TIEVA

FULCRUM IT PARTNERS

10 Real-World SD-WAN Gotchas to Consider Early

By Andy Jukes

About the Author

Andy Jukes is the Secure Networks Team Leader at TIEVA and has been a key part of the business for over five years. He brings deep technical expertise in the design, implementation, and support of secure networking solutions.

Andy specialises in SASE, SD-WAN, and next-generation firewalls - helping organisations modernise their network security to enable secure distributed workforces, reduce operational complexity, and support long-term cloud strategies. With a practical, hands-on approach, he works closely with IT teams to align technical decisions with real business outcomes.

Andy Jukes

10 Real-World SD-WAN Gotchas to Consider Early

Over the years, I've seen SD-WAN deliver on its promises: smarter routing, better cloud access, less reliance on legacy WANs. But I've also seen plenty of projects hit bumps in the road.

That's where the "**gotchas**" come in. By gotchas, I don't mean outright failures or catastrophic mistakes - I mean the small design decisions that feel fine on day one but create problems later.

The sort of things that don't show up in a demo, but surface six months down the line when you're trying to scale, troubleshoot, or keep users happy.

Here are 10 real SD-WAN gotchas - and, more importantly, the practical ways you can avoid them.

1. Start Simple - But Design SD-WAN for Scalability

Many of the SD-WAN projects I've worked on start simple - a hub-and-spoke design with 10 or 20 sites. And that's fine, until the business grows and suddenly 200 sites are trying to squeeze into a model that was never built for scale.

The challenge is to build a design that starts simple but scales without re-engineering.

Examples of where it breaks:

- **Routing choices –** decisions like relying solely on, eBGP or OSPF everywhere may feel simple, but they quickly become unmanageable at scale: too many sessions, unpredictable path control, and instability. Using iBGP with transitive attributes and route reflection as an example allows far greater scalability and predictability.

- **Rigid topologies –** a single hub may be fine at the start, but performance suffers without options to add regional hubs, cloud on-ramps, or SD-WAN hubs/spokes hosted in enterprise data centres to act as secure aggregation points both protecting sensitive workloads and providing controlled access to private and hybrid cloud resources.

- **DNS resolution –** using a single central resolver may be fine in small deployments, but at scale it adds unnecessary latency and often misdirects SaaS/CDN traffic to the wrong region. Distributed or cloud-based resolvers ensure queries are answered closer to the user, improving performance and routing accuracy.

- **Lack of traffic visibility –** skipping an upfront network assessment or ignoring traffic patterns and key applications means the design may not align with real-world usage, creating blind spots that impact scale and user experience later.

Designing it right:

- Build templates that allow iterative expansion and don't lock yourself into a fixed huband-spoke SD-WAN.
- Use cloud on-ramps for SaaS/IaaS access and natural application integration points.
- Leverage SD-WAN appliances in enterprise data centres for protecting private workloads and enabling controlled access to hybrid cloud.
- Don't introduce direct site-to-site tunnels where they're not appropriate use selective dynamic tunnels only where they add value. Avoiding unnecessary full mesh not only reduces complexity, but also helps control cost by limiting tunnel overhead and unnecessary license consumption.
- Begin with an assessment of existing network traffic to identify key applications, performance bottlenecks, and dependencies before committing to a final design.

Takeaway

Hub-and-spoke is a solid place to start - it's simple, proven, and gets you moving.
But the real test comes later. The designs that succeed are the ones that can grow gracefully into distributed, cloud, and hybrid models without needing to be ripped apart and rebuilt.

2. SD-WAN Overlays Don't Fix Bad Underlays

If you've ever dealt with flaky broadband at home, you'll know the frustration: no matter how good the Wi-Fi router is, a weak connection still drops out. SD-WAN works the same way. It can measure path health, steer traffic, and fail over between links - but it doesn't magically fix the quality of the connections underneath. SD-WAN builds an overlay - a virtual fabric of secure tunnels and dynamic policies - on top of your physical connectivity (broadband, MPLS, LTE/5G). And if that underlay is fragile, the overlay just makes those weaknesses show up faster.

