchestrator created the run ID and UTC time window, read the iperf JSON, queried the Starlink time series, read the Versa SLA event records, and produced correlated HTML, CSV, and JSON reports.




Figure 2. Starlink Versa IOT data collection

Terminology and operating rates

Item	Value
Access service	Two Starlink terminals on the same premises, referred to as terminal 1 and terminal 2
Overlay	Versa branch appliance to Versa SD-WAN hub in a public cloud region
Traffic generator	iperf3. TCP runs use default payload length; UDP load-balancing runs use a 1340 byte payload
Run length	120 seconds per profile per direction, except the failover runs at approximately 334 seconds
Runs per data point	1, except the distribution result in Section 4 which is 6 independent runs
Time of day	Afternoon, evening and overnight local time. Starlink capacity varies with time of day, so these figures are not representative of behavior or performance for all hours
Obstruction	Fraction obstructed held at or below 0.013 on both terminals for every run, which is a clear-sky condition. Nothing here should be extrapolated to a site with partial sky visibility

Table 2. Terminology and operating rates.

4. Throughput & load balancing measurements with both terminals

Six independent runs were performed spanning TCP and UDP, in both directions, with encrypted and unencrypted forwarding. The outcomes of these runs placed the split between the two terminals inside a 43.8 to 56.2 percent band as shown in Figure 3.

Both terminals carried traffic in every load balancing run

Six independent runs across TCP and UDP, both directions, encrypted and unencrypted. Dotted lines mark the 40 to 60 percent acceptance band.




Figure 3. Measured share of delivered traffic across the two terminals in each load-balancing run.

The makeup of each individual run shows that both terminals ramp up operation rates, settle at a comparable rate, and hold these rates for the remainder of the test. Neither is idle and neither is carrying a majority of the traffic, as shown in Figure 4.

Both terminals carry comparable rates at the same time

Two independent 40-stream UDP load-balancing runs. Each line is that terminal's own reported rate; the bold line is their sum.




Figure 4. Per-terminal reported rate and their sum through two independent 40-stream UDP load-balancing runs.

Delivered throughput offered load closely where the offered target was sized inside available capacity, and it repeated in each run. Two independent 40-stream UDP downloads delivered 278.3 and 282.1 Mbps of a 300 Mbps target, with per-terminal means of 172.2 and 148.8 Mbps in one and 183.0 and 150.2 Mbps in the other. A 32-stream UDP upload offering 32 Mbps delivered 30.9 Mbps at the receiver with 3.3 percent packet loss and 3.19 ms jitter.

Sizing the throughput target to the direction under operation matters, because uplink capacity is materially lower than downlink on this access service. An earlier trial run offering 55 Mbps across 40 upload streams exceeded available uplink capacity and produced a correspondingly uneven split; once the throughput load was set inside capacity, upload distribution behaved like download distribution. Acceptance criteria for a load-balancing deployment should therefore be defined per direction rather than as a single figure. UDP runs also took 20 to 30 seconds to settle at the expected throughput rate, so a warm-up window should be excluded from any acceptance measurement.

Per-packet load balancing is the exception to flow-based distribution and is scoped narrowly. It distributed a single UDP stream across both terminals at 457.9 Mbps with a 56 / 44 split. Because it distributes individual packets rather than flows it reorders by design, this suits UDP and does not suit TCP.

Mixed transport designs should be treated differently from the dual Starlink case measured here. When Starlink is combined with another path which has a materially different round trip time, such as MEO, GEO, or a terrestrial service, do not treat the paths as equal load balancing candidates by default. Assign path priority and use forwarding policy so traffic prefers the path that best matches the application requirement, with the other transport used for selected traffic or resilience. This avoids unnecessary reordering and timing effects that can appear when paths with very different latency carry the same traffic.

5. Latency measurements with both Versa and Starlink

Two latency measurements were collected in parallel throughout this testing. They sit in different places and they report different things, and the difference is the most useful observation in the dataset.

- The Starlink terminal measures its own round trip to its point of presence. It covers the access hop only, and the terminal generates the probe itself.
- The Versa SLA probe measures the full round trip from branch to the SD-WAN hub. It rides inside the tunnel, so it shares a queue with the traffic the tunnel is carrying.

Tunnel delay tracks throughput, and the priority class gets no protection

Same runs, same terminal. The access hop holds its range throughout while the tunnel through it queues.




Figure 5. Tunnel delay against delivered throughput on the loaded terminal, with the access hop shown for comparison.

