



Building the board-ready business case for SD-WAN

How to prove the commercial value of network transformation

ednx

Discover. Design. Deliver.

A note from our CEO	03
The fundamentals of SD-WAN	04
Key considerations when quantifying impact	06
Translating technical features into ROI.	09
Deciphering different delivery models.....	12
The importance of pilots and parallel rollouts.....	14
A finance leader's SD-WAN success story.....	17
The EDNX advantage	19

SD-WAN is, by no means, a new technology. In fact, it's been more than a decade since early adopters first brought it to the market. But as the demand for faster, more collaborative, and intuitive IT environments grows, it's certainly earned its place among the tech stack's most impactful investments today.

The challenge is, particularly if you're not well versed in this world, overhauling your network can be a complex and overwhelming task. And when every pound of spend is scrutinised by the board, you need to be confident every component delivers immediate value, and is destined to improve your bottom line long term. Building this ROI case is often the biggest bottleneck for modern businesses. So, we've created this guide to help make the conversation easier.

My co-founder, Wojtek Siodelski, and I bring more than 35 years of frontline networking expertise, backed by a growing team of specialists that operate across both public and private sectors. From complex upgrades to greenfield environments, single sites to global enterprises, we've seen it all. We understand the unique pressures these companies face at every stage and how to align network decisions with wider organisational goals. More importantly, we know how to communicate this value clearly to stakeholders.

In the following pages, you'll find specific focus areas that showcase SD-WAN's commercial worth – from 'hard' operational savings to 'soft' cost efficiencies driven by performance enhancements – and understand which technical features can be measured effectively. We also uncover the often-overlooked factors that can help reduce risks and improve control. Plus, with a real-world case study from one of our leading finance clients, you'll have a better view of this technology in action.

As ever, if you have any questions, we're here to help. Otherwise, enjoy.

David Ratcliffe
Co-founder and CEO
EDNX



The fundamentals of SD-WAN

Traditional wide area networks (WANs) were built for a different world. Offices were the centre of gravity, applications lived in the data centre, and multiprotocol label switching (MPLS) circuits provided private connectivity to tie it all together. It was a logical architecture, offering predictable performance, fast data transfer, and centralised network management. But as cloud adoption accelerated, hybrid working became the norm, and real-time collaboration tools welcomed new performance expectations, this once solid foundation quickly began to strain.

Capacity had to be added to compensate for increased bandwidth requirements, which could be a lengthy process given that every circuit is statically provisioned. At the same time, backhauling all cloud traffic through a central data centre created congestion and increased latency, impacting productivity and external quality of service (QoS). Even small changes became inherently cumbersome, requiring manual intervention and complex path reconfiguration to maintain predefined policies and application dependencies. What began as a stable and effective model soon became progressively more expensive and harder to justify – not to mention increasingly misaligned with the demands of modern business.


Enter software-defined wide area networks (SD-WANs). Built to stop business data from taking the scenic route, this architecture offers an infinitely more dynamic solution than traditional MPLS networks, with leaders like Cisco bringing greater visibility, agility, and control to complex IT environments.

Rather than routing all traffic through fixed hardware paths, SD-WAN introduces a software layer that sits above the physical network and makes real-time decisions about how traffic moves. It can span multiple connection types simultaneously – such as MPLS, broadband internet, and 4G/5G – and intelligently routes applications through the best possible path. For example, a live video call might be assigned a high-priority, low-latency path with dynamic failover to limit disruption, while bulk data backups might be sent through a cheaper and slower internet link.

Security sits at the heart of SD-WAN design too – and for good reason. As users, applications, and devices become increasingly distributed, relying on a fixed perimeter security model simply isn’t practical. Instead, SD-WAN pushes controls closer to the edge, allowing consistent security policies to follow users and traffic wherever they sit. Features such as segmentation can isolate sensitive systems and contain threats, while zero trust principles help ensure every user, device, and connection is continuously verified before access is granted.

Ultimately, SD-WAN’s impact is manifold. Companies get better performance and greater flexibility, at a much more palatable and predictable cost. And they benefit from security that moves with the network, rather than being bolted on around it. But the degree to which those promises are realised depends entirely on the quality of design and rigour of implementation. The solution can’t simply be treated as a plug-and-play technology upgrade. To reap maximum ROI, it must be directly aligned to the unique demands of the business. Increasingly, this is the part many organisations underestimate.

