

Post-Quantum login and authentication in SD-WAN

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

Afshan Mulla and Jaypal Baviskar,
Technical Leader, Cisco CX. (views our own)

Three facts about how SD-WAN logs in

1

Control connections run DTLS 1.2 over UDP by default

Cisco's IOS XE 17.x security guide (RFC 6347). TLS 1.2 over TCP is an option. Certificates use 2048-bit RSA keys.

2

The transport is deliberately the worst available

Broadband, LTE and satellite. That is the product. Loss and delay are the operating condition, not the exception.

3

Catalyst SD-WAN “dispenses with IKE altogether”

Each router generates an AES key per TLOC and the controller distributes it over that same control connection.


What fact 3 concentrates

Every tunnel key in the fabric rides the control connection.

Impersonating a controller needs a classical CA or controller key recovered (for example by a CRQC) plus an on-path position. Hard. But if it works, the attacker is not decrypting one tunnel; they are positioned to receive the fabric's keys.

Cisco Catalyst SD-WAN is the documented example. Other products differ: some run IKEv2 control channels, which have the same whole-message retransmission. The IKE-less key distribution is Cisco-specific, so check your own platform.


The handshake goes from 8 datagrams to 40



Why the count is what matters. DTLS 1.2 retransmits an *entire flight* when any datagram in it is lost (RFC 6347 §4.2.4). One lost packet out of forty re-sends about 20 kB, down the link that just dropped a packet.

At brownout loss rates, that becomes an outage

95th-percentile time to establish one control connection (s)



15.1 s

p95 establish time, ML-DSA-87 at 3% loss (model).
Today it is 1.1 s.

16 %

of attempts take longer than five seconds (model).
Today, 0.1%.

At 1% loss, a healthy circuit, this is a slow migration. **At 3% it starts to look like an outage, and that is exactly when the fabric most needs to re-converge.**

The answer is not a smaller certificate

ML-DSA and hybrid ML-KEM key exchange are specified for TLS 1.3; no specification defines either for TLS 1.2. So the transport has to move anyway. DTLS 1.3 acknowledges handshake records, and the sender should resend only what was lost.

ML-DSA-87 on DTLS 1.2

15.1 s _{p95}

16.3% of attempts over five seconds



40 datagrams · lose one, resend the flight

ML-DSA-87 on DTLS 1.3

2.1 s _{p95}

0.2% of attempts over five seconds



37 datagrams · lose one, resend one

Same credentials. Same loss. In our model, the transport change is worth more than any certificate optimisation.

Two things engineering alone cannot fix



The factory identity

The certificate in the Trust Anchor module is programmed at manufacture and cannot be replaced. 802.1AR LDevIDs enrolled with BRSKI can add a post-quantum operational identity, but that enrolment rests on the classical factory key. It must happen before a CRQC and be redone after every RMA.



The clock

Signature forgery is not retroactive. The data-plane keys riding a classically keyed control channel are, though: harvest now, decrypt later. And a 7–10 year refresh runs past CNSA 2.0's 2030 date for national-security network gear and NIST's proposed 2035 disallowal of RSA-2048.

For SD-WAN hardware shipping today, the decision is being made now.

Not because the threat is imminent, but because the refresh cycle outlasts the policy deadlines. That is the argument that moves budget holders.

Four things to do, in this order

1

Move the transport before the credential

Standards-based ML-DSA and hybrid key exchange need DTLS 1.3 anyway, and in our model its ACKs are worth more than any certificate optimisation. This is the finding I most want tested.

2

Keep the chain off the wire — where there is a chain

The manager already sends configuration to every router. Pre-provision intermediates there (Trust Anchor IDs §9.5) for every handshake after the first: 41.7 → 26.3 kB, 40 → 27 datagrams. If your PKI is flat, there is little to elide.

3

Make the handshake survive a flap

DTLS Connection IDs and PSK resumption already exist. SD-WAN should use them, with a policy for when resumption may replace re-authentication, so a flap or reload doesn't cost a full post-quantum handshake.

4

Make the failure attributable

Today an oversized handshake shows up as DCONFALL or a timeout, the same codes as a broken circuit. Report handshake bytes, datagram count, retransmissions and establish time per transport.

None of this needs new cryptography.

What I want from this room

Data

What is the real loss, burst and RTT distribution on production SD-WAN broadband, LTE and satellite? Vendors and managed service providers hold this. We have not found it published, and it would correct my weakest assumption.

Sequencing

Is DTLS 1.3 adoption the real gate on post-quantum SD-WAN, rather than certificate size? If so, the advice operators are being given is backwards.

Coupling

Data-plane keys ride the control connection. Does that argue for a stronger parameter set there, or for reducing the coupling?

Identity

Factory identities cannot be replaced and stay in service 7–10 years. LDevID enrolment via BRSKI rests on the classical factory key: by when must it happen, how is it redone after RMA, and who signs?

IETF SD-WAN work exists in BESS and IDR, focus to cover post-quantum authentication of SD-WAN control connections is necessary.

Thank You!

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

Afshan Mulla and Jaypal Baviskar ,
Technical Leader, Cisco CX, Professional Services