As traffic load on the Starlink path increased, the SD-WAN tunnel RTT increased with it. In Figure 5, the red line shows that tunnel RTT rose from 77 ms at approximately 17 Mbps to 121 ms at approximately 220 Mbps. Over the same runs, the Starlink terminal's own access-latency measurement remained between 17 and 36 ms.

The measurements are consistent with queueing on the loaded access path as utilization increased. The idle terminal, which used the same SD-WAN hub but was not carrying the test load, remained between 82 and 90 ms and provides a useful baseline.

The orange line in Figure 5 represents the higher priority forwarding class on the loaded terminal. That class carried no test traffic, yet its RTT increased along with the traffic carrying class, reaching 126 ms compared with 121 ms for the loaded traffic class. This indicates that the higher priority marking did not result in differentiated treatment across the measured access service. In this test, end to end underlay QoS treatment was not observed. Latency sensitive traffic therefore needs to be protected at the branch through per class queueing and egress shaping set below the terminal's usable rate.

6. Using Starlink telemetry and Versa SLA probes

Starlink terminal telemetry and Versa SLA probe measurements observe different parts of the path, so using them together helps isolate where a performance issue is occurring. For this analysis, Starlink terminal telemetry and Versa SLA probe measurements were collected separately and aligned on a common elapsed-time axis to correlate access behavior with end-to-end application performance.

In two UDP load-balancing runs, one terminal experienced reduced delivered throughput. By aligning the application results with the Starlink terminal telemetry and Versa SLA probe measurements, we were able to determine whether the issue was associated with the access link or with the SD-WAN tunnel path.

Correlating the two telemetry sources on a run that lost traffic

40 UDP streams, 300 Mbps offered across two terminals. Shaded band marks the window in which the application lost traffic.

Top two panels are per second from the traffic itself. Lower two are terminal telemetry, polled about every 10 seconds per terminal.




Figure 6. A 40-stream UDP load-balancing run that lost traffic, with both telemetry sources on one elapsed-time axis.

The application experienced lower throughput and higher packet loss for the first 28 seconds of this run. Receiver packet loss peaked at 67 percent at second 26 and delivered throughput fell to 96.5 Mbps against a 300 Mbps target throughput. Both performance and packet transmission recovered, and the run held its target for the remaining 90 seconds. The lower two graphs of Figure 6 explain why. Terminal 2 reported access latency of 66 ms against 26 ms on terminal 1, and its reported downlink was still at 44 Mbps at second 33 while terminal 1 had already reached 178 Mbps. Both the elevated access latency and the reduced reported downlink on terminal 2 resolve on the same timeline as the application loss cleared. The access link on that terminal was degraded and its own telemetry validated this state.

The Versa SLA probes recorded the same event from the other side, at 49.0 percent forward loss on terminal 2 against 15.2 on terminal 1. That is the pairing that matters. When both instruments move together, this indicates an access link degraded, which is an obstruction, hardware, or capacity problem at the terminal. When the Versa SLA measurement degrades while Starlink latency to the PoP remains normal, the issue is likely elsewhere along the SD-WAN path, such as branch shaping, transport congestion, or queueing at the hub or remote endpoint.

7. Measuring traffic behavior during failover and recovery

To understand the impact of a Starlink terminal outage, a steady group of eight UDP flows, offered at 12 Mbps in total, was held active over the SD-WAN tunnel while one terminal was administratively disabled for 60 seconds and then restored. Rolling flow groups arrived every 10 seconds throughout, taking the offered ceiling for the run to 120 Mbps. Placement of new sessions during and after the outage was exercised as well as continuity of the ones already running. The same outage was then applied with the same traffic pattern over local breakout on the same links. Figure 7 visualizes these results.

SD-WAN overlay: the terminal drops out, the application does not

Terminal 2 administratively disabled at 90 s and restored at 150 s. Terminal telemetry is sampled roughly every 10 seconds per terminal.




Figure 7. SD-WAN tunnel: delivered payload and per-terminal reported throughput on one elapsed-time axis.

For flows using the SD-WAN tunnel, the result is best described by what did not happen. The delivered payload held steady at its 120 Mbps ceiling straight through both terminal state changes. All eight steady flow identifiers carried data throughout and were still carrying it in the final measurement intervals. The run records full-flow recovery after restoration at 0.0 seconds with zero automatic replacement attempts, and all 26 rolling groups passed, 13 of them having started after restoration and re-established across both terminals.

