

Post-Quantum Login and Authentication Problems in SD-WAN

Afshan Mulla · Technical Leader, Cisco CX, Business Critical Services (BCS) · afsmulla@cisco.com

Jaypal Baviskar · Technical Leader, Cisco CX, Service Provider Technology Group (SPTG) · jbaviska@cisco.com

Position paper- IAB Workshop on Accelerating the Deployment of Post-Quantum Authentication (pqws), Prague, 11–12 October 2026

Position. SD-WAN is the hardest place in the enterprise to deploy post-quantum authentication, and it is where the consequences of getting it wrong are largest. Three things combine.

First, every SD-WAN router proves who it is over DTLS on UDP, where one lost datagram forces the whole handshake flight to be sent again and post-quantum certificates turn a 4-datagram handshake into a 36-datagram one.

Second, SD-WAN runs on the diverse links available: broadband, LTE and satellite, exactly where loss and delay are highest.

Third, SD-WAN does not use IKE. Data-plane encryption keys are generated by each router and distributed by the controller over that same control connection, so the security of every tunnel in the fabric rests on one certificate exchange.

This paper shows that with the algorithm CNSA 2.0 mandates, only about a third of control-connection attempts on a 3%-loss circuit would complete without retransmission, against nine in ten today. We propose three fixes that need no new cryptography.

1. How SD-WAN login actually works, and why it is different

A WAN Edge router joins the fabric by building mutually authenticated control connections to the orchestrator, the controllers and the manager. Cisco Catalyst SD-WAN documents the design: control connections run over DTLS on UDP by default, TLS on TCP is an option, and each device authenticates with a certificate while both ends also check the peer's serial and chassis number against an authorized-device file the controller has already distributed [3]. Hardware routers hold their factory identity in a tamper-resistant Trust Anchor chip that carries a pre-programmed SUDI certificate. Controllers today use RSA-2048 keys [3].

A single router does not build one control connection. It builds one persistent connection to the manager, one to each controller over every transport, and transient connections to the orchestrator over every transport during bring-up [4]. With two transports and two controllers that is five persistent connections and seven handshakes at cold start - per router, per reload, and again after every transport flap.

The last difference is the one that matters most. SD-WAN *dispenses with IKE altogether*. Each router generates its own AES key and sends it to the controller inside OMP route updates, and the controller distributes it to the other routers [3]. That removes an enormous amount of pairwise key-exchange work, which is why the design is good. It also means the confidentiality of every tunnel in the overlay depends on the control connection being authenticated correctly. An attacker who can forge a controller identity is not attacking one tunnel; they are handed the fabric.

2. What a post-quantum control connection costs

How this was worked out. Key and signature sizes are from FIPS 204 [1] and FIPS 205 [2]. Each certificate carries an assumed 500 bytes of X.509 structure beyond the key and signature, which is low for enterprise templates, measured minimal certificates from OpenSSL 3.5 are 3,987 / 5,516 / 7,474 bytes for ML-DSA-44 / 65 / 87 [12]. Each side sends a leaf plus one intermediate. Both sides authenticate. DTLS handshake records are assumed to fit 1,200 bytes, which is generous once NAT and tunnel overhead are counted. "Clean" is the chance that a handshake completes with no retransmission at all, which is $(1 - \text{loss})^{\text{datagrams}}$. DTLS retransmits an entire flight, not the lost datagram, so a single loss re-sends about 20 kB.

Signature algorithm	Controller sends	Edge sends	Total bytes	Datagrams	Clean @1% loss	Clean @3% loss
RSA-2048 (today)	2,388	2,356	4,744	4	96%	89%
ML-DSA-44	12,068	10,884	22,952	21	81%	53%
ML-DSA-65	16,015	14,831	30,846	27	76%	44%
ML-DSA-87 (CNSA 2.0)	21,633	20,065	41,698	36	70%	33%
SLH-DSA-128s	25,816	24,632	50,448	43	65%	27%

Table 1. One mutual DTLS control-connection handshake, in bytes and UDP datagrams, and the chance it completes without a retransmission on a lossy circuit.

- Loss stops being a nuisance and becomes the binding constraint.** A broadband or LTE TLOC at 3% loss today completes 89% of handshakes cleanly. At ML-DSA-87 that falls to 33%. Two out of three control connections would need at least one full flight retransmitted, and each retransmission re-sends about 20 kB down the link that just dropped a packet.
- A controller restart becomes a self-sustaining storm.** Take a 5,000-site fleet with five persistent control connections each. Re-homing after a controller upgrade means 25,000 handshakes: about 118 MB and 100,000 datagrams today, against roughly 1.04 GB and 900,000 datagrams at ML-DSA-87. With two thirds of those handshakes retransmitting at least once, the herd grows faster than it drains. Signature verification is not the problem, ML-DSA verifies in well under a millisecond on commodity hardware. The wire and the loss are the problem.
- Satellite and high-latency TLOCs stop working first.** DTLS retransmission is timer-based with exponential backoff. On a 600 ms path, one lost datagram in a 36-datagram handshake adds at least a second before anything is resent, and a second loss adds several more. A connection that comes up in under a second today can take tens of seconds, which sits badly with the timers driving fabric convergence.

