

SD-WAN ROI

A Financial Evaluation Guide
for CFOs



Dr.A.Shaji George



SD-WAN ROI-A Financial Evaluation Guide for CFOs

Dr.A.Shaji George

Independent Researcher, Chennai, Tamil Nadu, India.

Abstract – In this article, we will explore the financial aspects of Software Defined Wide Area Network (SDWAN) deployments through a critical cost analysis lens and discuss how uneducated IT heads and pushy salesmen often mislead organisational leadership in the cost savings of SDWAN deployments. The study indicates that the financial benefits of SDWAN are directly tied to how dependent an organisation is on cloud services and that SDWAN implementations can be financially beneficial only when there are significant workloads in cloud services. Studies have shown that organisations that are less than 40% cloud dependent are likely to see a net loss of money after implementing SDWAN, despite what the salesman would have them believe during the sales pitch. The article offers detailed calculation frameworks that can be used by financial teams to independently determine savings claims, red flags within IT proposals and concrete thresholds to make IT investment decisions for SDWAN. Readers learn about the various costs involved, such as MPLS contracts, bandwidth, equipment depreciation, and operational expenses, and thereby have the tools to question assumptions made by the IT leadership. The analysis sets the stage for the right SDWAN implementation timing, based on tangible cloud migration milestones instead of vendor marketing cycles or IT departmental preferences. Financial professionals who finish this article will have the information they need to avoid investing too early or incorrectly in SDWAN, thus avoiding major capital erosion.

Keywords: SD-WAN ROI, SD-WAN TCO, MPLS vs SD-WAN, SD-WAN cost, SD-WAN business case, Network cost analysis, SD-WAN pricing, Cloud migration costs.

1. INTRODUCTION

Across the globe, organizations are approving SDWAN projects every year, and they are based on financial projections that never come to fruition. The trend is the same in every sector. The CFO is being presented with some solid numbers by an IT Head. A network vendor's sales professional supports each assertion with slick presentations and case studies. Leadership approves major capital investments, with an expectation of network economics changing. Three years later, the savings have not materialized, the IT Head has moved on to another firm and the finance team are still trying to figure out how the projection and reality don't match.

This situation arises due to a basic miscommunication. Many IT leaders don't have the financial modeling skills to accurately project TCO. Sales people work on a commission only basis, where they are paid for a sale made, irrespective of the customer's result. Financial teams, who are not used to considering network architecture details, blindly believe technical recommendations and lack the means to validate them.

In this article, we want to give financial professionals specific and practical knowledge to assess SDWAN proposals on their own. This is not a technology **how to** manual. The emphasis is virtually entirely on monetary aspects. So, how does SDWAN save money. When do those savings become a reality. What calculations can financial teams do to check claims. But most importantly, how much of the service must be in the cloud for SDWAN to be cost effective. The risk of losing is high. Medium sized enterprise could invest in SDWAN implementation anywhere from five hundred thousand dollars to two million dollars. Ten million or more can

be invested by large enterprises. This financial fallout is when these investments don't produce the expected returns. Tons of capital is locked up in network infrastructure that offers marginal value over existing solutions, which could otherwise be used for revenue generating initiatives.




Fig -1: SD-WAN A Smarter, More Flexible Way to Connect Your Business

It is not a technical subject, and anyone can learn about it. The mathematics is simple. The frameworks outlined here put into perspective the complex ideas of networks in the language of finance that any capable analyst can assess. At the end of this article, readers will have a series of questions to pose to IT leadership, red flags to look for in proposals, and calculation techniques to test any savings claims prior to approving spend.

2. OBJECTIVES

The main goal of this analysis is to give financial teams the tools to independently assess SDWAN investment proposals, based on IT department assessments or vendor claims.

Other goals include setting clear financial points for when SDWAN makes sense, understanding the common manipulation tactics of uninformed IT leaders and commission-driven salespeople, developing calculation models that reveal inflated savings estimates, and establishing a correlation between cloud service dependency and SDWAN financial viability.

The article seeks to change the paradigm from technology assessment to financial analysis, because SDWAN decisions are basically about capital allocation decisions, which are outside of the purely technical realm and thus should be under financial governance.

3. CURRENT TRENDS IN SDWAN MARKET DYNAMICS

The global market for SDWAN in 2025 was about eight billion dollars, and is expected to grow over the decade. This growth is a true sign of real technology progress, and of the efforts of vendors who are looking to grab market share from products that are becoming commoditized.

What we are seeing is that vendors are directly going to Chief Financial Officers without going through the technical assessment process. Now marketing materials focus on ROI calculations, payback periods and TCO comparisons. This is a recognition that the ultimate decision maker in purchasing is the financial decision maker, but it also opens up the opportunity for manipulation when finance professionals don't have the context to assess claims.

CURRENT TRENDS IN SDWAN MARKET DYNAMICS

Bigger Market • Finance-Driven Decisions • Fewer Big Vendors • More Managed Services • Subscription Pricing




Fig -2: Current Trends in SD-WAN Market Dynamics

The vendor market has become quite consolidated with some key players such as Cisco, VMware, Fortinet, Palo Alto Networks and HPE Aruba having significant market share. This consolidation has created a strong competitive environment, which has resulted in more aggressive selling and more optimistic savings estimates in proposals. At the same time, the managed services model has become popular many organizations prefer to outsource SDWAN operations to avoid developing in-house expertise. This can help to simplify operations, but it can also help to bury the real costs in the bundled service agreements, which makes it harder to compare the costs. The pricing models that are in place these days are based on subscription and they complicate the cost analysis. When network services move from capital expenditure to operational expenditure, traditional capital expenditure calculations don't apply as cleanly. Financial teams need to adjust the evaluation frameworks accordingly, but many organizations are still using outdated assessment methods.

4. THE FUNDAMENTAL FINANCIAL PREMISE OF SDWAN

Before looking at manipulation tactics and calculation frameworks, it's important to understand the basic financial proposition of SDWAN. There are various ways in which technology can save, and with each there are certain conditions that have to be met for the savings to occur.

The main savings are from the cost of replacing expensive MPLS circuits with less expensive broadband circuits. The typical price range for MPLS is \$500 to \$3,000 per megabit per second per month, depending on geographic location and contract terms. Broadband internet may be priced at half the cost to twice the cost for the same bandwidth. The math's seems convincing. Save money and replace costlier private circuits with lower-cost Internet connections.

THE FUNDAMENTAL FINANCIAL PREMISE OF SDWAN

Use lower-cost Internet connections instead of expensive MPLS, and let SD-WAN intelligently route traffic to the best destination.




Fig -3: The Fundamental Financial Premise of SD-WAN

But there are key assumptions in this calculation. First, it is the assumption that your application traffic characteristics are tolerant of the characteristics of the internet, such as latency, packet loss and jitter. Second, it takes for granted that you have a security architecture that can defend traffic running through the public infrastructure. Third, and most critical to this discussion, it assumes that your traffic patterns will benefit from the flexibility SDWAN is able to offer.

That's where the key information is that many IT leaders don't understand or choose to overlook. SDWAN adds value by dynamically directing traffic through a variety of transport types to suit the needs of applications. For an application in the cloud, traffic is going to gravitate towards Internet based transports. The SDWAN device detects this traffic behavior and efficiently routes it without sending it through the costly MPLS backhaul in the corporate data center. This optimization opportunity is hardly relevant when applications are all in corporate data centers. The flow of traffic from branch sites should still be able to get to the corporate data center. It doesn't matter whether the traffic is going over MPLS or internet with VPN encryption, it's still going the same

distance to the same destination. There's nothing to optimize with SDWAN's routing intelligence. This correlation between cloud dependence and SDWAN value is the key financial fact that uninformed IT leaders don't understand and unscrupulous salespeople use to their advantage.

5. THE CLOUD DEPENDENCY THRESHOLD WHEN SDWAN MAKES FINANCIAL SENSE

The financial analysis of SDWAN deployments always shows a **tipping point**. At low levels of cloud service dependency, SDWAN investments don't get positive returns. Above these values, the returns may be significant.

5.1 Let me explain this through concrete calculations that any financial analyst can verify

Let's take an organization with 100 branch sites, all of which are connected using MPLS at \$2,000 per site, per month. The annual cost of MPLS is \$2.4 million. The IT manager recommends implementing SDWAN and states that it will save 50% or 1.2 million dollars a year.

5.2 Now examine the assumptions underlying that claim

With SDWAN, if 90% of application traffic is destined for corporate data centers for on premises applications, then the traffic patterns are the same. The traffic still has to get to the data centre. When the organization keeps MPLS for reliability and introduces internet circuits for SDWAN, costs actually go up. If the organization removes MPLS altogether and uses internet circuits optimized for SDWAN, they run the risk of reliability for traffic patterns that SDWAN can't meaningfully impact.




Fig -4: The Cloud Dependency Threshold When SD-WAN Makes Financial Sense

With 60% or more cloud traffic destinations, the math changes dramatically. Internet transport can be used to route directly from branch locations to cloud points of presence, such as Office 365, Salesforce, AWS workloads, and others. This traffic does not have to go over costly MPLS circuits. The organization is able to



scale down MPLS capacity, by completely removing circuits at certain sites, or by scaling bandwidth tier at others.

Based on research from several implementation studies, the following financial thresholds are suggested. For organizations that do not have a high dependence on cloud services less than 20 percent SDWAN implementations will result in net cost increases. With 20–40% cloud dependency, organizations can expect to break even or realize small savings, depending on the terms of existing contracts and the effectiveness of implementation. Companies that are 40–60% cloud dependent see a substantial reduction in their network expenses, typically between 15% and 30%. For organizations that are more than 60 percent cloud dependent, the savings that are touted in the sales pitch can be significant, up to 40 to 50 percent less network spend. These thresholds are based upon good implementation, fair contract terms, and suitable transport options. They also make the assumption of honest evaluation of cloud dependency, which leads us to the next critical issue.

6. HOW UNINFORMED IT LEADERS MISCALCULATE CLOUD DEPENDENCY

To accurately assess cloud dependency, there needs to be an understanding of traffic patterns at the application level. Many IT managers are not getting this visibility or are using incorrect methodologies that overestimate the amount of cloud they may be dependent upon. Cloud dependency percentage is likely the most critical factor impacting SDWAN financial results, and organizations often make investments with flawed dependency calculations. The success or failure of SDWAN implementation can often be attributed to the difference between the perceived and actual cloud dependency. To make sure investments based on cloud traffic assumptions are methodologically sound and approved by finance leaders, they need to be aware of common pitfalls in miscalculations..

6.1 Common Methodological Errors

Some of the common pitfalls are counting shadow IT cloud usage, which already exists outside of corporate networks counting cloud services that are planned to move to the cloud in current dependency calculations the assumption that hybrid cloud architectures, where cloud services still need to connect to data center resources, are the same as pure cloud services and the assumption that all SaaS applications are equal, irrespective of actual cloud usage traffic.

Shadow IT cloud usage refers to the services that employees use without the knowledge or permission of IT. These are usually personal productivity applications, file sharing services, consumer communication applications and departmental software that is not purchased through procurement processes. This traffic already bypasses enterprise connectivity and goes straight to internet destinations as shadow IT is outside corporate network infrastructure. SDWAN can only optimize traffic that passes through corporate networks. Shadow IT is a cloud dependency calculation that inflates the percentage and doesn't help realize any savings.

The shadow IT measurement problem is exacerbated when organizations try to survey their cloud services rather than analyze them using the network. Users of applications report the usage without distinguishing between services used via the corporate infrastructure and services used via direct internet connections from their personal devices and home networks. The resulting inventories vastly exaggerate the applications that are relevant to SDWAN traffic optimization. For dependency assessment, it is important to measure traffic over the network since only services accessed via corporate network infrastructure are eligible for SDWAN redirection.

When future planned cloud migrations are added, there is temporal confusion in current dependency calculations. Application migration schedules are usually part of an IT roadmap and plan for a gradual shift of on-premises applications to cloud platforms. Sometimes, enthusiastic IT leaders factor in projected future state instead of current reality, and they project favorable SDWAN economics based on a situation that doesn't currently exist and may never be true. As organizations face unforeseen complexity, compliance requirements and resource constraints, migration timelines continue to fall behind. When dependencies are based on projected migration, instead of actual migration completion, the savings estimates can only become a reality once migration is actually complete.

HOW UNINFORMED IT LEADERS MISCALCULATE CLOUD DEPENDENCY

Different methods. Different numbers. Different financial results.

Understand real traffic patterns, avoid common mistakes, and base SD-WAN decisions on facts, not assumptions.




Fig -5: How Uninformed IT Leaders Miscalculate Cloud Dependency

Hybrid architectures are especially problematic with regard to dependency classification. As more enterprise applications are deployed to the cloud, hybrid application deployment models are becoming more common, with parts of the application deployed on-premises and parts deployed in the cloud. A hybrid application could be one that runs transactions in the cloud, but keeps master data in enterprise data centers. However, traffic patterns are still very much reliant on data center connectivity, with the cloud component showing up in cloud application inventories, since cloud processing depends on constant communication with on-premises data sources. Hybrid applications are not true cloud applications as they do not have the same traffic pattern and potential for direct routing to the internet.

Pure cloud services are entirely contained within the cloud provider's infrastructure, and do not need to connect to enterprise data centers for normal operation. These services are good candidates for direct Internet access via SDWAN since there is no involvement of the enterprise data center between user locations and cloud destinations. To differentiate pure cloud services from hybrid architectures, IT managers need to understand application dependencies at a level that they may not have. It makes a huge difference for



SDWAN financial projections as hybrid applications can't be moved to direct internet routing no matter what SDWAN can do.