SD-WAN architecture



Source: Paloalto Networks

Key considerations when quantifying impact

Understanding how SD-WAN works is important for enabling digital transformation, but building a credible commercial case for it requires a much more deliberate approach. By translating technical capability into clear business value, leaders can build a more confident case for investment that resonates with decision makers beyond IT.

Minimising downtime costs

Traditional WANs can be slow to respond to failure. When a circuit goes down, the network has to detect the problem, recalculate the best alternative route, and propagate the change across connected devices. This can take anywhere from seconds to several minutes, and in many environments, it still requires someone to intervene manually – and this is disruption that adds up quickly.

SD-WAN takes a fundamentally different approach. It monitors all available connections simultaneously and shifts traffic the moment a link degrades, often before anyone even notices a problem. For businesses where downtime means halted operations, missed transactions, or frustrated customers, improving the speed of recovery has direct financial value. Quantifying it means knowing your average disruption frequency, typical duration, and the loaded cost of an hour’s outage across affected teams or revenue streams.

Improving application performance

Not all traffic deserves equal treatment, but most traditional WANs provide exactly that. A video call, for example, is highly sensitive to delay and instability, while a large file transfer is not. When both are competing for the same connection without any intelligent prioritisation, the performance can suffer and productivity takes the hit. Let’s say that call was with an important stakeholder who’s tight on time and can’t reschedule, or a prospect that’s deciding who’s the more reliable of you and a competitor, the impact can ripple further still.

Legacy approaches to managing this challenge are largely manual and static – rules set once and left untouched, unable to adapt as the application landscape evolves. SD-WAN, on the other hand, continuously monitors connection quality and routes each application’s traffic via the path best suited to its specific requirements. And it adjusts in real time as conditions change. The commercial case rests on two things: the productivity lost when critical tools underperform, and the customer impact when poor application performance spills into service delivery.

Reducing network spend

MPLS built its reputation on reliability, with dedicated private circuits that guaranteed performance and were largely immune to the variability of the public internet. For a long time, that premium was worth paying. But internet connectivity has matured significantly, and the performance gap that once justified MPLS costs has narrowed considerably for most traffic types. The issue is that MPLS contracts are typically rigid – fixed capacity, long lead times to provision or change, and costs that don’t flex with demand. Organisations often end up paying for headroom they rarely use, while simultaneously running connections that are underutilised.

SD-WAN allows a more intelligent split – keeping MPLS where its reliability genuinely matters, and routing everything else over lower-cost connections. The result is a better return on overall network spend, and a cost model that’s far easier to scale (and predict).

Accelerating deployment and change

Deploying a new site on a traditional WAN can be a slow, hands-on process. Hardware needs to be configured, shipped, and commissioned by a skilled engineer on location, and any mistake in that process can cause significant delays and spiralling costs. Multiply that across a growing estate, and the network starts to constrain the pace of the business rather than support it.

SD-WAN's zero-touch provisioning model changes the economics of deployment entirely. Devices arrive pre-authorised and configure automatically once connected to the internet, with no specialist needed to intervene on the ground, streamlining and standardising rollouts. And, because the same logic applies to ongoing change, teams can easily push updates and policy adjustments across the entire estate from a single interface. For organisations scaling through new locations, acquisitions, or shifting working patterns, this isn't just an operational convenience, but a meaningful reduction in long-term provisioning costs.

Enabling cloud and business flexibility

Most legacy network architectures were designed around a central hub – typically the data centre – through which all traffic was routed before reaching its destination. When everything lived on-premise, that made complete sense. But cloud changed the equation. As organisations moved workloads to platforms like Microsoft 365, Azure, and AWS, backhauling data created unnecessary detours – adding latency, increasing congestion, and placing pressure on infrastructure that was never designed to support the volume of cloud-first traffic.

SD-WAN removes those bottlenecks. Rather than sending data on a round trip through the network core, it allows traffic to take a more direct path to the cloud via the most efficient connection available. It also reduces dependency on any single network provider, giving organisations greater freedom to negotiate, adapt, and make infrastructure decisions based on performance and value rather than contractual obligation. Over time, this flexibility only becomes more valuable.