The reality at scale:

- A congested broadband link still drops packets, no matter how smart the overlay is.
- Without diverse underlays, multiple tunnels will fail together.
- Latency spikes, jitter, and loss bleed straight through to applications.

Designing it right:

- **Invest in diverse underlays –** DIA + DIA, DIA + 5G/LTE, or dual broadband from different ISPs.
- **Dual links unlock the real value of SD-WAN –** steering, prioritisation, and resilience only work when there's more than one path.
- **Test load balancing and failover in practice –** just configure, simulate outages and confirm apps survive.

Takeaway

Secure SD-WAN adds intelligence, but it's not wizardry. If the connections underneath are weak or all the same, the overlay won't save you. Real resilience comes from building on strong, diverse underlays first.

3. Stop Clicking, Start Orchestrating

It's tempting to set up those first few sites by hand in a GUI - point, click, done. And for a handful of branches, it works just fine. But fast-forward to 200 sites and the cracks show. Configs drift, bulk changes turn into a nightmare, and when something breaks no one's quite sure who changed what or when. It's like trying to manage an orchestra by asking each musician to play their part in isolation - sooner or later, you lose harmony.

Manual configuration slows everything down: rollouts take longer, sites behave inconsistently, and troubleshooting becomes a guessing game. At scale, clicking around in a GUI just isn't enough.

Designing it right:

- **Think automation first –** use APIs and automation pipelines to push, validate, and roll back configs at scale.
- **Adopt role-based templates early,** so branch offices, data centres, and cloud edges all follow a consistent pattern.
- **Centralise visibility and control through an orchestrator** (FortiManager, Cisco vManage, VMware Orchestrator), reducing human error and giving teams a single source of truth.
- **Audit and compliance built-in** - orchestration ensures changes are logged, consistent, and reversible.

Takeaway

A GUI is fine when you're testing in the lab or setting up a demo. But once you're rolling out at enterprise scale, you'll need automation, templates, and orchestration to keep things consistent - and to stop the whole thing turning into a mess of manual fixes.