The lower graph in Figure 7 shows how the Versa SD-WAN device redistributed traffic across the two Starlink terminals during the outage. Before the outage, terminal 1 was carrying about 68 Mbps alongside terminal 2. When terminal 2 was disabled, the Versa device shifted the affected traffic to terminal 1, increasing its throughput to about 136 Mbps for the duration of the 60-second outage. The delivered payload in the upper graph remained stable throughout.

After terminal 2 was restored, the Versa device gradually redistributed traffic across both terminals again. Traffic on terminal 2 increased while traffic on terminal 1 decreased as the load was rebalanced. The affected path reported 100 percent PDU loss during the outage, confirming that the path was fully unavailable. Traffic continued because the surviving terminal carried the affected traffic.

Local breakout: the same outage, and the sessions on that terminal are gone

Same 60-second outage on the same links. Aggregate throughput falls and recovers only as new sessions are created.




Figure 8. Local breakout: the same two measurements through the same 60-second outage.

As shown in Figure 8, local breakout sessions behaved differently from SD-WAN tunnel sessions during the terminal outage. With local breakout, each session was established using the public IP address associated with the Starlink terminal used for egress. When that terminal became unavailable, that source address was no longer reachable, so the existing session could not simply move to the other terminal and continue. The session ended and had to be re-established using the surviving path.

SD-WAN tunnel sessions behave differently because they are carried through the SD-WAN tunnel and are not tied to the public IP address of a specific Starlink terminal. When one terminal becomes unavailable, the Versa SD-WAN device can move the traffic to the surviving path while preserving the application session.

Of the ten sessions exposed to the shutdown, six continued and four were interrupted and did not recover. The throughput dip was caused by the sessions that had been using the failed terminal and could no longer continue with the same source address. While terminal 2 was unavailable, VOS continued forwarding new local breakout sessions through the surviving terminal. Twelve new sessions were established successfully during the outage. After terminal 2 was restored, VOS again distributed new sessions across the available paths, and 24 subsequent sessions completed successfully.

How long the control plane took to notice

The outage was induced through the Versa Director API, which is an administrative change that has to be committed on the appliance before the interface state changes. Measurements from the point at which the API call returned to the point where the SD-WAN data path to the affected terminal was declared down was approximately 13 seconds and approximately 8 seconds after undoing the state change. Roughly 10 seconds of the configured detection window and the remainder is commit and reporting time.

The class carrying the traffic is probed every 1 second with a loss threshold of 10 probes. Ten one second probes provide a standard detection window that tolerates short transient loss while maintaining a stable path state. Loss itself remains visible through the per path loss, delay and jitter measurements.

8. Understanding when to use local breakout or SD-WAN tunnel for workloads

Local breakout is a recommended design for a large share of branch traffic that does not need to go to the hub, and is measured here on the same links as a first-class configuration. At high load all three configurations below distributed flows across both terminals, so the differences between them are attributable to forwarding configuration.

Most of the measured gap is associated with overlay forwarding, not encryption

128 concurrent TCP sessions; same destination, underlay path, and RTT in all three cases.




Figure 9. Aggregate payload throughput by forwarding configuration, 128 concurrent TCP sessions.

As shown in Figure 9, local breakout delivered 305.0 Mbps download, the unencrypted SD-WAN tunnel delivered 249.4 Mbps, and the encrypted SD-WAN tunnel delivered 229.4 Mbps. Upload delivered 83.8, 60.6 and 61.6 Mbps respectively. The destination, underlay path, and round-trip time were the same across these tests. The primary variable was the forwarding mode. The largest measured difference was between local breakout and SD-WAN tunnel forwarding. Adding encryption for the SD-WAN tunnel produced a smaller incremental difference on download and no measurable reduction on upload. These later two results therefore separate the cost of SD-WAN tunnel forwarding from the incremental cost of encryption without attributing the difference to a change in network path.

The takeaway is therefore not which forwarding mode is universally the best, but which capabilities each application requires. Local breakout provides the most direct forwarding model for internet-destined traffic that does not require special handling provided by the SD-WAN tunnel. SD-WAN tunnel forwarding adds path selection, per-path SLA probe measurement, and session continuity across WAN path changes. Encryption can then be added where policy requires protection to the SD-WAN hub. Sites can use more than one of these modes at the same time, selected by application policy.

Note: An application that already encrypts itself end to end does not need the SD-WAN tunnel to encrypt it a second time. That traffic can take the unencrypted SD-WAN tunnel and maintain other path policy capabilities without paying the encryption performance penalty twice, subject to whatever the security policy requires of traffic in transit. Most sites will end up running all three positions at once, selected by policy.