Translating technical features into ROI

Every organisation's network is different, and the scale of opportunity with SD-WAN will vary depending on the size of the estate and the complexity of the existing infrastructure. But the areas where costs shift are consistent. Access circuits, hardware, licensing, operational overhead, and the hidden cost of downtime all move in a predictable direction when SD-WAN is introduced – and all can be modelled with reasonable accuracy against current spend.

The sections below break down each of these key components, before bringing them together in a simulated example that shows what the numbers can look like in practice.

Access circuits

Access circuits are often where the largest savings sit, making them one of the most defensible lines in the business case. Start with what’s currently contracted: how many MPLS circuits exist, at what capacity, and at what monthly cost. Then model a hybrid alternative – retaining private connectivity where performance genuinely demands it, and replacing the remainder with broadband or dedicated internet access (DIA) connections running under SD-WAN management. For most estates, the cost-per-Mbps comparison alone is enough to justify a conversation with the C-suite.

Support and licensing

A multi-vendor network estate generates a long tail of contracts that can vary significantly in scope. Auditing this spend in full – including maintenance agreements, software subscriptions, and any third-party management fees – often surfaces costs that have renewed automatically without scrutiny. This total becomes the baseline, while moving to a unified SD-WAN platform consolidates the majority of it into a clearer structure. The difference between the two is a direct saving, and the reduction in IT resource absorbed by vendor management adds further value that, while softer, can be reasonably estimated based on time currently spent.

Downtime and SLAs

Service level agreement (SLA) credits are a poor substitute for resilience. Not to mention, they can be a misleading benchmark when building a business case – the comparison shouldn’t be what the carrier pays back, but what an outage costs operationally. Take average disruption frequency and duration from incident records, apply a realistic hourly cost across affected headcount and any revenue impact, and compare that figure against what’s been recovered through credits. The gap is typically stark, and against that baseline, SD-WAN’s near-instant failover and continuous path monitoring represent a measurable reduction in financial exposure. Even when modelled conservatively, it can make a compelling case.

Cost per Mbps

Single-path WAN architectures are structurally inefficient. Traffic is constrained to one circuit regardless of what else is available, which means organisations routinely over-provision to protect against congestion peaks. As such, modelling requires an honest look at current utilisation – how much contracted bandwidth is being paid for versus what’s actually being used on average. Because SD-WAN actively distributes traffic across all available connections, it typically closes this gap significantly. Organisations can achieve the same or better effective throughput with a lower-cost circuit mix, which can have a major compounding effect over a three to five-year horizon.

Hardware and software estate

From routers to firewalls, the hardware underpinning a traditional WAN accumulates over time, each with its own procurement cost, maintenance contract, and refresh cycle. SD-WAN consolidates many of these functions into a single platform, which means a portion of that hardware becomes redundant at the point of migration and more at each subsequent upgrade. The best part is, the financial savings are easy to model: identify what can be decommissioned, remove the associated maintenance and licensing costs, and factor in the reduced capital requirement at the next hardware refresh. The savings are rarely headline-grabbing in isolation, but they contribute meaningfully to the overall picture.

A simulated SME example

Sometimes, it’s easy to let the figures speak for themselves. So, to make this tangible, the table below models a typical 10-site SME (circa 100–150 users) running a traditional MPLS-led WAN against an SD-WAN alternative. These figures assume a phased MPLS migration, retaining some private circuit capacity rather than implementing a full internet-only model.

Crucially, the intention of this table is not to replace an appropriately scoped commercial assessment, but to show the shape of the opportunity. Figures have been generated by AI to provide an illustrative framework only. Actual savings will vary – in larger or more MPLS-heavy estates, the circuit savings alone can dwarf every other line item. Meanwhile, in smaller organisations, consolidation of hardware and support costs often drives the strongest return.