3. Where the limits sit in an SD-WAN fabric

Where	What it is	Can it be changed?	What it means
DTLS over UDP (default control transport)	One lost datagram retransmits the whole flight; no reliable stream underneath [3][5]	Config: TLS/TCP instead	Handshake stalls, control connection flaps
Control connections per router	One to the manager, one to each controller per transport, plus transient orchestrator connections on every transport [4]	Architectural	Handshake cost multiplies five- to sevenfold per device
Transport quality	Broadband, LTE/5G and satellite TLOCs with 1–5% loss and up to 600 ms RTT	No, that is the point of SD-WAN	Loss and delay multiply the flight count
Controller certificate keys	RSA-2048 today [3]	Yes, once FIPS-validated PQ modules ship	RSA-2048 is disallowed by policy well before the fleet refreshes
SUDI in the Trust Anchor chip	Certificate pre-programmed at manufacture [3][14]	Never	Factory identity stays classical for the device's 7–10 year life
Data-plane keys	AES keys generated per router and distributed in OMP routes over the control connection; IKE is not used [3][16]	Architectural	The whole fabric's tunnel keys rest on one control-plane identity

Where	What it is	Can it be changed?	What it means
Authorized serial/chassis file	Already distributed by the controller to every device [3]	Yes - the opening	Proves the fabric can already ship trust material out of band

Table 2. SD-WAN-specific pressure points. The last row is not a problem; it is the mechanism the proposals below build on.

4. What we propose

(1) Stop putting the chain on the wire. SD-WAN already has the channel. Every SD-WAN fabric already distributes an authorized-device file from the controller to every router before any handshake happens [3]. The same channel can carry intermediates and trust anchors, so the handshake sends a short reference instead of a chain. Cached information [7], raw public keys [8] and trust-anchor negotiation [9] are barely deployed on the public web because nobody can coordinate both ends. In SD-WAN both ends and the CA belong to one organisation, so every objection disappears. Removing one intermediate from both directions saves 15.4 kB of the 41.7 kB in Table 1 : 37% and takes the handshake from 36 datagrams to 23, which lifts the clean-completion rate at 3% loss from 33% to 50%. Raw public keys would remove most of what is left.

(2) Make the handshake survive a flap. LTE and broadband TLOCs flap routinely, and every flap costs a full mutual handshake today. Post-quantum makes that fifteen times more expensive at the exact moment the link is at its worst. SD-WAN needs resumption that survives a transport change and a device reload, so the full certificate exchange happens once per device per epoch rather than once per flap. This is the difference between a fleet that recovers from a brownout and one that does not.

(3) Make the failure attributable, and account for it per TLOC. Control-connection state today reports a generic failure. An operator watching a post-quantum rollout cannot distinguish "the certificate is too big for this circuit" from "the circuit is broken", and will reasonably suspect the circuit. Add an explicit state for a handshake that exceeded what the transport carried, and report handshake bytes, datagram count and retransmissions per TLOC in telemetry. This is small, and it decides whether a migration takes weeks or quarters.

5. Questions for the workshop, and where the work belongs

- What is the real datagram-loss and RTT distribution on production SD-WAN broadband, LTE and satellite TLOCs, and what handshake size does it actually support? Vendors and managed service providers hold this telemetry. Nobody has published it.
- Should SD-WAN control planes change their default from DTLS to TLS for the post-quantum era, accepting TCP head-of-line blocking in exchange for reliable delivery of large flights? This is a live design question with no current answer.
- Because data-plane keys ride the control connection, control-plane authentication protects far more than itself. Does that argue for a stronger parameter set on the control plane, or for reducing the coupling?
- Factory SUDI certificates cannot be replaced and the routers holding them will be in service for 7–10 years. What bridges a classical factory identity to a post-quantum fabric identity, and who signs it?
- In multi-tenant managed SD-WAN, tenants will migrate at different rates. Must a fabric support mixed classical and post-quantum control connections indefinitely, and what does that do to the authorized-device model and to composite certificate sizes [15]?

Where the follow-up belongs. We have multiple IETF groups who look after: DTLS resumption and trust-anchor negotiation in TLS, certificate-size work in LAMPS and PLANTS, and a per-TLOC handshake telemetry model in OPSAWG. We need an attention of respective working groups on SD-WAN control-plane security. SD-WAN service definitions sit with MEF, and factory identity needs liaison with IEEE 802.1 and TCG. The precedent is worth noting: the EAP community documented the same large-certificate problem years ago [17] and it is still unresolved. SD-WAN has less margin than EAP and a shorter policy deadline [11], so it should not wait its turn.

References

- [1] FIPS 204, Module-Lattice-Based Digital Signature Standard, NIST 2024.
- [2] FIPS 205, Stateless Hash-Based Digital Signature Standard, NIST 2024.
- [3] Cisco, Catalyst SD-WAN Security Configuration Guide, IOS XE 17.x - Security Overview.
- [4] Cisco, Catalyst SD-WAN Getting Started Guide - Overlay Network Bring-Up Process.
- [5] RFC 9147, DTLS 1.3, 2022.
- [6] RFC 6347, DTLS 1.2, 2012.
- [7] RFC 7924, TLS Cached Information Extension, 2016.
- [8] RFC 7250, Using Raw Public Keys in TLS and DTLS, 2014.
- [9] draft-ietf-tls-trust-anchor-ids-04, May 2026.
- [10] draft-ietf-plants-merkle-tree-certs-05, Jul 2026.
- [11] NSA, Commercial National Security Algorithm Suite 2.0, updated 2024-2025.
- [12] Cisco, The Great Certificate Size Explosion (IPsec Series 7), Aug 2026.
- [13] Cisco, Mutual Post-Quantum Auth over IKEv2 (IPsec Series 8), Aug 2026.
- [14] IEEE Std 802.1AR, Secure Device Identity, 2018.
- [15] draft-ietf-lamps-pq-composite-sigs-19.
- [16] Cisco, Catalyst SD-WAN Security Configuration Guide, IOS XE 17.x - IPsec Pairwise Keys.
- [17] RFC 9191, Handling Large Certificates and Long Certificate Chains in TLS-Based EAP Methods, 2022.