SaaS application equivalence assumptions assume that all SaaS applications have the same network demands. This is a very wrong assumption to make when looking at real traffic patterns. In addition to the lightweight applications such as electronic signature service, survey platforms, or project tracking software, email and collaboration platforms may be included in application inventories. Dependency percentages calculated by simply counting applications, without taking into account the volume of traffic, yield dependency percentages that do not reflect reality. Application counting methodologies are the same for 10 lightweight, megabyte SaaS applications as they are for a single heavy application that generates gigabytes of traffic.

6.2 The Application Count Versus Traffic Volume Distinction

Let's take an example. An IT manager states that he is 60 percent dependent on the cloud, by number of applications. Forty cloud-based applications and twenty-six on-premises applications are in use. When it comes to applications, the dependency on cloud seems to be high.

This approach is attractive to IT managers because it is fairly simple to create an application inventory. In most enterprises, there are application portfolios that contain a list of systems in the enterprise. The classification as cloud or on-premises applications is quite easy and does not involve much technical analysis. Percentages are easily computed from inventories that have been categorized with a basic arithmetic calculation. This methodology is fundamental flawed for SDWAN financial evaluation, but is easy to use and therefore widely adopted.

The application count methodology assumes that all applications are equal in terms of traffic. While this may be the case when applications impose similar network demands, enterprise application traffic patterns can be extremely variable. The volume of traffic produced by core business systems that transact thousands of times a day is orders of magnitude greater than that created by non-core systems that are used infrequently and by a small number of users. Unlike the annual performance review applications, which are accessed in brief spurts during evaluation times, inventory management systems are constantly updating across locations, creating ongoing bandwidth requirements.

But the numbers of vehicles on the road are different. These 40 cloud applications could include such things as project management software, collaboration platforms, and HR systems that have minimal network traffic. These 26 on-premises applications could be the ERP system, manufacturing execution systems, data warehouses, and other applications that produce large amounts of traffic. In terms of traffic, the dependency could be as high as 15 percent.

The reality vs. perception 15% disparity completely changes SDWAN financial projections. Organizations can expect a significant reduction in the number of MPLS circuits when majority traffic moves to direct internet routing at a sixty percent cloud dependency. With 15% cloud dependency, MPLS demands barely change as the predominant traffic flows still need data center connectivity. A single organization can come to different conclusions about the same SDWAN proposal, based on different methodologies that drive the evaluation.

Typical on-premises applications that are traffic-intensive include enterprise resource planning (ERP) systems handling core business transactions, manufacturing execution systems coordinating production processes, data warehouse and analytics platforms processing large volumes of data, legacy systems with batch processing needs, and industry-specific applications with extensive data exchange needs. These



systems are often enterprise investments that have been in place for years, and are tightly coupled with business processes. Their traffic volumes are far greater than any amount of cloud applications that may be in inventories, and they are lightweight cloud tools.

Some of the typical cloud applications that can be considered low traffic are human resources information systems used occasionally to record transactions, project management applications that share tasks and status updates, collaboration applications that produce mostly text-based communications, electronic signature applications and document management applications, and departmental productivity applications with small user bases. These applications, while useful to the operation, do not create significant network bandwidth requirements compared to the core transaction processing applications.

6.3 Requirements for Sound Assessment

Financial teams that are looking at SDWAN proposals should insist on a cloud dependency assessment based on traffic, not on applications. Ask for reports of network traffic analysis that reveal true bandwidth usage to destinations. If the IT department isn't able to give you this information, be highly skeptical of any savings projections.

Not all organizations have network monitoring capabilities, which are needed for traffic volume assessment. Correct measurement includes using measurement tools at the egress points of the network, collecting traffic flows for representative time periods, categorizing traffic based on destination type and determining the percentage distributions correctly. If these capabilities do not exist, it may be necessary to invest in monitoring tools or hire experts to conduct assessments prior to evaluating SDWAN.

The representative time periods are important in terms of accurate measurement. Network traffic is different during different times of day, days of the week, and stages of the business cycle. The volume and transaction flow at the end of the month is different from the mid-month volume. For industries that have cyclical demand, traffic characteristics vary with the season. Holiday times, vacation seasons affect user behavior and application usage. The dependency calculations are incorrect if the assessments are made on nonrepresentative periods of time. Ask for evidence of measurement times and that the periods captured are representative of typical operating time.

The classification of destinations is a complex methodology that needs to be well carried out in order to obtain meaningful results. Traffic to public cloud provider IP ranges may be for cloud applications, but also could be hybrid application cloud components talking to on-premises systems. If the applications running at SaaS vendor locations are dependent on enterprise data center resources, then the traffic to SaaS vendor locations is not considered to be true cloud dependency. Classification methodologies should include information on how hybrid patterns were identified and handled in percentage calculations.

Financial teams that don't have technical expertise can't check the accuracy of the traffic analysis themselves, but they can judge if IT teams are using the proper analytical rigor. Ask for evidence of measurement processes, time periods used, categories used, and assumptions made. These elements can be clearly communicated by teams that are doing thorough analysis. Many teams using inappropriate methodologies are unable to provide a clear explanation of their methodology, and recognize limitations in their evaluations. The failure to clearly explain methodology is an indication of analytical weakness and should raise doubts about conclusions.

If IT departments are unable to deliver traffic-based dependency assessments, then organizations must either invest in the right measurement capabilities or proceed with the understanding that there is some



degree of uncertainty. Proper measurement costs are usually a small percentage of proposed SDWAN spending. If organizations aren't willing to invest in learning about their own traffic patterns, then they should ask themselves whether they have enough operational visibility to effectively manage network transformation, no matter what SDWAN financial projections may be.

7. THE SALESPERSON PLAYBOOK RECOGNIZING MANIPULATION TACTICS

Networking sales people are under a lot of pressure to make sales. Commission models have a tendency to encourage sales, not results. This provides an incentive to make optimistic forecasts and present only selected information. Knowing how the system works and what makes salespeople tick can help the finance team assess proposals with the right level of skepticism, without implying that the salespeople have bad intentions. Salespeople work in systems that are geared toward making money and their sales techniques are a reflection of the system and not of their own dishonesty. Common patterns can be spotted and this can help in negotiations and safeguard the interests of the organization from systematic optimism in standard sales approaches.

There are a number of industry characteristics that make the networking industry have particularly aggressive sales practices. There are significant commission opportunities for successful representatives, particularly with high value contracts that have multi-year terms. Customers can't evaluate claims independently because of the complex technical requirements. The constant change of technology makes it difficult to rely on past experience, which decreases customer trust in their own judgment. As markets mature, competitive pressure among vendors increases as market growth demands displacing incumbent solutions. All these forces add up to sales situations where manipulation thrives, whether or not the individual vendor is ethical.

7.1 Understanding Common Manipulation Tactics

By knowing the common manipulation tricks, financial teams can safeguard the interests of the organization. The following tactics seem to be common among all vendors, irrespective of their reputation and market position. Even the most well-meaning vendors have salespeople who are under quota pressure and are willing to go to any lengths to make the sale. Individual representatives may sincerely believe their forecasts, but unconsciously use optimistic assumptions that are in their own best interests. Identification of the patterns leads to the right kind of skepticism without the need for assumptions about intent.

Financial teams may not have technical expertise to make their own assessments of networking proposals. This information imbalance is an opportunity for manipulation techniques that use technical complexity to hide bad economics. The following tactics are specifically aimed at this lack of knowledge, and create impressions that are divorced from reality with the help of technical jargon and complex calculations. The best way to fight these tactics is for finance professionals to demand transparency, documentation and verification and not try to be technically adept enough to evaluate on their own.

7.1.1 The Benchmark Bait and Switch

Often, salespeople will share case studies with you demonstrating huge savings at other organizations. These case studies always include businesses in which SDWAN value can be realized to the maximum extent they are often already heavily dependent on cloud, have costly existing MPLS contracts, and have several international locations with complex routing requirements.

There is a systematic selection bias for cases. Vendors gather success stories from implementations that are having exceptional results and unsuccessful implementations produce no promotional content. The resultant case study libraries are most extreme positive experiences, not typical experiences. If these extraordinary instances are offered as typical examples, it will deceive customers about what they can expect to happen in their own situations. A case study that demonstrates a 70% cost savings in a very cloud-centric company doesn't tell you what will happen in a moderately cloudy company that is predominantly on-premises.



Fig -6: The Salesperson Play book Recognizing Manipulation Tactics

Salesmen strategically use case studies to establish anchoring effects on customers' expectations. Starting with hugely successful savings numbers sets a baseline that makes subsequent projections look paltry. A customer who was originally shown 70% savings examples might think that a projected 40% savings for their own organization is a **safe** number and not an optimistic one. This anchoring manipulation is irrelevant of the case study organizations and the prospect's real situation.

Case study information can be so specific that it's easy to miss out on the elements that lead to outstanding outcomes. Presentations highlight percentage savings, but fail to mention that the baseline cost was very high, that the contract was timed to get the maximum benefit, or that something was different about the organization that allowed it to achieve a result that could not be replicated at other organizations. The impressive percentage was the product of conditions that are not likely to be found in your organisation, but the presentation formats don't allow you to see these critical differences. Comparisons give no valid predictive information without detailed knowledge of the characteristics of the case study organizations.

Financial teams should ask to see case studies from organisations that have a similar percentage of cloud dependency, geographic distribution and similar application portfolios. Case studies of generics are not predictive of your situation. If sales reps are unable to offer you any case studies that are relevant, then that



tells you something about the likely results of other businesses like yours. The vendor's lack of success with similar organizations indicates that your profile may not match those that allow for value to be realized.

Be specific when asking for information that can be compared to something. Ask case study organizations for cloud dependency percentages, number of locations, cost of the baseline circuit, and implementation time. Check these attributes against your profile and determine if they're relevant. If the impressive headline numbers are from a case study from an organization with fundamentally different characteristics, that case should be given zero weight in evaluation. Case studies are the only valid evidence for investment decisions, if they are successful in situations similar to yours.

7.1.2 The Hidden Cost Omission

Many standard SDWAN proposals don't include all the cost categories that should be included in total cost of ownership. Frequently overlooked are professional services for implementation, training of IT staff, management platform licensing fees, security service subscriptions needed for breakout to the internet, additional bandwidth costs for backup circuits, and vendor support fees.

SDWAN implementation professional services costs range widely depending on the complexity of the organization, existing infrastructure and internal capabilities. Simple deployments where the organizations have standard configurations and experienced staff may not need a lot of external help. Professional services engagement may be extensive when dealing with complex environments with legacy equipment, non-standard configurations, and limited internal expertise. Sales proposals often represent a low level of professional service's needs, based on the assumption that the situation in the organization is minimal, and understate the total implementation costs.

Training investments help staff to use new platforms effectively, and are not often included in initial proposals. SDWAN platforms are very different from traditional networking equipment, and they require new skills to configure, monitor and troubleshoot. If a staff member is not properly trained, he or she will not be able to take full advantage of the platform's capabilities and could cause operational issues as a result of misconfiguration. Training costs are the direct costs of the training program, staff time away from program duties, and productivity losses during learning curves. These investments are true costs and should be factored into total cost of ownership calculations.

Licensing fees for management platforms are often independent of SDWAN licensing. Vendors often split the fundamental connectivity functionality from the more complicated management, analytics and automation tools. Initial proposals can be for only base licensing with the assumption that customers will need advanced capabilities later. In considering proposals, seek clear assurance that the licensing quoted will cover all the capabilities demonstrated during sales presentations and all the features assumed in operational improvement projections.

Security service subscriptions are required when SDWAN deployments are used to allow direct breakout to the internet at branch sites. In the past, all traffic was forced through data center security stacks, which protected traffic without any security investments at the branch level. For direct internet routing, security needs to be provided at every breakout point, usually in the form of cloud-based security services, and the cost of these services is based on subscription. These costs are a direct consequence of SDWAN architectural change and should be included in true total cost calculations.

In proposals that focus on saving money on primary circuits, the bandwidth costs of backup circuits are given little consideration. A common characteristic of SDWAN is that backup links at every location are required to



provide failover capability. Backup circuits can be implemented with less costly technology than the primary circuits, but they are incremental costs not included in traditional single-circuit implementations. Proposals that indicate significant savings per location from cost reduction in primary circuits may not account for the addition of backup circuits, which can significantly affect the net savings.

Ongoing vendor support fees are paid throughout the life of the operation but may not be emphasized in the initial proposals that are based on acquisition costs. Support agreements include software updates, technical support, and hardware replacement. These fees are more of a necessity than an option when it comes to enterprise networking, as it is important to have strong support arrangements in place. When properly accounted for, multi-year support costs can be as high or higher than initial acquisition costs, and have a dramatic impact on total cost of ownership calculations.

If comparing SDWAN proposals to current network costs, make sure that the cost categories are the same. Ask for a detailed breakdown of costs and check if they're complete using a standard format. Create detailed cost category checklists that include all categories of expenses and ask vendors to specifically answer each one. Costs of zeroes should be specifically asked for in categories where this occurs. If costs are "included" it must be identified where the costs are listed on the itemized breakdowns.

7.1.3 The Capacity Planning Deception

Many times, salespeople will offer bandwidth tiers based on the average bandwidth utilization and not the peak bandwidth. Average utilization for existing MPLS circuits could be 30%, which results in the suggestion that internet circuits be reduced in size. But networks need to have the capacity for peak times, not the average.