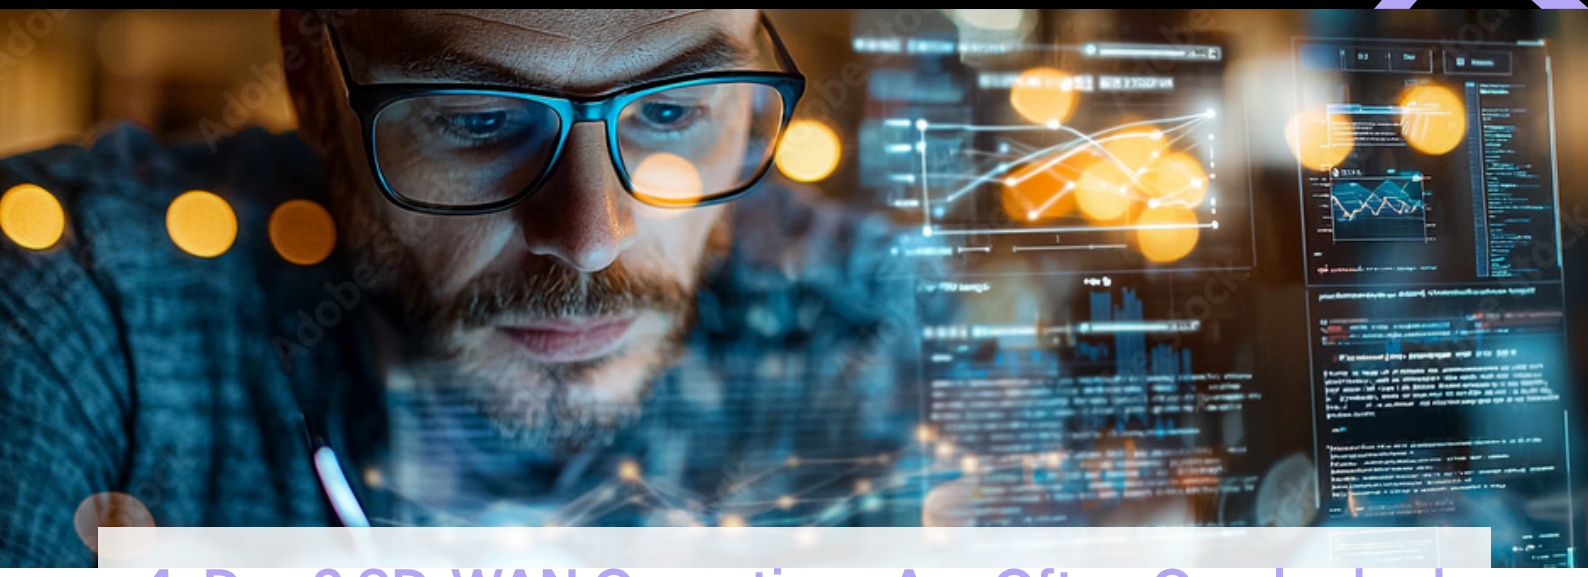

4. Day 2 SD-WAN Operations Are Often Overlooked

It's easy to get caught up in the excitement of go-live: the shiny new dashboards, zero-touch deployments working as promised, and those first tunnels coming online. But the honeymoon period doesn't last. The real test starts on Day 2 - when updates need patching, policies need fine-tuning, and new apps or sites suddenly need to slot into the fabric.

Without planning for that day-to-day reality, what felt simple at launch quickly turns into firefighting. Configs drift, odd user experience issues creep in, and keeping the SD-WAN aligned with business growth becomes a constant struggle. The result? Slower responses, rising complexity, and users wondering why things don't actually feel any better.

Designing it right:

- **Bake observability in from the start –** not just tunnel up/down status, but application health, user experience, and underlay performance.
- **Integrate with automation pipelines** so updates, patches, and monitoring are consistent and repeatable.
- **Plan for lifecycle management** - upgrades, new services, and capacity expansions should be expected, not disruptive.
- **Link ops to business outcomes** - measure not just uptime, but whether apps stay performant and users remain productive.

Takeaway

Getting to go-live is the easy part. It's what happens after - the updates, the tweaks, the everyday running - that decides whether SD-WAN keeps delivering value or slowly drags you back into firefighting mode.