9. Forwarding profile behavior and TCP optimization

To evaluate forwarding profile behavior and TCP optimization with Starlink terminals, flows were configured with one TCP stream per forwarding profile, lasting 120 seconds each, directed to one terminal, with download and upload measured in separate runs on the same day. A single stream is the workload shape where SD-WAN policy features would be expected to assist most, because a single congestion window on a long path is typically the problem condition SD-WAN policy features address.

TCP optimization is the one feature that changes a single flow

One TCP stream per profile, 120 s each, one terminal. Download and upload measured in separate runs on the same day.




Figure 10. Single-flow TCP payload throughput by forwarding profile, one terminal.

TCP optimization was enabled for selected flows through the matching SD-WAN rule. When TCP optimization is active, VOS terminates and proxies the TCP session. As a result, the original sender does not see loss on the full end-to-end path in the same way and, in these tests, reported zero TCP retransmissions.

TCP optimization produced the largest measured change in single-flow throughput. As shown in figure 10 on the encrypted forwarding profile, a flow matching an SD-WAN rule without TCP optimization delivered 15.1 Mbps download, while the equivalent flow with TCP optimization enabled delivered 129.8 Mbps. On the unencrypted profile, throughput increased from 8.4 Mbps to 121.1 Mbps. Upload showed the same pattern, increasing from 4.8 Mbps to 24.3 Mbps on the encrypted profile.

Forward error correction did not produce a measurable single-flow throughput improvement over its baseline in either direction. Packet cloning also did not consistently improve single-flow TCP throughput across the measured runs. That is expected because packet cloning is primarily a resiliency mechanism: it sends duplicate packets across multiple paths to improve the probability that one copy arrives successfully. It does not create additional TCP congestion windows or increase the throughput potential of a single TCP flow. Packet cloning is therefore best used selectively for small volumes of loss-sensitive traffic, with the tradeoff of approximately double the WAN bandwidth consumption.

When to use TCP optimization for a workload

TCP optimization should be applied selectively, not to all traffic. Its main value is for applications that use one or a small number of long-lived TCP flows and would otherwise be unable to make full use of the available bandwidth. When a site already carries many concurrent TCP sessions, those sessions collectively create enough independent congestion windows to fill the available link without TCP optimization. In that case, enabling TCP optimization broadly provides little additional benefit.

The practical approach is therefore to identify the specific flows that matter most and apply TCP optimization only where it improves application performance, such as selected file transfers, backups, replication jobs, or software distribution.

Concurrency reaches the same place without the feature

One terminal throughput. The dashed line is the single optimized session from Figure 8, measured on the same terminal.




Figure 11. Throughput against concurrent session count with no optimization, against the single optimized session from Figure 10.

The measurements in Figure 11 show this clearly. On download, 16 unoptimized sessions delivered 111.2 Mbps compared with 129.8 Mbps for one optimized session, while 64 unoptimized sessions reached 167.5 Mbps. On upload, the unoptimized traffic exceeded the optimized single-flow result between 4 and 16 sessions.

This shows that TCP optimization is not required simply to fill the available link when the workload already contains enough concurrent sessions. Its value is in accelerating selected flows that would otherwise be constrained individually.

10. Configuration reference

Every measurement in this paper was produced using the configuration documented in this section. These values are provided to make the validation reproducible. They should be treated as the settings used in this test, not as universal recommendations for every Starlink deployment.

Path liveness

For the forwarding class carrying the test traffic, Versa SLA probes were sent every 1 second with a loss threshold of 10 consecutive probes before declaring the path down. This created an approximately 10 second detection window.

That setting was used in this validation to avoid unnecessary path state changes during short transient loss while continuing to measure per path loss, delay, and jitter. The remaining path liveness parameters used their default values.

TCP optimization

TCP optimization is applied within a SD-WAN rule rather than in a forwarding profile, in mode Auto, referencing one LAN-side and one WAN-side TCP profile. Because it is bound at the rule, a rule named for a baseline will still be optimized if the configuration is present. This can be verified with the sender retransmission count as described in Section 9.

Parameter	WAN profile	LAN profile
Congestion control	bbr	cubic
Loss detection	rack	duplicate-ack
Loss recovery	proportional-rate-reduction	default
Send buffer	4096	128
Receive buffer	4096	128
Rate pacing	Enabled, with auto rate pacing limit	Disabled
Hybrid slow start	Disabled	Disabled

Table 3. TCP optimization profiles. The WAN profile faces the satellite path and carries the settings that matter.