Average utilization statistics are misleading since network demands fluctuate significantly over the course of operation. Business applications may be running during the day and backup applications might be running at night. Peaks above normal day to day processing because of month-end processing. Many industries have periodic demand surges due to seasonal fluctuations in business. There are short-term but high bandwidth needs for video conferencing, transferring large files and deploying software. If the circuits are sized for average utilization, then during peak periods, there isn't enough capacity, and performance suffers when it is most needed.

Thirty percent is the average utilization sales guys point to as proof of overprovisioning, but ninety percent could be the peak utilization during critical times in your business. Undersized replacement circuits would be fine for average workloads but would not be able to meet the demand during peak workloads. This failure manifests itself for the user as slowdowns, dropped connections and poor video quality during critical meetings. These performance effects lead to productivity losses, user dissatisfaction and ultimately to circuit upgrade, thus negating the projected savings and causing disruption to the operation.

The sizing of circuits is based on the network engineering principles and involves designing circuits for maximum demand with a margin for unforeseen demands. The industry standard for keeping utilization below a certain threshold in peak times is seventy percent or less, to ensure that the performance is responsive and it can handle bursts of traffic. Circuits of this size will need significantly more capacity than average utilization calculations indicate, and thus the dramatic savings projected by average-based proposals is reduced.

Ask for analysis of 95th percentile usage trends. Scale size replacement circuits, and re-estimate savings using realistic bandwidth requirements. The 95th percentile measurements represent the loadings that are



not met 95% of the time and are good indicators of what is required for peak operation, excluding unusual spikes. This approach can give realistic capacity requirements which will guarantee sufficient performance in normal peak periods.

Initial savings proposals based on the proper circuit sizing may be significantly different from recalculated savings. If circuits are sized according to the average utilisation, a fifty percent circuit cost reduction can be indicated on the proposal, but if the circuits are sized for peak demand, then only a twenty percent circuit cost reduction can be indicated on the proposal. This re-calculation provides real financial opportunity and not a false savings that would need to be corrected in the future by upgrading circuits. Be wary of any capacity savings claims without realistic capacity planning.

7.1.4 The Contract Term Arbitrage

MPLS contracts are frequently volume based and have penalties for early termination. Sales proposals often assume that MPLS can be eliminated right away, but do not consider these contractual obligations.

The contracts of telecommunications services are usually of a 3-5 year duration, and based on volume commitments. Carriers will provide good rates per circuit for guaranteed revenue in the form of minimum purchase agreements. When these contracts are terminated early, penalty clauses are in place to make up for lost revenue the carrier would have expected. These penalties can be as high as the balance of the contract and if a contract is terminated two years early then the circuit costs for the two years may need to be paid.

Sales proposals that predict the immediate elimination of MPLS are oblivious to these contractual details. The proposal reflects year one savings based on the assumption that all MPLS circuits end up immediately, but in fact based on the contracts, it is not possible to terminate without paying a significant penalty payment. Year one costs are not only the new SDWAN costs, but also the ongoing MPLS costs possibly higher than the former costs and not necessarily savings. Proposals that disregard existing contracts put forward seriously misleading financial projections.

Transition planning is further complicated by volume discount structures. MPLS contracts will generally have tiered pricing, where the per-circuit pricing will be lower at higher volume levels. The reduction of circuits can result in circuits being moved to higher price bands, which can lead to higher per-circuit costs for circuits that are retained. A contract with a great rate at 50 circuits could be much less favorable if there is a transition to 30 circuits. These pricing tier impacts diminish the net savings from circuit elimination.

Check any MPLS contracts that are in place for minimum contract terms, early termination fees, and volume discounts. Simulate transition scenarios that meet current commitments and maximize savings, while staying within contractual limits. For realistic transition planning, MPLS circuits may need to be kept longer than ideal to avoid penalties or minimum volumes may need to be kept to maintain pricing tiers. As a result of these constraints, the savings that can be achieved is less than that proposed under the assumption of unconstrained optimization.

SDWAN investment timing decisions should be impacted by contract timing. Organizations that are nearing the end of their MPLS contract can shift without penalties, and realize full projected savings. Companies with long-term contracts might prefer to wait until the end of the contract to invest in SDWAN, as that will yield a more favorable financial result than implementing SDWAN now with the existing MPLS contracts in place. Explicitly consider timing options, rather than taking proposals for granted, assuming immediate unconstrained transition.



7.1.5 The Phantom Feature Attribution

When it comes to SDWAN, many sales presentations claim that SDWAN can save you money due to the features that offer marginal value in your environment. Quality of service optimization is good, but if your network isn't congested, it isn't going to help very much. Application visibility dashboards are helpful but don't necessarily save money.

The basic principle of feature attribution manipulation is to generate the impression of value, but not necessarily the financial impact. Sales presentations are filled with amazing capabilities, backed up with pictures and technical explanations. Features seem to be important in abstract and customers think that there must be a monetary benefit, but it's not stated. The net effect of many helpful features is to build up hope that there will be meaningful savings, but there's no particular feature that can be shown to save any dollar amount.

Quality of service optimization allows to prioritize critical applications during congestion periods thus giving bandwidth preference to important traffic flows over less critical ones. This can be very useful in busy networks where many applications compete for bandwidth and thus affect application performance. But, networks which are not congested do not suffer from performance degradation which needs optimization. Quality of service attributes for uncongested networks are not of value, but of insurance against future congestion. Sales presentations that focus on the quality of service benefits, but don't confirm what congestion is like today, are assigning value that may not exist in your environment.

Application visibility dashboards provide traffic patterns, performance and usage metrics in graphical format. These tools provide a greater understanding of the operation and can be used to aid in troubleshooting. But there is no direct cost savings from visibility. Knowing the traffic pattern is only useful if that knowledge can be used to take action for savings. When proposing savings from visibility features, it is important to detail what visibility enables and how that translates into measurable financial benefits. When visibility isn't followed by action, interesting information is created, but there's no reduction in cost.

Centralized management features make it easy to manage distributed locations with a single interface and automated configuration deployment. These capabilities decrease the amount of effort required to manage the network. But, savings in money only occur if the reduction in effort allows staff to be reduced or if major reallocation to higher value activities can occur. Organizations that are not reducing staffing levels in spite of the increased efficiency will not reap any monetary benefit. Any proposals that are claiming management efficiency savings should state the amount of staff reduction or quantifiable reallocation value.

For every feature identified as a savings opportunity, ask for specific financial savings. Total savings projections should be equal to the sum of the feature claims. Ask for specific details that demonstrate how each feature will help achieve savings and the assumptions that were used to calculate those savings. If the benefits are not quantifiable, then features should be given zero weight in financial evaluation no matter how they may appear. When evaluating investment returns, only features that have been shown to have financial impact should be considered and documented.

The reconciliation requirement provides transparency because it creates an obligation to report on the overall forecast. Sales presentations can mention many great things about the product, but not present any of the numbers, making it seem like there is a lot more value. Reconciliation makes it necessary to show explicitly how much of the savings is due to what feature; this makes it possible to verify if the benefits claimed justify the savings projected. Claims that can't be matched up with totals in the projections show that projections are not based on demonstrable value they should be ignored even with impressive features.

8. THE FINANCIAL CALCULATION FRAMEWORK

Financial teams need to have concrete methodologies to validate SDWAN savings claims. The cost analysis can be done using the following framework.

8.1 Step One Establish Baseline Network Costs

Develop a comprehensive list of existing network spending, such as circuit costs MPLS, internet, point to point, etc., equipment depreciation and maintenance, network operations staff assigned to WAN management, vendor support contracts, security services that guard network traffic, and bandwidth overage fees. Add up these categories to determine actual current state annual network cost. This is more than what IT departments would report, as they might not include labor, depreciation or adjacent costs in their calculations.




Fig -7: The Financial Calculation Framework

8.2 Step Two Assess Cloud Dependency Accurately

Ask network traffic analysis to display the amount of traffic by the destination category. Determine the percentage of total WAN traffic to corporate data centers, cloud destinations and internet endpoints. If IT is not able to provide this analysis, then use a third party to do traffic assessment prior to SDWAN assessment. The price of a correct assessment is unimportant compared to the price of a wrong investment.

8.3 Step Three Model Future State Costs Comprehensively

Project SDWAN scenario costs, such as new equipment capital expenditure or subscription fees, implementation professional services, ongoing management platform licensing, Internet circuit costs at the appropriate bandwidth tiers, any retained MPLS circuits, new security service subscriptions, increased IT labor

for transition and ongoing management, and training expenditure. Ensure cost modeling covers the same time horizon as savings projections, typically three to five years.

8.4 Step Four Calculate Net Present Value

Use appropriate cost of capital to discount future cash flows to present value. The return expectations for SDWAN investments should be similar to those of other capital investments. Many IT proposals fail to take into account the time value of money and offer total savings figures that are undiscounted and make the proposal appear more attractive financially. Use the same financial evaluation techniques for all capital requests.

8.5 Step Five Sensitivity Analysis

Simulate scenarios with different percentage of cloud dependency. What are the outcomes if cloud migration is not as rapid as expected? What if bandwidth needs are more than originally thought? How about if MPLS termination is slower than scheduled. If the proposals only pay off under optimistic conditions, then it should be a cause for concern. Favorable investment returns for a wide range of scenarios.

9. CASE ANALYSIS WHEN SDWAN FAILS FINANCIALLY

Looking at some of the real-life examples in which SDWAN investments did not yield the expected returns will highlight the typical patterns of failure.




Fig -8 : Case Analysis When SD-WAN Fails Financially

9.1 Manufacturing Organization with Legacy Systems

A manufacturing business with seventy five sites authorized an SDWAN project that is projected to save forty percent on network costs. The IT manager shared some great analysis that demonstrated that MPLS could be



eliminated at most sites within eighteen months. But manufacturing execution systems, quality management applications and inventory systems all were located in corporate data centers. These applications had to have guaranteed latency and zero packet loss, which internet circuits could not guarantee. After implementation, the organization continued to operate its MPLS at 60 out of 75 sites and deployed internet circuits and SDWAN equipment at all of the sites. The total network costs grew by thirty two per cent, not down. The IT manager said that the applications had an unforeseen requirement, but the applications should have been assessed to find the requirement prior to its approval.

9.2 Financial Services Firm with Regulatory Constraints

A financial services organization estimated a sixty per cent savings with SDWAN implementation. The proposal was based on the assumption that Office 365 and other cloud applications would be directly broken out to the Internet. All traffic inspection was to be centralized via the regulatory requirements. Breakout to the Internet was not allowed, and is the main SDWAN cost saving mechanism. The organization deployed SDWAN, but kept all traffic going through the data center security stacks. The real savings amounted to four percent, mainly due to bandwidth optimization features. Three years of savings were used to cover the implementation costs, leaving a negative net present value.

9.3 Retail Chain with Bandwidth Misjudgment

A retail business with 400 sites was looking at replacing 20 Mbps MPLS circuits with consumer grade broadband and decided to try SDWAN. The annual savings were more than two million dollars. The average MPLS utilization was eighteen percent, indicating that there was a lot of over provisioning. But during holiday times, burst capacity was needed during point of sale transactions. Consumer broadband does not offer the same committed information rate guarantees as does MPLS. The first holiday season after the implementation had transaction processing failures. Higher tier business Internet circuits were procured in emergency mode, which meant that projected savings were eliminated. Failure of transactions led to additional unmeasured costs as a result of reputational damage.

10. CASE ANALYSIS WHEN SDWAN SUCCEEDS FINANCIALLY

It is also essential to know the success patterns for evaluation.

10.1 Technology Company with Cloud First Architecture

A tech firm with two hundred locations all over the world, which had seventy three percent of its traffic volume dependent on the cloud, approved SDWAN. The primary infrastructure was on AWS, CRM on Salesforce, HR on Workday and productivity on Office 365. The data center traffic was mainly composed of legacy financial systems that were slated for cloud migration in the next 18 months. It deployed SDWAN with internet primary transport at all locations, and reduced capacity MPLS only at the regional hub sites. Realized savings were 51% of prior network investments. Fourteen months payback on implementation costs. The remaining MPLS circuits were eliminated, furthering savings as cloud migration was completed. Some of the key success factors involved realistic cloud migration roadmap with reduced remaining data center traffic, accurate cloud dependency assessment, and committed cloud dependency assessment.

10.2 Professional Services Firm with Distributed Workforce

With one hundred twenty offices around the world serving customers around the globe, a consulting firm analyzed and determined that it had sixty five percent cloud dependency and had significant international traffic patterns, which is why it chose SDWAN. The old network design involved backhauling all traffic to the US




Fig -9 : Case Analysis When SD-WAN Succeeds Financially

data center to connect with cloud destinations and internet resources. Poor application performance for international locations because of latency. With SDWAN in place, local internet breakout at all locations was achieved. The performance of cloud application was greatly enhanced. Ninety out of one hundred twenty locations had their MPLS circuits removed. Forty two per cent cost reduction was achieved in financial results and significant non-financial benefits from performance improvement. The payback for implementation was after 11 months.

11. COLLABORATIVE EVALUATION BUILDING FINANCE AND IT PARTNERSHIP

The frameworks discussed in this paper are for verification, but the best results come from a collaborative evaluation process, not an adversarial review process. Financial teams have capital allocation experience and IT professionals have technical and operational insight. Both views are not enough to make the best decisions. Effective organisations set up joint evaluation processes in which finance and IT add unique value. Financial scrutiny should be seen as a validation of good proposals and not as an obstacle or criticism for IT teams. Financial teams must acknowledge that there is a real degree of technical complexity and that IT recommendations may be a valid organisational need.