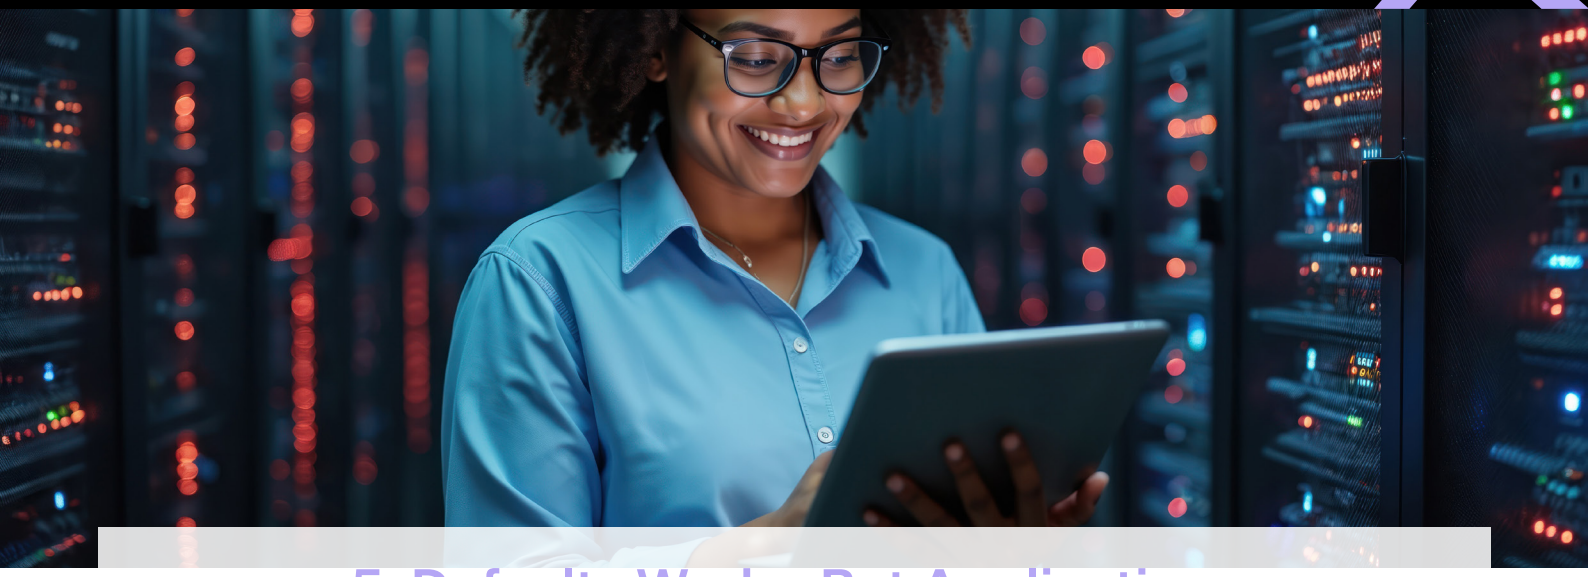

5. Defaults Work - But Application-Aware Policies Protect the Experience

Out-of-the-box SD-WAN QoS and latency-based steering get you up and running, but they're blunt instruments. Defaults can't tell the difference between a Teams call that needs crystalclear audio and a background software update quietly eating bandwidth. On paper, the network might look "**healthy,**" yet users are still complaining about choppy calls or slow apps.

At scale, this becomes a real operational headache. The business-critical apps don't get the protection they deserve, while low-priority traffic sneaks onto your best paths - which pretty much defeats the whole point of having intelligent steering in the first place.

Designing it right:

- **Run an SD-WAN assessment before rollout** - know which apps matter most, and what traffic patterns they generate.
- **Apply custom QoS profiles aligned to app needs** - low-latency and high-priority for voice/video, balanced for SaaS, best-effort for bulk transfers.
- **Use application-aware steering** - don't rely only on latency metrics; steer based on business importance.
- **Continuously review and adjust policies** - new apps and patterns will always emerge.

Takeaway

Defaults are fine for a demo, but real-world networks aren't demos. At enterprise scale, you need accurate app classification with custom QoS and smart steering - that's what keeps users happy and makes SD-WAN live up to its promise.

6. Standard Routing Works - But Vendor Features Maximise Value

It's easy to fall back on what you know: BGP communities, OSPF costs, MEDs. They're familiar, they're predictable, and yes, they get the job done. But if that's all you use, you've basically bought yourself a very expensive tunnel fabric. You miss out on the automation, simplicity, and intelligence that SD-WAN is built to deliver.

At scale, that choice comes back to bite. Instead of letting SD-WAN handle the heavy lifting, network teams end up buried in route-maps and policy hacks, burning time on workarounds the platform could do for them. And let's be honest: if you're paying for SD-WAN but treating it like a traditional router, you're not getting the value you signed up for.

Designing it right:

- **Use SD-WAN rules for intent-based traffic steering** - describe outcomes ("keep Teams on the lowest-latency path") instead of managing hop-by-hop routing metrics.
- **Leverage integrated features** - for example, application-aware routing alongside assurance capabilities such as packet duplication and forward error correction (FEC), which can be applied selectively to protect sensitive traffic like voice and video.
- **Blend features where appropriate** - use routing attributes where they add clarity, but let SD-WAN handle the heavy lifting.

Takeaway

Traditional routing still has its place, but don't stop there. The real wins come when you lean on SD-WAN's native features - things like application-aware routing and assurance tools. They simplify life for your team, make the network more resilient, and make sure you're actually getting your money's worth from the platform.