- TCP Optimization operates on the TCP header, which is unencrypted even when the application payload is TLS encrypted.
- Concurrent optimized sessions are a bounded platform resource that scales with the appliance. On a small branch appliance it is in the low hundreds, so the feature should be scoped by policy to the traffic that needs it rather than enabled for all traffic.
- It must be configured at both ends of the SD-WAN tunnel. On a shared SD-WAN hub the policy can be scoped to a defined set of remote sites, so enabling it for one site does not change behavior for the others.

Forward error correction, packet cloning and header compression

Feature	Parameter	Value used
Forward error correction	Packets per FEC packet	8, giving roughly 12.5 percent overhead. The default of 4 costs 25 percent, which is more than Starlink loss rates in this environment justify
Forward error correction	FEC packet placement	Alternate circuit, with the duplicate on the same circuit
Forward error correction	Maximum FEC packet size	1400 bytes
Forward error correction	Receiver	Recovery enabled, preserve order enabled
Packet cloning	Replication factor	2, matching the two terminals
Header compression	Level	High, HMAC skipped, unidirectional

Table 4. Forwarding profile parameters.

One design constraint on packet cloning is that the receiver holds a finite reordering window while it waits for the second copy of a packet. That works when the two paths have comparable delay, which two Starlink terminals do. Pairing a satellite path with a terrestrial path whose delay is an order of magnitude lower defeats the packet cloning mechanism, because the late copy arrives after the window has closed. Forward error correction has no such constraint, since the redundancy travels on the same path as the data.

11. Conclusions and configuration guidance

Objective	Recommended configuration
Use both terminals as one service	Perform SD-WAN tunnel load balancing across both terminals, with per-path acceptance criteria defined separately for each direction and the offered target sized inside available capacity for that direction
Choose forwarding by application requirement	Use local breakout where direct internet forwarding is the priority. Use SD-WAN tunnel forwarding where path selection, SLA visibility, session continuity, or connectivity to a defined SD-WAN hub is required. Apply encryption according to security policy. Use TCP optimization selectively for prioritized workloads that depend on one or a small number of TCP flows.
Scope TCP optimization	Enable it by policy for single large transfers, backups and replication. Do not enable it site-wide: a workload of many concurrent sessions reaches the same throughput without it.
Choose between SD-WAN tunnel and local breakout	Use local breakout for internet-destined bulk traffic at attended sites. SD-WAN tunnel for anything bound for headquarters or the data center, for business-critical applications, and wherever sessions must survive the loss of a terminal, because a local breakout session is anchored to the address of the terminal it left through and an SD-WAN tunnel session is not. Most sites run both, selected by policy

Protect latency-sensitive traffic	Higher-priority marking did not produce differentiated treatment on the access service, so protect latency-sensitive traffic through branch-side queuing and egress shaping.
Monitoring a site	Establish a per-path SLA measurement on the SD-WAN tunnel alongside terminal telemetry, not alone. Both measurements moving together means the Starlink terminal is degraded. When the Versa SLA measurement degrades while Starlink latency to the PoP remains normal, the issue is likely elsewhere along the SD-WAN path, such as branch shaping, transport congestion, or queueing at the hub or remote endpoint.
Set expectations on failover	Set seconds for the control plane to report, and no interruption to sessions in flight. Detection follows the configured SLA probe interval and loss threshold, which was 1 second and 10 probes
Protect a small volume of loss-intolerant traffic	Enable packet cloning or FEC, scoped narrowly by policy. Confirm behavior on the specific traffic
Combine Starlink with unlike transports	Use path priority and application policy when combining Starlink with MEO, GEO, or other terrestrial links whose latency differs materially. Do not treat unlike paths as equal load balancing candidates by default

Table 5. Configuration guidance derived from the measured results.



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

About Versa

Versa, the global leader in unified networking and security, enables organizations to create self-protecting networks that radically simplify and automate their network and security infrastructure. Powered by AI, the VersaONE Universal SASE Platform delivers converged SSE, SD-WAN, and SD-LAN solutions that protect data and defend against cyberthreats while providing a superior digital experience. Thousands of customers globally, with hundreds of thousands of sites and millions of users, trust Versa with their mission-critical networks and security. Versa is privately held and funded by investors such as Sequoia Capital, Mayfield, and BlackRock. For more information, visit <https://www.versa-networks.com> and follow Versa on LinkedIn and X (Twitter) @versanetworks.

©2026 Versa Networks, Inc. All rights reserved. Portions of Versa products are protected under Versa patents, as well as patents pending. Versa Networks and FlexVNF are trademarks or registered trademarks of Versa Networks, Inc. All other trademarks used or mentioned herein belong to their respective owners.