There are a number of practices that enable productive collaboration. Joint requirement gathering involves taking all financial and technical aspects into account from the start of the project. Joint ownership of forecasts is a way of making people more accountable across functions, which helps to minimize the incentive for optimistic or pessimistic bias. Organizations learn and enhance their estimation accuracy with regular calibration sessions that allow them to compare projections to outcomes. In the case of disagreement, try to concentrate on assumptions, not conclusions. The main reason for most projection conflicts is the differing assumptions regarding cloud migration timelines, bandwidth needs, or

implementation complexity. Explicit assumptions allow for productive discussion of the validity of the assumptions, instead of unproductive discussion about the conclusions.




Fig -10 : Collaborative Evaluation Building Finance And IT Partnership

It is also important for organizations to be aware that IT leaders are frequently under pressure from a number of sources. Vendors, peers, and technology media all provide environments where some type of technology adoption is a must, regardless of the fit of the technology in the organization. When the financial verification is in place, IT leaders have the cover they need to resist the inappropriate pressures; they are not being personally reluctant but are doing so because it is a requirement of the organization. It's not about catching incompetence or thwarting deception. The aim is to make sure that technology investments are evaluated as carefully as other capital allocation decisions. Working together, finance and IT achieve better results both for themselves and for the organisation as a whole.

12. SECURITY INTEGRATION AND TOTAL RISK-ADJUSTED COST ANALYSIS THE HIDDEN FINANCIAL DIMENSIONS OF SD-WAN INVESTMENT

Traditional SD-WAN financial assessments only consider circuit costs and ignore security integration costs and risk exposure costs. This oversight results in incomplete ROI calculations, and unexpected budget overruns, which can turn what would otherwise be a good investment into a financial disappointment. When considering SD-WAN proposals, CFOs need to take financial considerations related to security into account to come up with realistic total costs that reflect the reality of the organization, not the vendor's best-case scenario. The threat landscape is complex and constantly changing, and so are the regulatory requirements, which means that there are financial risks that can't be accurately represented in superficial comparisons of cost, and that security-integrated analysis is crucial to making good investment decisions.

12.1 The SASE Convergence Factor Understanding the Security–Network Integration Imperative

New SD-WAN implementations are more often now becoming integrated with Secure Access Service Edge (SASE), where networking and security are combined and deployed in a single platform, which changes the way that networking and security are traditionally purchased and budgeted. This convergence is not just a sign of the times; it's a testament to the understanding that network transformation without security transformation results in vulnerabilities that are ripe for threat actors to exploit. Those that take SD-WAN as just a networking project are surprised to find that security needs have a significant impact on the overall investment size and operating expenses.




Fig -11: The Hidden Financial Dimensions of SD-WAN Investment

SASE is a combination of SD-WAN networking and cloud-based security such as ZTNA, CASB, and FWaaS. 60% of SD-WAN deployments will be SASE by 2026. This convergence trajectory implies that those organizations that are deploying basic SD-WAN solutions today will almost certainly need to integrate security into the lifecycle of their investments, and even if they are not planning to deploy full SASE, security cost considerations will be a factor in their SD-WAN investments.

The cost ramifications of this convergence are starkly evident in cost differences per site. SASE costs \$300–\$600 per site per month, vs \$100–\$300 per site per month for SD-WAN alone, and combines multiple security solutions. In many cases, this apparent cost premium can hide the actual economics as organizations choosing basic SD-WAN need to plan for a full-featured security infrastructure at each branch. Many organizations end up having to deploy more security appliances at every location, which can cancel out cost savings and create management challenges. This is frequently what spurs interest in SASE solutions that include full security from the ground up.

Distributed security requirement includes several categories of technology, which add up to a significant amount of money when purchased individually. Organizations need to allocate funds for next-generation



firewalls that offer more than just packet filtering and intrusion prevention systems that can detect and block more advanced attack patterns; they also need to invest in secure web gateway solutions that can secure users accessing internet resources directly, and zero-trust network access solutions that can provide identity-centric security models instead of the traditional VPN approach. Basic SD-WAN features limited inbuilt security, but it was not designed with security in mind, even though it does offer encrypted connections. Full protection may necessitates using the solution in conjunction with third-party security services, further increasing costs.

The total cost of implementing security components individually is significant when all components are considered. The TCO of SASE is neutral or positive when replacing 3+ security point solutions, although it can be 40-80% higher per site than SD-WAN alone. But if organizations find out about security requirements after they have selected the SD-WAN platform, they must decide whether to live with the security gaps and risk a potential breach of the organization, ask for an additional budget increase, thereby losing credibility with the finance organization, or abandon the partially completed SD-WAN deployment, thereby incurring sunk costs with no return.

When financial teams are considering SD-WAN proposals, they should require full security cost integration from the beginning. What matters is making sure that you've taken stock of where you're currently spending money on security. This inventory should include all the security components that are currently securing the network traffic, the cost of these components, and the impact SD-WAN architecture changes have on security requirements. Any proposal that touts SD-WAN savings without considering the costs of integrating security should be considered incomplete and rejected for investment decision making.

The long-term cost trajectories are greatly influenced by the integrated versus disaggregated architecture decision. Point solutions like SD-WAN, NGFW, SWG and VPN cater to particular networking and security needs. The complexity and cost of IT infrastructure is driven by the need to purchase, size, scale and maintain each solution individually. SASE is a transformative alternative to those legacy technological silos, offering a globally distributed cloud service that eliminates the need for on-premises and virtual point solutions, and can be cost-effective, scalable and agile.

Cumulative spending of all SSE and SD-WAN is expected to total \$97B between 2025 and 2030, according to Dell'Oro Group. The spending is almost triple the total SASE outlays in the previous five years, and indicates a fundamental change in how enterprises are designing networks. Security policy is being defined more and more upstream through risk governance requirements, and connectivity is designed to implement this policy uniformly across users, applications, and locations. This market trend validates the prevailing architectural trend of security-network convergence, and the need for security cost integration for any SD-WAN investment with multi-year timeframes.

12.2 Risk-Adjusted Cost Modeling Quantifying Security Exposure in Financial Terms

A key SD-WAN architectural feature that allows for cost savings by displacing MPLS is direct internet access at branch locations, which opens up security exposure that traditional MPLS architectures do not have by isolating networks. This architectural change necessitates explicit risk quantification to make sure that the circuit cost savings are not disproportionate to the breach exposure, which could result in overall investment economics being undermined.

In 2026, the cost of a single data breach is \$4.44 million on average around the world, and \$10.22 million in the US, which is the highest cost recorded. It takes on average 241 days for organizations to identify and contain a breach and 68% of breaches still have a human element. These breach cost figures are key for assessing



investments in SD-WAN security. Organizations that are considering security cost savings need to factor in the potential cost savings, as well as changes in the likelihood of breach if they do not configure SD-WAN security correctly.

The latest findings reveal that the cost of breaches is still rising, despite the increase in detection. In 2026, one in four malicious breaches was AI-enabled, representing a 56% increase, which averaged US\$6 million. The average breach cost for 2026 was a record US\$4.99M, 12% higher year-over-year. The global average breach cost was US\$5.5M, whereas the US average cost was over double this figure at US\$11.5M. As these costs increase, so do the stakes in security investments on organizational financial performance.

The architecture of the network can directly impact the likelihood of a breach and the size of the cost. In 2025, breaches that occurred across multiple environments were particularly expensive, averaging USD 5.05 million, while on-premises breaches averaged USD 4.01 million and hybrid cloud breaches averaged USD 3.80 million. By its very nature, SD-WAN deployments result in multi-environment architectures between branch locations, cloud services, and data centers, making organisations more vulnerable and exposed to increased risk levels that demand greater security investment to keep risk levels acceptable.

SD-WAN is not inherently a risk to breaches when implemented correctly with the proper security controls, but vulnerabilities can certainly exist in improper implementation that result in quantifiable financial risks. IBM's 2025 data measures that cost at each day of access that goes undetected. That's an average of 241 days per life cycle, and \$4.44 million per average life cycle, or about \$18,400 per day of dwell time. Without proper security monitoring and threat detection, organizations adopting SD-WAN risk longer dwell times, which can significantly add to the cost of breaches.

CFOs should mandate that IT teams document security architecture changes that happen as part of SD-WAN deployment and formally risk assess the changes by qualified security leadership. These evaluations need to measure the impact of SD-WAN architectural shifts on threat surface, detection and incident response processes. Those breaches that were contained within 200 days cost an average of \$3.87 million, while those that took more than 200 days to detect resulted in an average penalty of \$5.01 million, which amounts to \$1.14 million for each day it takes to detect. Security investments that speed up the detection process have a measurable ROI in terms of the cost exposure of the breach.

There are also financial implications that need to be taken into account, such as insurance implications. Insurers are also looking at the security posture of an organization when determining coverage terms and premiums, and cyber liability insurance premiums are increasingly based on network architecture, security controls and incident response capabilities. SD-WAN deployments that result in a loss of security visibility or leave the internet exposed can result in premium increases or coverage restrictions that can offset the apparent network cost savings. It is important to check with insurance carriers prior to making SD-WAN architectural decisions to understand coverage implications.

The risk adjusted cost model should include probability weighted breach exposure, leveraging the breach frequency data from the industry and organizational security maturity assessments. US organizations faced an average cost of \$10.22 million per data breach in 2025, marking the highest average cost of data breach to date, up 9% from \$9.36 million in 2024. For 15 years running, the US has been the top country when it comes to breach costs. The worldwide average decreased to \$4.44 million, which was 2.3× the cost in the US. The breach cost exposure is especially high for organizations based in the U.S., which makes it even more important to make investments in security integration.



The evidence is clear that security investments can make a difference when breaches happen. In 2025, breach costs and resolution times were greatly diminished by AI and automation. Those that did not have these tools had the highest average cost of USD 5.52 million per breach, while those organizations that had extensive use of these tools reduced costs to USD 3.62 million per breach. When it comes to ROI, the \$1.9 million difference in cost between organizations with and without security automation should be enough to justify investments in security for SD-WAN deployments.

12.3 Hidden Implementation Costs Beyond Platform Licensing

Implementation complexity is a cost factor that is consistently underestimated and has been identified as a key factor that undermines SD-WAN savings in industry research. While most vendors focus on zero touch provisioning and quick deployment, in most organizations, deployments are staged and involve a high level of professional services investment in addition to platform licensing fees. Often, financial assessments based on vendor-provided projections underestimate implementation costs because vendors have an incentive to underestimate total cost of ownership.

The often touted 50% cost savings is achieved by converting non-critical traffic from MPLS to broadband. But additional costs such as installation, professional services, and training can increase initial budgets by 30–40%. This significant implementation overhead has a significant impact on the first year's economics and a significant impact on the time to ROI based on platform only cost projections.

The managed services decision has significant impact on overall cost paths through the investment lifecycle. Fully managed SD-WAN is priced between \$125–\$375 per site per month, and includes full outsourcing and 24/7 NOC support, which is ideal for lean IT teams. Managed SD-WAN services cost \$125–\$375 per site per month, based on the level of support. When considering self-managed vs managed deployment options, organisations need to be realistic about the management complexity and if they underestimate, they will end up with a reduction in service quality or unplanned managed services procurement at an unattractive price.

Design, proof of concept, deployment and migration can be costly professional services. The cost of staff training for new operational models such as policy-driven management and troubleshooting hybrid paths should be planned for to prevent hidden costs at cutover. Often these professional services expenses are overlooked when procurement is calculating the cost of the platform, as these expenses are one-time implementation costs.

Implementation will usually be charged as a percentage of first year cost, usually 10–20%, and not as a flat fee. The final number will vary based on the number of sites and complexity of deployment. This pricing by percentage can result in implementation costs that increase with the size of the deployment, and that superficial cost comparisons don't reveal, especially in large enterprise deployments.

Another often overlooked cost item is underlay circuit upgrades. The performance of SD-WAN also relies on the quality of the underlying Internet service, and if the region has a poor broadband service, it may cost more than originally thought to use dedicated circuits. Some of the hidden costs of SD-WAN are hardware costs ranging from \$1500 to \$12,000 per site based on bandwidth requirements implementation and professional services ranging from \$5000 to \$25,000+ training IT staff additional security features, which are often sold separately and internet circuit upgrades, which may be required.

SD-WAN deployments have hidden costs that can undermine the cost savings if not dealt with early. Planning and often professional services will be required for the integration of SD-WAN with on-premise security, identity providers and cloud routing. For full feature compatibility, legacy appliances might require



replacement. Things get complicated when organizations have a hybrid environment during transition periods, where legacy and SD-WAN infrastructure need to be supported simultaneously, and operational cost is doubled.

Often forgotten costs include training and skills development, which can impact implementation success and operational efficiency. SD-WAN platforms demand new skills of the network practitioner, and current network staff members might not have the skills needed to manage them effectively. Changes to policy driven networking means retraining operations staff, and updating runbooks. This is a factor that can be underestimated, which can extend the stabilization period and increase incident costs.

Implementation and professional services fees are dependent on the vendor, and are usually 10-20% of the first year's costs. Circuit installation costs are \$500-\$2000 per location, depending on location and access needs. If you're self-managing SD-WAN, expect to spend about \$1,000 per site on hardware. These incremental costs add up and have a significant impact on the investment economics of multi-site deployments, compared to the platform-only projections.