Cost area	Traditional WAN (annual)	SD-WAN (annual)	Estimated saving (annual)
Access circuits	£42,000 (10 MPLS circuits at mixed capacities and contracted rates)	£24,000 (reduced MPLS footprint with broadband/DIA))	£18,000
Downtime/SLA impact	£10,000 (productivity loss and operational disruption beyond SLA credits recovered)	£3,000 (reduced impact through path monitoring and near-instant failover)	£7,000
Hardware estate	£8,000 (router maintenance, appliance support, and hardware refresh allocation)	£4,500 (consolidated edge devices and reduced maintenance burden)	£3,500
Supporting and licensing	£12,000 (multiple maintenance agreements, subscriptions, and management contracts)	£8,000 (single-platform licensing and simplified support structure)	£4,000
Bandwidth efficiency	£7,000 (cost of over-provisioned bandwidth to absorb utilisation spikes)	£2,500 (more effective traffic balancing across multiple links)	£4,500
IT operational overhead	£15,000 (engineering time spent on routing changes, troubleshooting, and vendor management)	£8,000 (centralised orchestration and reduced manual intervention)	£7,000
Total	£94,000	£50,000	£44,000 (~47%)

Deciphering different delivery models

The commercial case for SD-WAN doesn't end with the technology decision – how it's delivered has a direct bearing on cost structure, internal capability, and long-term value. It's a dimension that often receives less scrutiny than it deserves at board level, yet the choice between delivery models can materially change both the upfront investment and the ongoing financial commitment. The two primary models are broadly defined by where responsibility sits once the network is live.

Installation in isolation

An installation-only engagement covers assessment, design, deployment, and commissioning. The partner delivers a fully functional SD-WAN environment, and operational responsibility transfers to the internal team once live. From a commercial standpoint, this model has clear advantages. The cost is largely fixed and defined upfront, making it straightforward to present to a board and easier to evaluate against projected savings. And, because there’s no ongoing service fee eroding the return, the cost model is stable once the investment is made.

Crucially, this approach doesn’t equate to abandonment. If your partner is credible, they’ll be there to support any genuine architectural issues as they arise – not as a general IT service desk, but as a company that takes genuine pride in their work. Thorough documentation, structured knowledge transfer, and access to expert support where it’s genuinely needed are all part of a good quality engagement. The goal is to leave internal teams genuinely confident and capable of managing infrastructure, which can deliver the strongest long-term ROI if an organisation has the right resources.

Retained support with managed services

A managed service extends the partner’s involvement beyond deployment, covering ongoing monitoring, change management, and day-to-day operational support under a recurring contract. For organisations without the internal resource or appetite to manage a transformed network estate, this can be a justified investment. It reduces the burden on IT teams and provides a defined escalation path when issues arise.

The commercial trade-off, however, deserves honest scrutiny. Managed services introduce ongoing OpEx that continues indefinitely, and that recurring spend needs to be factored into any ROI model from the outset. It’s also worth recognising the structural dynamic at play. A managed service provider has an inherent commercial interest in maintaining involvement, which can – if left unexamined – lead to a scope of service that extends beyond what the organisation actually requires. This isn’t an argument against managed services, nor is it a universal statement, but it is a case for entering into any agreement with clear eyes. Define the scope tightly, understand what’s included and what isn’t, and be realistic about whether the ongoing cost is genuinely justified by the resource gap or risk.

Making a sound commercial decision

The bottom line is, there’s no one-size-fits-all route. Organisations with strong and competent IT resources are often well-positioned to take full ownership after a well-executed installation – avoiding the ongoing cost while retaining the flexibility to manage the network on their own terms. Where resource is genuinely constrained, or where the complexity of the environment warrants it, a managed service can be a savvy commercial decision – provided the scope and cost are proportionate. What matters most is that the choice is fully informed and based on an honest assessment of capability, appetite, and risk.

The importance of pilots and parallel rollouts

Every SD-WAN deployment carries some degree of commercial and operational risk. The network is business-critical infrastructure, and if the transition is poorly planned or rushed, the consequences often ripple far and wide. That's why a meticulously detailed approach is key.

Proving the case before full commitment

A proof of concept (PoC) or structured pilot is one of the most effective tools available when building and validating a business case for SD-WAN. Rather than relying solely on vendor benchmarks or projected performance figures, it allows organisations to test the solution against their own applications, users, and network conditions in a controlled environment, before any significant investment is locked in.


Virtual platforms make this increasingly accessible. SD-WAN environments can be simulated and stress tested without the need for full physical deployment, allowing teams to model specific scenarios, probe edge cases, and surface issues that wouldn't otherwise emerge until go-live. More than technical confidence, this forms a crucial evidence base for the board. Real performance data gathered from a pilot strengthens the business case materially, giving decision makers something concrete to evaluate.