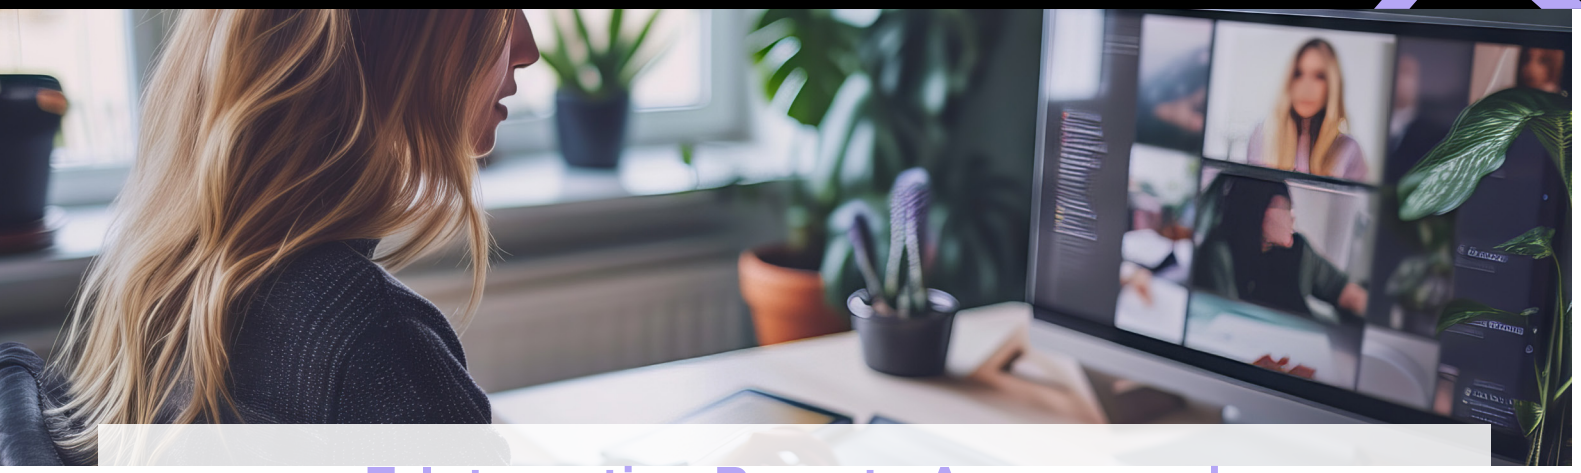

7. Integrating Remote Access and Partner VPNs into SD-WAN

It's common to treat remote access and partner VPNs as a separate stack from SD-WAN. On paper, it feels simpler: one setup for sites, another for users and partners. But in practice, that split creates headaches. Suddenly, VPN users don't follow the same traffic steering or security policies as your branches, and partners end up with access paths that work one day and break the next.

At scale, those silos really show. Troubleshooting becomes a nightmare because half your traffic is managed dynamically and half isn't - and the two worlds don't share the same visibility or enforcement. The end result? A network that looks neat on a diagram, but behaves inconsistently in the real world.

Designing it right:

- **Converge VPNs into SD-WAN hubs or cloud gateways** - remote workers and partners follow the same paths, policies, and security as branches.
- **Advertise upstream and VPN routes into the SD-WAN fabric** - connectivity is dynamic and resilient, rather than manually defined.
- **Ensure third-party partners integrate with SD-WAN routing updates -** if a site fails over to another hub, the partner must learn that new path automatically, not keep sending traffic to a dead endpoint.
- **Keep routing and policy consistent across all connections** - sites, users, and partners should all be part of the same policy framework.
- **Integrate with identity and access controls** so VPN users get the same segmentation and zero-trust principles as on-site users.

Takeaway

Keeping remote access and partner VPNs separate might feel tidy at first, but it only adds silos and complexity. Bringing them into the SD-WAN fabric means everyone - sites, users, and partners - follows the same rules, gets the same visibility, and enjoys a network that actually behaves the way the diagram says it should.