Cloud-based SD-WAN deployments come with extra hidden expenses, often revealed after the deployment. Another hidden cost to keep in mind is that you may have to pay two egress fees. The northbound connections will be via internet IPsec tunnels and there will be charges for each GB traversing the internet, from the application's VPCs to the Transit VPC. If the traffic is going to a branch site, then there will be an extra per-GB charge when the traffic leaves the Transit VPC southbound. Effectively, the customer is getting charged double egress fees for each GB reaching a branch site. These egress fees can have a significant impact on organizations' total cost of ownership when they have high cloud-to-branch traffic patterns.

12.4 Vendor Lock-In Financial Implications Long-Term Cost Exposure

Switching costs, due to proprietary orchestration platforms and management tools, should be considered in long-term financial planning, but are often not evaluated during scrutiny. The initial deployment cost and projected savings may be the only concerns of organizations, but exit cost implications can have a major impact on future flexibility and negotiating leverage. Proprietary features can be a barrier to migration. Include potential future switching costs and interoperability testing in TCO modelling.

Before signing on for multi-year contracts, organizations should consider terms of the contract, data portability requirements, and competition. Projected migration costs at the end of the contract should be included in the total cost of ownership calculation, even if the migration is not planned now, as these costs have an impact on the vendor negotiating power during the contract. Vendors have found out that customers have significant switching costs, so they are not as inclined to keep the prices or quality of service competitive when it comes time to renew their contracts.

The termination fees for an MPLS contract are typically based on the amount of money still owed under the contract, and may be as high as the entire contract balance. Appliances have a lifespan of about 5-7 years and a replacement cost is within most contract lifetimes. Renewal rates are often higher than first term as first term contracts are discounted to secure the business. These contract dynamics can result in financial exposure that goes far beyond the costs of deployment, and which need to be dealt with by comprehensive financial analysis.

SD-WAN is still a rapidly changing competitive market, with various vendors consolidating, capabilities expanding and pricing models changing the economics of long-term investments. There can be a wide range of SD-WAN pricing. The top speed is 3x the lowest at 200Mbps. The difference is the same on a multiples



basis (2.5x) but much more significant on a dollars basis (\$496 vs \$1256 per month). This price volatility is a testament to the importance of having competition and not relying on one vendor.

The use of multiple clouds adds to the issue of vendor lock-in, as the architectures of such environments can vary, and proprietary solutions may have different approaches to handling them. If a second cloud provider is added to prevent vendor lock-in, to plan for business continuity or to use a new internal application, the architecture becomes much more complicated, with all the branch sites needing to be connected to both cloud providers. Either they have to hairpin traffic back to the on-premises data center, resulting in increased latency, or they have to create new tunnels, which will incur additional egress fees and have speed limitations. When considering SD-WAN platforms, organizations need to understand the impact of their selection on multi-cloud flexibility, and avoid solutions that might limit future architectures.

12.5 Recommended Financial Safeguards Comprehensive Evaluation Framework

CFOs should make sure to include systematic financial safeguards when assessing SD-WAN offerings, and that the evaluations reflect a full cost picture and not the rosy picture of the platform that the vendors might want to show. These safeguards ensure that the interests of the organization are safeguarded, and that there is full disclosure before making any investment commitment.

Comprehensive Security Cost Integration: Ensure detailed security cost projections, including all security components and costs throughout the deployment lifecycle. This integration should include the cost of hardware for security devices in case of using disaggregated security, the cost of all security software licenses across all locations, the ongoing management overhead of security operations and the integration cost of the security elements with SD-WAN platforms. Basic SD-WAN provides routing and optimization. Implementing integrated security via SASE firewall, CASB and ZTNA can double or even triple the costs, but it does away with the need for standalone security appliances at every location.

Risk Assessment Documentation: Formal documentation from qualified security personnel risk assessment documentation of how SD-WAN architectural changes impact the organization threat exposure. This assessment should measure changes to attack surface from direct Internet access, analyze the impact on detection capability from architectural changes, document changes to incident response procedures, and estimate changes in cost exposure for breach, based on the probability of the breach.

Implementation Cost Breakdowns: Provide implementation cost breakdowns that clearly separate platform costs from professional services, training, circuit upgrades and integration costs. Costs include appliance and hardware costs, which vary in their capabilities and throughput and security features in some deployments, extra firewall appliances may be needed. Software features, cloud management portals and updates are typically included in the licensing and subscription fees. Recurring fees. Support and maintenance, such as technical support, SLAs and regular patching. Outside expertise may be required for deployment and integration such as implementation assistance, employee training, LAN integrations etc.

Long-Term Cost Projections: Need cost projections beyond initial contract terms to include the costs of renewal pricing, hardware refresh cycles, and migration costs. There is a need for continuous investment in software updates, hardware refresh cycles and changing security threats. Subscription models can level out upgrade costs, while perpetual licenses can result in upgrade costs in the future. Sensitivity analysis for price changes, technological advances and different scenarios of organizational growth should be incorporated into these long-term projections.



Managed Services Evaluation: Realistically consider internal capabilities when evaluating self-managed versus managed deployment options. The true cost of self-managed SD-WAN is in staff time and expertise, as license and hardware costs are lower. Monitoring, change and fault fixing can consume a lot of a network engineer's time and for smaller teams, this time can be more valuable than the managed service that it replaces.

Vendor Comparison Requirements: Make sure to compare several vendors before committing to one, as the difference in capabilities and pricing will impact long-term value. At the same bandwidth, the price of SD-WAN can differ by as much as 2.5X between vendors (\$496 to \$1,256 per month at 2Gbps). Always shop around from 3 vendors and get the best price for your overall number of sites.

Total Cost of Ownership Framework: Adopt a holistic TCO approach that includes hardware, licensing, transportation, security, operational, and other hidden migration expenses. SD-WAN cost assessment is multi-dimensional, as it needs to factor in hardware, licensing, transport, security, operations and hidden migration costs. The CAPEX/OPEX ratio is altered by selecting self-managed or managed offerings. To make an informed decision conduct a thorough requirements analysis, develop a multi-year TCO model with sensitivity analysis and validate assumptions with a targeted pilot.

These safeguards provide a fuller picture of the cost that will be reflected in the financial evaluation, and will facilitate good investment decisions based on organizational needs, not vendor sales goals. Those organizations that deploy these frameworks avoid the budget busters and ROI failures that define incomplete SD-WAN financial evaluations.

13. THE ROLE OF IT LEADERSHIP KNOWLEDGE GAPS

Many IT managers approve or recommend SDWAN without the knowledge of the financial mechanisms that drive value. This lack of knowledge results in a number of unwanted trends. Traditional technology leadership development focuses on technical skills rather than financial skills and produces professionals who know a lot about the architecture of the system and how it works, but do not have the tools to rigorously analyze the economics of an investment. This imbalance can cause problems when IT management takes ownership of recommendations for major capital investments without the ability to analyse. The problem is when technology investments are made based on technical enthusiasm, instead of sound financial analysis. Finance leaders need to be aware of these knowledge gaps and put in place governance processes to overcome the gaps in analytical skills whilst honouring the legitimate technical expertise.

13.1 Vendor Dependency for Analysis

Many IT managers don't have analytical skills and instead have their vendors do total cost of ownership calculations. This obviously results in conflicts of interest. Vendors have a vested interest in providing optimistic forecasts, no matter how accurate they are. When IT teams don't have modeling capabilities, the pattern of vendor dependent analysis comes naturally. This is where vendors step in and offer sophisticated-looking analyses with detailed spreadsheets, professional presentations and compelling visualizations. These materials provide the illusion of thorough assessment, but really are for selling purposes. When vendors give analysis, it is always positive no matter the situation in the organization, because if it was negative, it wouldn't move the sales process.

The analysis provided by vendors is always unreliable due to structural conflicts of interest, no matter how good the vendor's integrity. The sales representatives are subjected to incentive plans that are based on the

number of sales made and not on the proper forecasts. They are dependent on bringing in money, not optimizing customers. Well-intentioned representatives will work in systems that encourage positive projections. In addition, there is organizational pressure to hit quotas and show pipeline progress which further undermines analytical objectivity. This is true, whether a vendor is known or unknown, or whether a vendor is a representative or not.

There are generally predictable vendor analysis trends, and finance teams ought to understand what they are. The savings rates always hit levels that warrant investment, no matter what the initial savings rate is. Assumptions are optimistic about the timing of circuit elimination, bandwidth pricing, and simplicity of implementation. Risk factors are given very little elaboration, but benefits are given a lot of elaboration. Sensitivity analysis to show robustness of projections under different assumptions is seldom presented. Once these patterns are recognized, it allows for the correct skepticism, without the need to have any deep technical knowledge.




Fig -12: The Role Of IT Leadership Knowledge Gaps

Financial teams need to insist on independent analysis. When IT is unable to make calculations in-house, use third party (neutral) assessments instead of vendor assessments. Independent analysis is an investment, but it can save much greater losses due to wrongful investments. Outside consultants, without vendor relationships, can take a more objective look at proposals, pointing out unrealistic assumptions and quantifying what can be expected in terms of actual outcomes. This investment in analytical integrity can pay off many times in losses avoided and negotiating leverage. Organizations that do not want to invest in an independent analysis should ask themselves whether investments make sense at all, if the evaluation is inadequate.

13.2 Technology Focus Over Financial Focus



Sometimes, technically driven IT leaders are looking for SDWAN solutions for capabilities and not necessarily for costs. While all of these will be beneficial application visibility, management simplicity, and security they all need to be measured and put in perspective.

Capability evaluation is a logical area for technical professionals to focus on, given their training and experience which focuses on functional assessment. Engineering education is about solving technical problems, optimizing the performance of systems and elegant solutions. Financial analysis is not well covered in technical courses and therefore, the graduates are not well equipped to analyze investment economics even though they are technically competent. This imbalance is exacerbated by career development, where technical accomplishments are rewarded but financial performance is not necessarily.

The SDWAN platforms provide real technical features that are attractive to network professionals. Distributed locations are easily configured with centralized management. Application-aware routing optimizes traffic that was previously not possible with traditional equipment. Integrated security features combine previously disparate security measures. Real-time visibility gives operational insights for quicker trouble shooting and proactive management. These are useful features that techies can truly appreciate.

But capability value cannot be assumed to be worth investing in, it must be quantified. The value of simplified management depends on the management burden imposed by the current management, in terms of staffing requirements, delays in change or operational errors that result in measurable costs. Unless there are tangible business impacts due to existing blindness, application visibility is not valuable. The only justification for investing in security integration is to have measurable risk exposure or operational overhead due to the current fragmentation. Improvements in capability without specific quantification are assertions and not justified benefits.

If IT proposals are made with a focus on the benefits of the technology, insist on a cost-benefit analysis. Finance departments must demand a quantification for every claimed benefit. What is the staff hour savings of simplified management at what fully loaded labour costs? How many business impacts are prevented by improved visibility and how often and how much does it cost? What security incidents are prevented by integration at what loss reduction. These questions help to bring technical enthusiasm into monetary terms for proper evaluation. Any idea that cannot pass this test should be thrown out, no matter how **cool** it is.

13.3 Peer Pressure and Market Trends

There are many times when IT leaders make technology decisions based on the trend of adoption in the industry. Statements such as everyone is moving to SDWAN replace for a thorough analysis. Market adoption does not mean that it is necessarily financially appropriate for certain organizations.

There are a number of reasons why industry trend arguments are useful to IT leaders. Trend citations give the impression of outside validation that many bright minds have come to similar conclusions. It is safer to follow the herd than to make a decision on our own which could be wrong. Trend narratives are constantly being reinforced with professional networks and industry publications, which give the perception that going against the trend is being contrarian. Adoption momentum is highlighted in conference presentations, vendor marketing and analyst reports, and a world in which non-adoption becomes the exception.

But general market trends don't give any indication of fit for a particular organisation. If SDWAN is successful for an organization, there is a chance they may possess some traits that make it easier for them to realize the value of SDWAN that your organization does not. The level of high cloud dependency, number of locations, existing contracts and other enablers vary significantly from organization to organization. Market adoption



statistics do not distinguish circumstances but are a sum of all these. If you're an organization with five sites and little to no cloud adoption, SDWAN may not be for you, but if you're a company with fifty locations and a lot of cloud dependency, SDWAN is highly beneficial.

Trend arguments also fail to take into account timing issues that impact the decision of the individual organization. Early adopters will have greater risks and costs of implementation compared to later adopters who will be able to take advantage of the market maturation. The economics for organizations with recently signed MPLS contracts are different from organizations that are nearing MPLS contract renewal. When it comes to investing now, percentages of cloud dependency are more important than industry trajectory. Market trends are indications of the overall direction of the market and not necessarily of the best timing for an individual investor.

Critically analyse demand in the context of your organization's situation rather than take industry trend arguments on faith. If the IT leaders are quoting numbers like adoption rates either peer organization implementations, steer the conversation back to your situation. How much traffic do you have going to cloud destinations? How many sites would implement SDWAN with the proposed architecture. What terms and termination clauses are in your current MPLS contracts. These are organization specific factors which influence the financial results independent of other organisations. Trend arguments should lead to further reflection, not to take the place of analysis.

13.4 Resume-Driven Development

Some IT leaders are interested in new technologies to make themselves more marketable, not necessarily for their employers. SDWAN implementation is a great addition to the resume, irrespective of the financial results.

Labor markets for technology professionals value proven experience with technology. SDWAN skills are becoming more common in job postings as desired skills. Professional compensation surveys illustrate the premium for professionals that have backgrounds in implementing emerging technology. Portfolios of experiences with progressively more advanced technologies are often necessary for career promotion, showing that someone is up to date and has increased their capabilities. Such market conditions induce personal incentives that may be inconsistent with those of the employer.