Removing the risk of the 'big bang' rollout

Even with a successful pilot, full deployment carries risk if it's approached as a single, simultaneous switchover. A 'big bang' migration – cutting over the entire estate in one move – leaves no margin for error and hangs business continuity in the balance. It's an approach that tends to be driven by pressure to move quickly rather than a considered view of risk, and it's one of the more common sources of avoidable disruption during network transformation.

Deploying SD-WAN in parallel with existing infrastructure addresses this challenge directly. The new environment is built and validated alongside what's already in place, testing performance, resilience, and security with live traffic before the legacy network is decommissioned. Sites can be migrated in a controlled sequence, with each one validated before the next is moved, and if something unexpected surfaces, it can be resolved without impacting the broader business. By the time the legacy infrastructure is switched off, the SD-WAN environment has already proven itself in production.

But there's also a secondary commercial benefit that many companies overlook. The evidence gathered during a staged rollout further strengthens the business case as it progresses, giving the board visibility of tangible results rather than projections. It removes the risk-driven urgency that can otherwise push organisations into decisions they're not fully ready to make.



“The organisations that get the most out of SD-WAN are those that resist the pressure to move fast and instead invest time in structured testing and a controlled rollout. Running pilot programmes and parallel rollouts are not caution for caution’s own sake – it’s how you protect the business while building a long-term evidence base that’s impossible to argue with.”

Wojtek Siodelski
Co-founder and CTO
EDNX



A finance leader's SD-WAN success story


The principles outlined in this guide only carry so much weight in the abstract. This real-world example shows how those principles were applied for one of our clients in practice – including the challenges that prompted the transformation, the approach taken to deliver it, and the outcomes achieved on the other side.

A well-established financial services organisation came to EDNX with a familiar problem. Its MPLS-based network was expensive, rigid, and increasingly at odds with an ambitious cloud-first strategy – tying the business to a single cloud provider and making every change a slow, resource-intensive process dependent on service provider involvement.

EDNX designed and deployed a Cisco SD-WAN environment to replace it. Following a thorough diagnostic and detailed design stage, migration was phased and deliberate – offices moved across in a controlled sequence with performance validated at every step. Cisco Cloud OnRamp then unified the organisation’s cloud environments, intelligently routing traffic across AWS, Azure, and OCI based on performance and security needs, with consistent policy management across all three through a single interface. The legacy network ran in parallel throughout, until confidence in the new environment was complete – a transition approach that ensured zero unplanned downtime from start to finish.



Reduced circuit costs
Migrating from MPLS to an internet underlay delivered material savings.



Multi-cloud freedom
Optimised connectivity across AWS, Azure, and OCI eliminates lock-in.



Full operational agility
Changes can be made internally without service provider dependency.



Consistent security
Policy enforcement is unified across all locations and cloud platforms.



Zero downtime
The new environment was fully proven and validated before switching.

The EDNX advantage

A compelling SD-WAN business case requires more than a technology breakdown. It demands a partner that understands the commercial imperatives behind the investment and can translate them into a solution that delivers on both sides of the equation – superior performance for the tech team, and measurable ROI for finance. For ambitious organisations, that partner is EDNX.

As a trusted network specialist with deep engineering expertise, our team brings hands-on experience across a broad range of public and private sector environments. Beyond adding technical depth to ensure superior infrastructure performance, our knowledge ensures investments make as much sense on a balance sheet as they do on a network diagram.

We're also tech agnostic in everything we do. This means, rather than leading with a preferred product, we start with an organisation's operational and commercial goals before recommending the architecture that best fits. Where Cisco SD-WAN is the right answer, EDNX has the partnership depth to deliver it with confidence. Where the picture is more nuanced, our independence ensures the recommendation is driven by what's right for your business. No square-peg-round-hole solutions.

Whether you want to learn more about the fundamentals of SD-WAN, need help building a board-ready business case, or want to understand how we help, from assessment and design to deployment and optimisation, we're here to support.



Get in touch for a free consultation

info@ednx.co.uk
ednx.co.uk
0800 0336 186

ednx Discover. Design. Deliver.