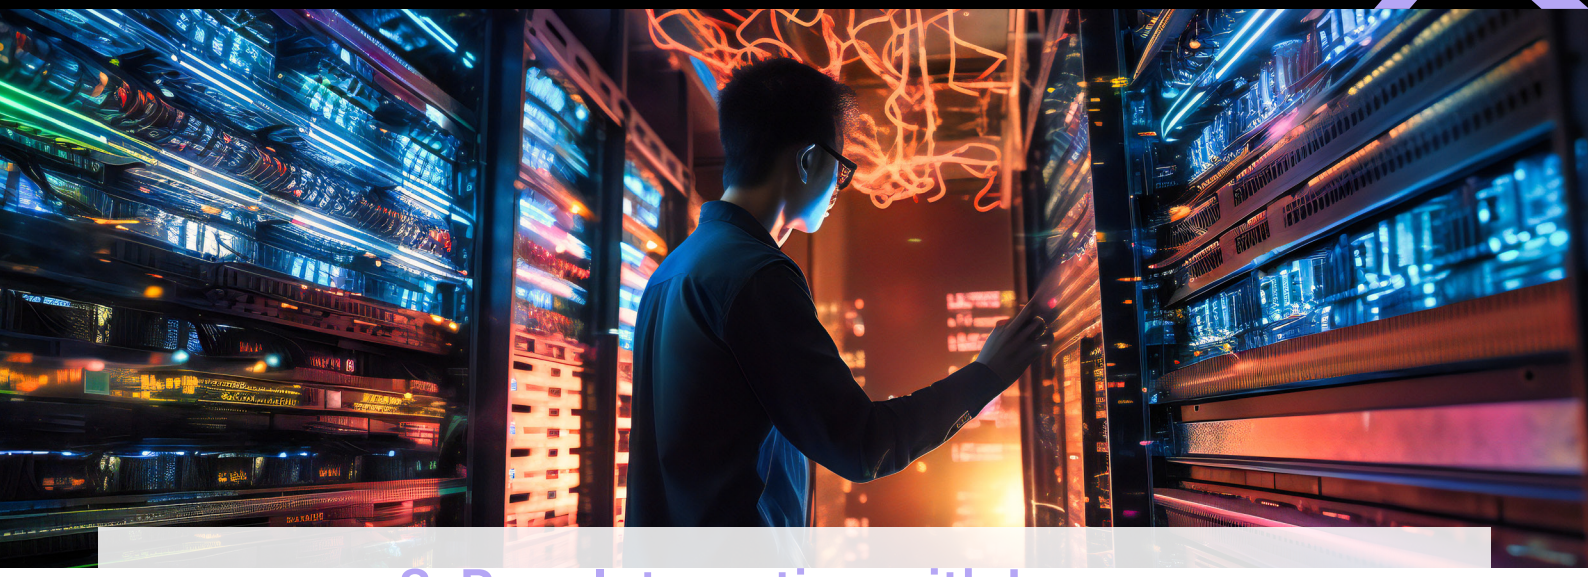

8. Poor Integration with Legacy Infrastructure and Overlooked Migrations

Most WANs aren't greenfield. They're usually a patchwork of what's come before - legacy firewalls, old VPN concentrators, overlapping routing domains, and sometimes an MPLS circuit or two still hanging around on contract. Ignore that messy reality during an SD-WAN rollout, and you're setting yourself up for outages, delays, and some very unhappy stakeholders.

SD-WAN doesn't live in a bubble. For months or even years, it has to run alongside what's already there. If you don't map and test those interdependencies properly, even a simple cutover can cause big problems - suddenly a critical app is unreachable, or a security control gets bypassed without anyone noticing until it's too late.

Designing it right:

- **Plan for dual control periods** - expect SD-WAN and the legacy WAN to coexist. Document how routing, security, and monitoring will work during the overlap.
- **Use phased migrations, not "big bang" cutovers** - move in manageable steps (per region, per site type, per application) to reduce risk.
- **Build interop gateways between SD-WAN and legacy domains** - so both environments can exchange routes and maintain reachability.
- **Audit firewalls and ACLs early** - legacy rules can silently block SD-WAN traffic or create asymmetric paths if overlooked.
- **Communicate with carriers and partners** - many dependencies sit outside your direct control, so coordination matters.

Takeaway

Pretending the legacy stuff doesn't exist is a recipe for failure. Plan for overlap, interoperability, and phased migrations right from the start. It might look slower on paper, but in reality it's the fastest and safest way to get SD-WAN up and running without breaking things.

9. SD-WAN Isn't Just WAN - It's End-to-End

If you think of SD-WAN as just a WAN replacement, you're selling it short. Yes, it can reduce MPLS costs and simplify branch connectivity, but that's only scratching the surface. The real power is in creating an end-to-end fabric - one that ties everything together: Wi-Fi, LAN, WAN, cloud, and even security services.

With that approach, you close the gaps. Users get the same experience no matter where they connect. Apps follow the same rules whether they're in your data centre or the cloud. Security is applied consistently, instead of in silos. And when you add ZTNA into the mix, SD-WAN extends zero-trust principles beyond the branch - so remote workers and partners get the same protection as on-site users, without the weak spots of legacy VPNs.