The themes of network transformation, cloud enablement and security integration are all very popular in today's markets, and SDWAN adds all of these elements to a resume, making it especially appealing. The leading enterprise-wide SDWAN deployment showcases strategic thinking, cross-functional leadership and technical sophistication that employers are looking for. The size and scope of the project further adds to the credentials, irrespective of the financial results that seldom make it into resumes or interview conversations.

Few technology implementations have a significant impact on the individual career path. IT leaders who then leave after putting in place SDWAN do not face any repercussions if the financial performance is not good. New employers are more interested in past implementation skills than in the bottom lines of preceding employers. As mentioned elsewhere, IT leaders who stay on with organizations are seldom held accountable for the accuracy of their projections. This lack of accountability allows for technology selection based on resumes, with no personal repercussions.

This dynamic is hard to tackle head-on but can help account for some of the cases of fervent promotion of technologies that don't align with the organization's requirements. It is not easy for organizations to tell the difference between a person who is genuinely interested in helping the organization and one who is advocating for it for his or her career. Whether or not it's true, if someone is accused of resuming, the

relationship is damaged. But knowledge of this scenario allows for proper skepticism of uninformed, enthusiastic advocacy, which has no real financial basis. Governance structures offer indirect mitigation in the sense that they set up a need for financial justification irrespective of advocate motivation. Projects which are driven by resumes do not get passed because they lack explicit cost-benefit analysis, independent verification and post-implementation accountability. Identifying motivations is not a requirement for strong governance, but rather that poorly explained investments are not made, irrespective of the advocacy dynamics. Rigorous evaluation frameworks are being put in place by finance teams to deal with resume-driven development without hitting it head on.

14. QUESTIONS FINANCIAL TEAMS SHOULD ASK

When assessing SDWAN offers, financial teams need to ask specific questions that will reveal poor analysis and question assumptions that have not been backed up. Cloud Dependency Assessment Inquire about the methodology for calculating cloud dependency percentage. Ask for the real traffic analysis data that is used to do the calculation. Inquire about the figure and if it is current or if there are planned migrations. Challenge the basis of the calculation traffic volume or number of applications.



Fig -13: Questions Financial Teams Should Ask

When it comes to cost calculations, inquire about what kind of costs are considered as baseline network costs. Ask for details of the total costs, including line items. Inquire about labor costs in the operational expense projections. Challenge the way equipment was depreciated for current and proposed states. With respect to savings projections, inquire about what cost reductions are behind the savings projections. Ask for the time frame that MPLS circuits will be eliminated. Discuss the early termination fees of existing contracts. Challenge the bandwidth requirement of proposed circuits.

When it comes to implementation, inquire about the cost of the professional services and if it is included in the projections. Ask for a timeline and milestones for implementation. Inquire about the basis for determining staff training needs. Ask about the contingency in case of implementation delays. When it comes to risk factors, inquire about what can happen if cloud migration doesn't go as expected. Ask for sensitivity analysis in various scenarios. Inquire about the requirements for the application and if they will need dedicated connectivity. Ask about guarantees for performance on internet based transport.

15. ESTABLISHING APPROVAL THRESHOLDS

Organizations need to have clear guidelines and metrics to approve SDWAN investment, as opposed to subjective assessment. In the absence of systematic evaluation processes, technology investment decisions are frequently made on the basis of vendor presentations or internal politics or pressure from competition, not on financial analysis. Formal approval limits provide accountability and consistent evaluation criteria, and prevent organizations from spending money that is not well justified. The following criteria give a detailed outline of SDWAN investment governance.




Fig -14: Establishing Approval Thresholds For SD-WAN Investment

15.1 Cloud Dependency Minimum

Organisations should set a minimum cloud dependency threshold prior to approving SDWAN investment, based on research patterns. A sensible criterion could be 50 percent or more of traffic volume is dependent on cloud, and based on an independent analysis, not the IT department's say so. Cloud dependency is the most important factor when it comes to SDWAN financial viability. SDWAN can route traffic directly from organizations to cloud providers via the internet, rather than using expensive MPLS backhaul via centralized



data centers, which is beneficial for organizations that have a large number of cloud workloads. But for organizations with primarily on-premises applications, there's not much benefit to this architecture. The 50% mark represents when SDWAN platform costs, management overhead and transition costs are typically balanced by the savings on direct cloud connectivity.

The methodology of verification is of great importance. IT departments can measure cloud dependency in a variety of ways that overstate the amount of cloud usage, such as measuring backup traffic, planned migrations, or measuring connections instead of bandwidth consumption. Finance should demand traffic analysis from other network monitoring tools or third-party assessment. Real traffic logs from representative periods are the best source of information as opposed to theory or vendor estimates. For organizations that are not at the 50% point, it is best to wait until the cloud migration is at a point where it can support positive returns to invest in SDWAN. Early investment based on migration forecasts that often fall short of the mark, means financial exposure with no financial benefit.

15.2 Return Requirements

The SDWAN investments should be consistent with the return on capital allocation that is expected across the investments. Use the same 15 percent IRR for capital projects as your organization, when evaluating SDWAN proposals. Technology investments tend to be exempt from the usual financial constraints on the grounds that they are a strategic need, a competitive imperative, or simply a matter of technical obsolescence. These exemptions weaken the capital allocation discipline and allow for poorly justified spending. The SDWAN is an investment in infrastructure similar to facilities or equipment or other capital investments and should be treated in the same manner.

When calculating internal rate of return for SDWAN, take into account all costs such as platform licensing, implementation services, in-house labor, training, parallel time and management overheads. Revenue and savings estimates should include realistic timeframes and not be based on assumptions of immediate benefit realization. Organisations should use the same discount rates based on their cost of capital and not those offered by vendors based on favourable assumptions. If IT departments or vendors say that the returns are not at the typical level, but still make sense because of strategic benefits, insist that those benefits be quantified. If the agility, security posture or competitive positioning is promised, then there must be a financial impact. Strategic benefits should not be used as a justification for not achieving the return requirements if they cannot be quantified. This field guarantees that technology investments are fairly competitive with other uses of organizational capital.

15.3 Payback Period Limits

The payback period of investments in network infrastructure should be reasonable. The maximum acceptable payback period could be 30 months. Longer payback periods should be subjected to further investigation. The payback period analysis is useful in conjunction with return calculations, because it demonstrates the timing of cash flows and the risk of investment. The longer payback periods will expose more to the variables of technology change, organizational changes, vendor viability issues, and business needs. The thirty month threshold is designed to allow for a reasonable timeframe for implementation, but not too much uncertainty exposure.

Conservative assumptions should be used in benefit realization timing for payback calculations. The typical SDWAN deployment takes 6-12 months to complete at distributed locations, and optimization and issue resolution can take several months after deployment. Often benefits come slowly and steadily as sites change and users begin to adjust to new architectures. Proposals which assume immediate full benefit



realization on the signing of the contract do not accurately reflect realistic cash flow patterns. When assessing payback periods, organizations should also take into account the alignment of contracts. SDWAN agreements are usually long term (3-5 years) with high early termination fees. Thirty-six month investments with a thirty-six month contract have little margin for error in projections or delays in implementation. Favorable payback period should be well inside the contract period, allowing for some cushion against variances in execution.

15.5 Sensitivity Testing Requirements

Approval should be based on positive returns under pessimistic scenarios, not just in base case scenarios. With moderate assumption changes, if returns become negative, the investment is deemed to be too risky. Sensitivity analysis shows the impact on investment results of the reliability of projections. SDWAN business cases are usually based on assumptions about bandwidth growth rates, the time to move to the cloud, labor efficiency savings, and circuit cost trends. There are uncertainties associated with each of the assumptions which are cumulative over multi-year projections. Evaluating returns under different assumptions in the model helps to determine investments with an unacceptable risk profile.

Effective sensitivity testing looks at individual and combined key variables. Organizations should enforce analysis demonstrating results when bandwidth growth is 20% below expectations, cloud migration is delayed by 12 months, costs of implementation are 25% higher than expected, and circuit cost reduction does not occur as expected. Combined pessimistic scenarios (multiple assumptions that don't turn out as expected) show worst case exposures. Investment should be approved if it has positive returns in reasonable pessimistic scenarios, not just in base case scenarios. Business cases that are fragile with modest changes to assumptions, and that look appealing when they assume a base case return, are those with proposals that go negative on modest assumption changes. Even if base case is attractive, this type of investment is not acceptable as it is risky. On the other hand, schemes that have a positive return even in the worst case scenario provide a strong financial argument for confident approval. Sensitivity analysis should be done by the finance teams and not taken from the vendor or the IT. Internal advocates choose the sensitivity ranges that will ensure positive conclusions. The variance ranges of similar projects can be used to perform independent analysis to get a more reliable risk assessment.

16. IDEAL TIMING FOR SDWAN IMPLEMENTATION

The right time to invest in SDWAN is when there are quantifiable conditions in the organization, not when the vendor is promoting it or the IT department wants it. Timing decisions have a huge influence on financial results, success of implementation, and value realization. When SDWAN is implemented without the proper prerequisites, it can result in less than expected results, operational issues, and longer payback periods. In contrast, those organisations that time the implementation to the favourable conditions maximise benefits and minimize exposure to risk. The framework below offers structured guidance to determine SDWAN implementation timing.

16.1 Prerequisite Conditions

The time of SDWAN implementation should be correlated with the achievement of certain milestones, such as becoming cloud dependent beyond a certain threshold, which can be determined by current traffic analysis. MPLS contracts that are expiring or have good termination clauses. Internet transport characteristics were assessed and tolerance was confirmed. Implementation resources that are available and not in conflict with other priorities. Created cloud migration roadmap with firm dates for on-premises applications.

IDEAL TIMING FOR SD-WAN IMPLEMENTATION

Invest at the right time. Use real data. Meet the right conditions. Maximize value. Minimize risk.




Fig -15: Ideal Timing For SD-WAN Implementation

Each precondition is focused on a set of risk factors that can prevent SDWAN success if they are not met. Cloud dependency verification makes sure the basic value proposition holds true for your organization's traffic patterns. Without a high concentration of cloud workloads, organizations cannot realize the direct-to-cloud routing benefits to justify the cost of the platform if implementing SDWAN. When it comes to traffic analysis, this should be based on what is happening today, not what it is expected to be in the future, and the cost of SDWAN starts at deployment.

The financial impact is directly related to MPLS contract status. With multi-year MPLS contracts with large termination fees, organisations cannot remove circuit costs, with or without SDWAN deployment. If parallel networks are run during long contract run-off periods, it will double the costs without saving anything. SDWAN is ideally placed to be deployed at the time when MPLS contracts are up for renewal or where there are acceptable termination clauses to minimize financial risk. Finance departments should do a deep-dive into the existing telecom contracts to understand the termination clauses, notice periods and penalty amounts.

Application tolerance assessment avoids disruption in the application and helps in maintaining the business performance and confidence of the stakeholders. Not every application will work well on internet transport, especially if it has a requirement for low latency, low jitter, or a guaranteed bandwidth. Even with SDWAN, some applications, like real-time voice and video communications, industrial control systems, financial trading platforms, and some database replication scenarios, may still need dedicated circuits. There is a need for organizations to perform a thorough assessment of the application requirements prior to adopting architectures for universal internet transport suitability.

Quality of implementation and timeline are influenced by the availability of resources. Deployments of SDWAN are not something that can be done without much network engineering, security, operations and project management involvement. Organizations that try to implement it at the same time other critical resources



are working on other priorities, have delays, configuration mistakes and insufficient testing. Technical debt from rushed deployments to arbitrary deadlines, which must be remediated at great cost. Timing should provide for the commitment of resources during implementation and stabilization periods.

With a roadmap for cloud migration, there's confidence in the future state assumptions for financial projections. Before organizations start to count remaining on-premises application migrations as SDWAN benefits, they should have set timelines, budgets and executive sponsorship for the migrations. Tacit plans to eventually migrate applications are not a good basis for investing in infrastructure. Roadmaps should include specific applications, target dates for migration, teams that will be responsible for the migration, and success metrics that will allow accountability for the migration.

16.2 Timing Red Flags

Some indicators that SDWAN may be too early are ongoing data center consolidation or construction projects, which indicate continued commitment to on-premises data centers. MPLS contract renewals with longer terms and termination penalties in recent times. Acquisitions that are scheduled to add sites that have unknown network needs. Major application upgrades either or both replacements which could impact traffic patterns. Leadership change in IT, which leaves gaps in accountability for results of implementation.

The data center investment narrative doesn't align with the cloud migration narrative that is used to justify SDWAN. The fact that organizations are building new data center facilities, expanding current capacity or consolidating workloads into enterprise owned data centers is proof of their continued investment in on-premises computing. The investments indicate that the dependency on cloud will not be significant even with the migration plans that are being announced. SDWAN proposals coming in at the same time as data center investments show mixed tech strategy, which needs to be sorted out before any of the investments are made.

MPLS contract renewals in recent times show that there is not a good coordination between network strategy and telecom procurement. Organizations that are in the process of evaluating SDWAN and are also looking to renew their MPLS contracts for longer periods of time are unnecessarily exposing themselves to financial risk. This renewal decision will be a cost that SDWAN cannot remove, which will limit savings and extend payback periods. This is typically because the procurement teams are working in silos and renewing contracts without taking into account the network architecture planners. If finance groups are seeing this trend, they should hold off on approving SDWAN until contract terms improve or explicitly account for costs of parallel circuits during the remaining MPLS terms.

Uncertainty resulting from pending acquisitions detracts from investment planning confidence. Acquired organizations introduce new unknown network requirements, such as incompatible technologies, legacy contract requirements, different application portfolios and cloud adoption rates. Pre-acquisition SDWAN designs might need extensive changes to address the new locations and requirements. Investing in SDWAN after the acquisition integration planning process offers a better understanding of requirements and more accurate financial projections.

When there are significant changes in applications, traffic flow changes in unpredictable ways. Bandwidth usage, latency requirements, geographic traffic patterns, and other factors change when you're replacing enterprise resource planning systems, migrating customer relationship management systems, moving to a new collaboration platform, and more. SDWAN designs that are optimized for today's application profiles can be suboptimal following major application changes. Significant application initiatives should be undertaken prior to SDWAN designs or application roadmaps should be explicitly included in architecture planning.



When there is a change in IT leadership, it can leave accountability gaps that allow for investments that are not well justified and oversight of implementation that is not sufficient. New leaders might be approving investments to show innovation but don't know what is going on in the organization. Leaving leaders might push through projects before they go without being held responsible for the results. Stable leadership with clear accountability for the successful implementation and continuous operation of significant infrastructure investments, provides better governance.

16.3 Sequential Approach

Organizations that fall below cloud dependency thresholds should take into account phased approaches. Start at places that are already above cloud dependency thresholds. Show value at a few sites prior to widespread use. Continue to expand as cloud migration continues and dependency percentage grows at additional locations.

This is done to minimize risk, learn about the process and match the time of investment with the time of value realization.

Phased implementation has several benefits over the comprehensive implementation approach. Early deployment at high cloud dependency sites allows for value demonstration before making a full investment. Organizations can compare the actual savings to the projections, understand the challenges of implementation, and make changes to processes prior to scaling. Experience from pilot sites contributes to the quality of execution of subsequent deployments and mitigates the risk of implementation.

The sites for the initial phases should be chosen based on high cloud dependency and reasonable implementation complexity. Ideal pilot locations are those with simple network needs, local staff that are helpful, and representative applications. Choosing sites that do not have special needs, difficult physical conditions or complicated legacy sites for the initial deployment minimizes variables that can impact pilot success assessment.

Pilot learnings and the latest cloud migration status need to be included in expansion criteria. Organizations should have defined metrics that lead to expansion decisions such as proven savings percentages, operational stability metrics, and user satisfaction measures. Expansion should be limited to when pilot locations are successful, and other locations are successful with the prerequisite conditions. This disciplined approach helps to not expand momentum before assumptions are confirmed.

Financial structuring should consider staged solutions modular contracts with the option to expand in stages without a penalty. Vendors that want to sign a long-term contract in advance of proving their worth may be more concerned about their revenue stability than about customer risk management. Contract arrangements should be flexible, with the opportunity to expand, pause or terminate contracts based on actual results, not expected ones.

17. CONTRACT NEGOTIATION CONSIDERATIONS

Contract negotiation plays a critical role in the financial results when SDWAN implementation takes place. SDWAN can mean hundreds of thousands of dollars saved or lost over contract terms for mid-sized organizations and millions of dollars saved or lost for large businesses when the SDWAN agreements are negotiated. When it comes to negotiating, both the vendor and the customer have teams that are experienced in working to maximize their financial position; however, the customer may not have the same level of experience in telecommunications contracting. This imbalance can allow for less favorable terms

which can reduce the expected savings and impose operating restrictions that reduce future flexibility. As part of the contract negotiation process, finance teams should be actively involved with their IT and procurement colleagues, with financial interests adequately represented.



Fig –16: SD-WAN Contract Negotiation Considerations

17.1 Vendor Selection Criteria

Consider TCO, not just up-front costs. Take into account implementation services quality, with reference checks. Evaluate management platform features that minimize management burdens. Assess the financial strength and market standing of vendors.

The TCO is not just the platform licensing fees that are listed in the initial proposals. Implementation services fees such as design, configuration, testing, migration support and training should be included. Recurring operational costs include management platform licensing, support contract fees, software maintenance and periodic hardware refresh. Professional services are often needed to make configuration changes, troubleshooting and feature enablement that are not covered by base support agreements, and these services can add up. A detailed cost breakdown should be requested over the entire expected life cycle of the organization, not just an annual summary of the costs.

The quality of implementation services is highly variable from vendor to vendor and partner ecosystem to partner ecosystem. Technical debt is the result of poor implementation and costs a lot to fix, lengthens deployment times and delays benefits. The question for reference checks should be directed towards the experience with implementation and not the satisfaction with the product. Ask for contacts at organizations that share the same level of size, scope and deployment dates. Inquire with references regarding timeliness, quality of resources, responsiveness to problem solving and effectiveness of knowledge transfer. Some



references may indicate that there were implementation issues, but the project was still successful, which is a good indicator of vendor implementation capabilities.

The features of the management platform directly affect the cost of operations. The sophisticated orchestration platforms that facilitate centralized management of policies, automated provisioning, and simplified troubleshooting, cut down on the amount of labor required to operate the system on a day-to-day basis. In contrast, platforms that need to be manually configured per site, do not offer a full picture of monitoring visibility, or offer poor diagnostics tools add to the operational burden and staffing costs. Assess platforms based on real-life operational scenarios and not feature checklists, preferably by conducting proof-of-concept deployments in scenarios that mimic real-life operational scenarios.

Long-term investment protection is influenced by the financial stability and market position of the vendor. SDWAN is still consolidating, with acquisitions threatening to impact product roadmaps, support quality and price stability. Less financially stable vendors may cut back on development spending, find it hard to keep engineering talent or leave the business altogether. Financial statements, market share, acquisition history and competitive positioning should be considered when evaluating vendors. Choosing vendors that have proven stability minimizes the risk of platform obsolescence and guarantees that vendors will be available for support during contract and beyond.

17.2 Contract Term Optimization

Unit economics are generally better with longer contract terms, but flexibility is decreased. Optimize costs while considering potential future needs. Look into 1+1+1 contracts with renewal options at pre-negotiated rates annually. The choice of contract term is a basic compromise between cost and flexibility. There are significant discounts for longer contracts, sometimes as much as 20-30% for five years contracts over annual contracts. The discounts are due to vendors' preference for steady revenue streams and the lower cost to acquire customers when they renew. But long-term contracts force companies to stay with a certain platform in a fast-changing technology landscape and in a fast-changing business environment.

Short-term commitments are warranted by technology uncertainty. SDWAN capabilities are constantly evolving and new features, performance and security integration capabilities are constantly coming out. Today's platforms may be replaced before longer term contracts are up for renewal. Likewise, organizations evolve via acquisitions, divestitures, geographic growth, application portfolio changes, and workforce distribution changes. Present needs can be met with contracts that are not adequate for future needs. These considerations indicate a need for conservative initial commitments to allow flexibility in adaptation.

The one plus one plus one architecture offers a good compromise between cost optimisation and flexibility maintenance. This is a one year term with two renewal years at pre-negotiated rates. Organisations secure best rates reflecting vendors' confidence in the likelihood of renewal, and have annual decision points to allow for course corrections if circumstances change. When negotiating renewal rates, it is important to do so when the initial contract is being negotiated and competitive forces give the greatest leverage, rather than at renewal time when the switching costs give the least leverage.

Volume commitment flexibility should be considered in addition to term considerations. If you are an organization that is unsure about their growth pattern, then you want to negotiate a contract that will provide for additional sites at the same rate without minimum volume commitments. On the other hand, organizations with a strong deployment confidence can get volume discounts based on the committed scale. Both expansion and contraction provisions should be included in contracts to allow growth at favorable terms and removal of the site without penalty when a business change calls for a reduction of footprint.



17.3 Performance Guarantees

Discuss service level agreements with financial penalties for poor service. Make sure SLAs include availability, latency, packet loss and support responsiveness. Determine the value of SLA credits against the impact on business in case of failures.

SLAs establish a framework of responsibility for vendor performance and financial penalties for non-performance. However, SLA value is solely dependent on particular terms, measurement methods and credit structures. Badly written SLAs don't mean much in and of themselves, but good SLAs can establish a set of meaningful performance incentives and fair remuneration for poor performance.

Complete SLAs should cover several aspects of performance that are indicative of the true needs for service quality. Availability guarantees are minimum uptime commitments, usually in the form of a percentage, e.g. ninety-nine point nine percent availability, per month. Latency commitments define the maximum amount of delay that is acceptable for traffic passing through the SDWAN infrastructure, which is important for real-time applications that are sensitive to transmission delay. The Packet Loss Thresholds are the percentages of the amount of data lost that impact application performance and user experience. Support responsiveness ensures that reported problems have maximum response times and resolution times according to the severity classification. For each dimension, there are specific numerical commitments, clear measurement methods and credit structures for non-compliance.

The details of the methodology of measurement have a great influence on the practicality of SLA. Scheduled maintenance periods, force majeure events, outages due to actions of the customer and other situations that may affect the actual protection may be excluded from availability calculations. Latency measurements can be between vendor infrastructure points, and do not necessarily show end-to-end user experience. Care should be taken in reading the definitions of measurement, the exclusion provisions and the reporting mechanisms. Ask vendors for a history of SLA performance, showing their ability to meet the SLA commitments. Committed vendors are more likely to share history of their successful performance and reluctance to do so will indicate that they have not performed well.

Credit structures determine financial value of SLA protections. Credits for outages are generally small percentages of monthly fees, and often many of these agreements have little financial incentive for the vendor's performance. Work out credit amounts based on actual business impact due to failures. An SLA credit of \$500 is nothing if an outage for four hours is worth \$50,000 in lost productivity and customer impact for your organization. Discuss credit arrangements in line with business impact, e.g. direct damage provisions for repeated or extended failures. Know the credit restrictions, claims and expiry conditions for practical value realization.

17.4 Exit Provisions

Before signing, negotiate reasonable termination provisions. Know ownership implications of equipment at end of contract. If a change of vendor is required, make sure that the data is portable.

Exit clauses are important for organisations to be able to change direction if necessary, but they are often overlooked when initial negotiations are around deployment planning. Lock-in effects are caused by unfavorable exit provisions that remain long after the initial contract expires, and thus reduce competitive leverage at the time of renewal and strategic options. When exit concerns are taken into account during initial contracting, when bargaining power is greatest, the results are much more favorable than if it is attempted when switching costs have been incurred.



The termination provisions should cover convenience termination situations as well as cause termination situations. Convenience termination clauses allow for the termination of a contract if the organizational needs change, for example if there is an acquisition, divestment or strategic redirection, without the vendor being at fault. These provisions usually have notice periods and can include early termination fees that take the form of percentages that decrease as the contract nears its end. Discuss fee structures that are reasonable to the vendor and based on the cost of the vendor, not on punitive amounts that are meant to keep the vendor from leaving. Cause-based termination clauses allow for a termination if a vendor does not meet performance obligations, fails to uphold contractual obligations or has material changes in its ability to deliver the service. These provisions should clearly outline triggering conditions, notice requirements, and cure periods, but be free of procedural hurdles which would prevent practical exercise.

Agreements with vendor provided equipment raise issues of equipment ownership. Know if ownership of equipment passes to your organization at the end of the contract, if it is to be returned to the vendor or if it is offered for purchase at specific prices. There is also a switching barrier with vendor retained ownership because hardware removal will impact operations irrespective of software platform decisions. Organisation owned equipment can still be useful with other management platforms or may end up being unusable due to lack of vendor software support. Agree on ownership arrangements, update supply details for equipment that has aged during the length of the contract and remove details on how the equipment will be removed at the end of the contract to prevent any unexpected costs or disruption of operations.

Data portability provisions guarantee that, should a vendor relationship come to an end, the organization still has access to configuration data, historical performance information and operational documentation. SDWAN platforms capture important configuration information that is built up over operational time and represents network architecture decisions, policy definitions and optimization tuning. This information is lost when moving to other platforms, and may need to be rebuilt based on incomplete documentation. Contracts should include details on data formats, the way data is exported, how long it will be held after the contract ends, and vendor assistance with orderly transitions. These provisions are free to the vendors, but give them needed protection if changes are needed.

18. BUILDING INTERNAL VERIFICATION CAPABILITY

Organizations can gain from establishing internal resources to assess network investment opportunities without relying on the advocacy of the IT department. Today, many technology investments are becoming a major capital allocation decision and should be evaluated as carefully as facilities, equipment, and other significant investments. But many organisations simply do not have the internal skills necessary to critically review network investment proposals, and are dependent on the IT department which also sells and reviews network investment proposals. It is a conflict of interest built into the structure which allows for poorly justified investments to move forward without proper consideration. Independent verification capability fills this governance gap, enhances investment quality and establishes a healthy accountability process throughout the technology planning process.

18.1 Financial Team Training

Invest in financial analyst training in network economics basics. Understanding the concepts of bandwidth pricing, circuit types, traffic analysis, etc., does not need to be at a deep technical level to be understood well enough to evaluate investment.

BUILDING INTERNAL VERIFICATION CAPABILITY

Evaluate network investments with facts, not just IT's opinion. Independent review. Better decisions. Greater value.



Fig -17: Building Internal Verification Capability

Financial analysts routinely review complex investments in a variety of areas such as manufacturing equipment, real estate development and specialized machinery, which do not necessarily involve their own operational experience in these areas. Evaluation of network infrastructure investments also needs the same basic knowledge and not necessarily engineering level technical expertise. Training programs should not be based on technical aspects of implementation, but rather on economic concepts, so that meaningful financial analysis can be done without unrealistic expectations on technical competency development.