Designing it right:

- **Extend consistency of policy and segmentation** from Wi-Fi through to WAN and cloud, so identity and traffic types are treated the same everywhere.
- **Take advantage of ZTNA and SASE/SSE integration** to close gaps left by traditional VPNs or backhaul-heavy security models.
- **Use cloud on-ramps and multi-cloud interconnects** to reach applications where they actually live, not where the WAN thinks they should be.
- **Align with IT and security teams** so SD-WAN becomes part of a broader enterprise fabric strategy, not a standalone networking tool.

Takeaway

SD-WAN is most powerful when you stop treating it as just a WAN replacement.
Used end-to-end, it becomes the backbone of your whole enterprise fabric - connecting access, transport, and security in a way that keeps everything consistent, secure, and free of gaps.

10. Skipping Real Pilot Testing Leads to Pain Later

It's tempting to rush through pilots or keep them to a "happy path" demo - one branch, one link, a few test pings. The trouble is, that doesn't look anything like how your network will actually be used. Then the first real test ends up happening in production, and that's when QoS issues, failover problems, or app performance hiccups show up - often right in front of your end users or leadership team.

A solid pilot isn't about proving that SD-WAN works - the technology itself is mature. It's about proving that your design works in your world: with your apps, your legacy kit, your security controls, and under your conditions. That means testing brownouts, remote access, coexistence with the old WAN, and security enforcement. Skipping that step usually means discovering problems later, when they're harder (and more painful) to fix.

Designing it right:

- **Pilot at representative sites -** include different regions, connectivity types (broadband, LTE, DIA), and actual user traffic patterns.
- **Validate end-to-end behaviour** - not just tunnels coming up, but QoS, application classification, and traffic steering in practice.
- **Run failover drills** - simulate link failures, packet loss, and brownouts to see how apps behave under stress.
- **Include security enforcement and cloud integration** - test that policies, ZTNA/SASE integration, and SaaS access perform as expected.
- **Document and iterate** - use pilot results to refine templates and operational processes before scaling.

Takeaway

A pilot isn't just a checkbox. Real-world testing with real users, apps, and even a few failure scenarios is what saves you from nasty surprises in production - and gives everyone confidence to scale up smoothly.

The Big SD-WAN Takeaway!

These are 10 of the real SD-WAN “gotchas” I’ve seen in the field. In my experience, SD-WAN designs don’t usually fail outright. What happens instead is they hit walls: they don’t scale, they don’t adapt, or they don’t deliver the experience the business was expecting.

The technology itself is mature - it does what it says on the tin. The difference between success and failure comes down to the design choices you make early on and how you run the network day to day.

The organisations that succeed with SD-WAN are the ones who:

- Keep centralisation where it fits, but distribute where it adds value, avoiding the trap of rigid, one-size-fits-all models.
- Invest in underlay diversity as the foundation of reliability, knowing overlays can only perform as well as the paths beneath them.
- Bake SaaS, QoS, and application-awareness into the design, so the network makes intelligent decisions that protect user experience.
- Integrate VPNs, SASE, and end-to-end fabric consistency, ensuring no gaps between access, transport, and security.
- Plan for lifecycle operations, treating SD-WAN as a living system that evolves with applications, users, and the business.

“Good enough” will get you by. But it’s thoughtful design that makes SD-WAN truly future-ready - not just for today’s connectivity challenges, but as the backbone of a secure, agile, end-to-end enterprise fabric.

Thinking about SD-WAN - or already deploying it?

Whether you’re just starting out or you’ve already begun rolling out SD-WAN and are running into some of these challenges, it always helps to compare notes. At TIEVA, we work with organisations to design, deploy, and manage Secure SD-WAN in a way that avoids the common pitfalls and keeps your network aligned with your business goals.

If you’d like a conversation about where you are now - and where you want to get to - we’d be happy to chat.

Explore Secure SD-WAN with TIEVA

TIEVA

BY FULCRUM IT PARTNERS

Contact us

Email TIEVA
hello@tieva.co.uk

Call TIEVA
+44(0)333 0430 333

tieva.co.uk