Key training areas include understanding the cost and capability of different types of circuits, bandwidth pricing structures and volume economics, fundamentals of traffic analysis to evaluate consumption patterns, and typical vendor pricing models and the financial implications. Knowledge of the bandwidth cost as a function of bandwidth capacity, the price-performance tradeoffs between various circuit technologies, and the traffic patterns and their impact on the infrastructure requirements are sufficient bases to evaluate the investment. The analysts don't have to know anything about routing protocols, encryption algorithms or hardware specifications to determine whether financial projections are based on reasonable assumptions.

Training delivery can vary from formal training to structured mentoring. There are a number of institutions providing telecommunications economics training to a non-technical audience, and in a concentrated format. Organizations may also create internal training initiatives that incorporate external resources with the context of the organization. Assigning financial analysts to technically literate colleagues for continued mentorship helps speed up learning and fosters cross-functional relationships that help improve future collaboration. Investment in training pays off in better investment decisions, in the avoidance of losses from investing in inadequately justified projects, and in better organization for continued technology governance.

18.2 Third-Party Advisory Relationships



Develop connections with independent network consultants to give unbiased evaluation of IT proposals. Make sure that there are independence requirements without vendor payments.

No matter how much you invest in training, there is no substitute for the specialized expertise and objective perspective an independent advisor can offer. Network architecture and telecommunications economics consultants face a variety of organizational scenarios, witness many vendor implementations, and carefully follow the industry developments. This wide range of experience allows problems, opportunities and risks to be identified that may not be noticed by internal teams with less exposure. External perspective also has political benefits in that it gives independence from internal advocacy dynamics which can influence objective assessment.

When forming advisory relationships, Independence requirements should be considered. There are many consultants working in this area who are paid by technology vendors via referral programs, implementation partnerships or sponsorship deals. Such financial links result in conflicts of interest that can either be intentional or unintentional and can affect recommendations. All vendor relationships should be disclosed and engagements should be limited to advisors who are not paid for by vendors. In general, the more independent the advisor, with higher direct consulting fees, the more value you will be getting from them than if you are paying less for the advice of an advisor who has a vested interest in the vendor.

The structure of engagement for advisories should align with organizational needs and patterns of investment. An organization that is making an occasional major network investment can hire advisors on a project-by-project basis to evaluate proposals and support in network negotiations. Retainer arrangements offer the opportunity to organizations with ongoing technology investment programs to have access to advisory expertise throughout the planning, evaluation, and implementation cycles. No matter what the engagement structure is, clearly define scope, deliverables and success criteria for value realization from advisory investments.

18.3 Benchmarking Data Access

Get on a network cost comparison service from an industry benchmarking service. Verify state costs and future state projections using benchmarks.

Benchmarking data serves as an objective reference point to determine if cost projections are based on realistic market conditions and/or optimistic assumptions that are not based on industry experience. If there is no benchmark context, then organizations will not be able to tell if the pricing is truly good or whether the pricing is unrealistic and will lead to disappointment during implementation. Benchmark data can be used to validate at many stages of evaluation, including initial business case review, to post-implementation outcome assessment.

There are multiple organizations that release telecommunications benchmarking information on circuit pricing, managed service costs and technology adoption trends by industry and geographic region. Subscription services offer access to up-to-date data that shows the current market conditions, and not historical data. Choose the sources that benchmark for relevance to your industry vertical, geography and size of your organization. Large enterprise benchmarks aren't really applicable to mid-market organizations with different vendor dynamics and pricing.

When making a benchmark application, it is important to interpret the application, not mechanically compare it. There are valid differences depending on contract conditions, quantity, service level and when it is implemented. But when the proposals are for costs that are significantly less than the benchmarks, it is



appropriate to question whether the assumptions are valid. Likewise, existing state evaluations that reveal costs that are well above the benchmark can imply potential cost-saving opportunities that do not require a major change in architecture, which could be a negative factor for incremental SDWAN savings claims. Benchmarks are not standards, but are used as a reasonableness check and indicate projections that require further investigation, but do not necessarily disqualify unusual projections.

18.4 Post-Implementation Review Requirements

Set standards for financial post implementation review of actual results vs approved projections. Establish a formal variance analysis to hold people accountable for accuracy in projections. The accountability loop from investment approval to outcomes is closed with the post-implementation review. If systematic comparisons of projected and realized results do not occur, organizations will not learn from experience, be aware of the weaknesses in the methodology of making the projections, or be able to hold proposers accountable for delivery of commitments. Formal review requirements, prior to approving investments, make it clear that projections have consequences, which promotes realistic estimates and not optimistic advocacy.

Review timing should take account of benefit realisation expectations in the approved business cases. Too early reviews reflect implementation costs with no benefits, and lead to falsely negative reviews. Late reviews enable memory to fade, people to move on, and excuses to build up, making it difficult to be held accountable. The 12-month post-implementation reviews tend to reflect initial steady state operations where the participants are still available and circumstances are still similar. Twenty-four or thirty-six months follow-up reviews re-evaluate the projected benefits to determine if they are long-term or short-term.

Variance analysis should be done systematically to compare actual results with specific results of approved business cases. Record actual savings versus expected savings by category, including actual bandwidth costs, circuit elimination, labor savings, and other benefit components, versus original estimates. Recognize and clarify major deviations, between deviations of execution within the control of the organization and deviations of outcome outside of the control of the organization. This analysis offers learning opportunities and allows for suitable differentiation of accountability.

There needs to be a balance between learning orientation and consequences. Organizations that want to make honest predictions into the future must not take a "whipping" for a reasonable deviation, but must correct for systematic over-projections. Monitoring the accuracy of projections over time by proposer allows to identify individuals or teams that tend to use inaccurate projections in their proposals and adjust their calibration accordingly. If there is a pattern of consistently overestimating results by large amounts, seeking further clarification, or having contingency reserves for future proposals from those sources. The graduated accountability approach enhances the quality of projections without introducing unnecessary conservatism that stifles good investment opportunities.

19. RECOGNIZING WHEN SDWAN IS NOT APPROPRIATE

A few organization profiles show that SDWAN implementation won't provide financial benefit, no matter how well it's implemented. The technology market has been touting SDWAN as a one-size-fits-all solution that is suitable for practically every organization looking to modernize its network. This is a sales-oriented view, and may mislead organizations where there are no real-life scenarios where benefit realization is possible. By accurately judging the fit, you avoid investing in technology that isn't going to pay off. Finance leaders need to be aware of the circumstances that would suggest SDWAN is not the right option, and not be pressured to make investments that do not have a solid financial basis. Here are some typical scenarios that should not be

used as an argument for SDWAN implementation even if the vendor is excited or internal advocates are pushing for it.

RECOGNIZING WHEN SDWAN IS NOT APPROPRIATE

Not every organization will gain financial benefits from SDWAN. Evaluate the fit based on facts, not hype.




Fig -18: Recognizing When SD-WAN is Not Appropriate

19.1 Heavy On-Premises Application Portfolio

If an organization has no intention of moving to the cloud and has little cloud adoption, then SDWAN shouldn't be pursued for cost savings. While there may be benefits of the technology, there will be no financial gains without basic changes to the traffic pattern.

The value of SDWAN is largely in the ability to route traffic to the cloud directly over the internet and not backhaul it through a centralized data center using costly MPLS circuits. Organizations with a majority of on-premises applications do not have traffic patterns that take advantage of this architecture. Whether you're using traditional routers or SDWAN platforms to manage traffic, the same connectivity is needed between branch locations and enterprise data centers. When the cloud dependency is not that high, the SDWAN platform provides no savings and is just another cost.

Technology enthusiasts can make the case that SDWAN offers more than just cost savings, such as better visibility, easier management and enhanced security integration. These capabilities are present but they aren't often worth the price of the platform for organizations that don't have cloud traffic to optimize. There is a lower cost of visibility and control with traditional network management tools. Security enhancements that can be made with SDWAN usually come with extra security service subscriptions that come with extra costs. These non-financial benefits should be critically assessed and given a monetary value, and not taken on faith that they will improve an organization's operation.

Cloud migration plans are not a replacement for the cloud reality. There are a lot of organizations that have ongoing migration plans that fail to meet their projected migration dates. Unexpected dependencies,



compliance issues or performance demands keep applications that are considered cloud-ready from being moved to the cloud. SDWAN investments based on future cloud migration have benefits that may not materialize, while costs are incurred right away. Any organization must prove their dependency on the cloud by analysing their traffic today and not by the migration plan with the uncertain delivery date before approving SDWAN investment.

19.2 Small Location Count

For organizations that have less than 10 locations, SDWAN cost savings might not justify the investment in implementation. Limited locations of SDWAN platforms lead to high per location costs.

Fixed costs of SDWAN platforms include management platform licensing, implementation services, training and administrative overhead. These costs are mostly fixed and do not change with scale of deployment, leading to sub-optimal unit economics for smaller deployments. An organization deploying SDWAN at five locations will incur the same platform cost as an organization deploying across fifty locations but only spread the cost across 1/10 the base of locations. At a small scale, per-location costs are far too high and often outweigh any possible circuit savings.

This sometimes is hidden in vendor pricing models, such as bundled pricing or minimum contract terms. True unit economics are obtained by careful analysis of total costs over realistic deployment scenarios. Ask for the cost of your actual number of locations, not the cost of a location with the assumption of larger numbers. Determine per-location costs that include proportional platform fees, implementation services and ongoing management overhead. Calculate the savings per site and compare these to these comprehensive costs to determine real financial viability.

Small organizations also have a higher burden of implementation, relative to their capacity to operate. No matter how many locations you have, SDWAN deployment demands a lot of network engineering, security, and operations teams. Small technical teams might find it difficult to support implementation, and continue running the system, which may lead to the need to buy external support, which can further reduce financial returns. Implementation effort is spread over wider teams and with less disruption to the operation in larger organizations.

19.3 Favorable Existing Contracts

Organizations that have just negotiated MPLS contracts at favorable rates may not be able to afford SDWAN savings and the cost of transition and termination penalties. Don't evaluate SDWAN until after the contracts are renewed.

Cost savings from SDWAN are usually based on the cost of the MPLS circuits that will be eliminated or reduced. Whether deployed now or in the future, organizations with multi-year MPLS deals will never be able to achieve these cost savings. Termination penalties, minimum commitment obligations and notice period requirements can result in financial exposure that is far greater than the potential SDWAN benefits that can be realized during the remainder of the contract terms. Having parallel networks for contract run-out periods is not a savings measure, but a cost of double connectivity.

The timing of contracts should be evaluated first before any serious SDWAN discussion. Examine current telecommunications contracts to see what terms and termination clauses are in place and what minimum commitments exist. Determine total exposure due to early termination for all circuits and locations. Consider this exposure vs expected SDWAN benefits over the same period. In many situations, a natural contract expiration will yield much more rewarding monetary results than a forced transition.



There are various scenarios when organizations are facing MPLS contract renewal. Natural evaluation points will occur with upcoming renewals, which is when SDWAN alternatives should be seriously considered. If you are a vendor who needs to renew, you can present better terms, which can alter the financial comparisons. Organizations have the ability to negotiate month-to-month extensions to maintain flexibility while evaluating SDWAN, and not enter into longer-term renewals that tie them up. Evaluation of SDWAN at the time of contract cycles gives maximum strategic flexibility and financial optimization opportunity.

19.4 Unique Application Requirements

Companies that have applications that need specific performance guarantees that the internet transport can't offer should carefully consider which applications can benefit from SDWAN and which need dedicated transport, whatever it costs.

While SDWAN platforms offer sophisticated application capabilities, some applications demand network performance attributes that Internet transport can't always deliver. Applications such as real-time industrial control, high-frequency financial trading, medical device communications and others require low latency, minimal jitter and guaranteed bandwidth availability. Even with SDWAN optimization, internet transport can be subject to congestion, routing changes, and provider problems that can cause the transport to drop below an acceptable level of performance at times. Organizations that operate such applications need to have their own connectivity with locations that run these applications.

By accurately evaluating application needs, deployment failures that hurt business operations and undermine stakeholder confidence can be avoided. IT departments that are keen to show off the success of SDWAN may downplay the issue of application sensitivity or overestimate what SDWAN can do to meet performance needs. Finance should ask for detailed documentation of application performance requirements and SDWAN platform capabilities to meet those requirements. Before going live, proof-of-concept tests in realistic conditions including simulated degradation scenarios can be performed to see if applications can withstand internet transport characteristics.

Hybrid options can be used by organizations that have a combination of applications. Locations that are only using cloud applications and general business productivity applications can successfully migrate to SDWAN with internet transport. Performance sensitive applications may need to remain dedicated connected in locations where they are deployed. This evaluation, which is done on a location-by-location basis, results in more realistic financial projections than assumptions made for all different operational environments. But this means that many sites would have to remain connected, which can greatly diminish the overall savings potential and jeopardize the viability of the business case, despite some sites being able to transition.

20. THE LEADERSHIP COMPETENCY GAP HOW WEAK IT LEADERSHIP CREATES FINANCIAL VULNERABILITY

The errors and the manipulation weaknesses discussed in this analysis have a common denominator a lack of IT leadership competence both in terms of financial and technical expertise. If IT Heads or CIOs aren't skilled in either area, then they leave their enterprises vulnerable to vendor manipulation, bad investment decisions and significant financial losses. By learning about these leadership gaps, boards and executive teams can focus on root causes and not symptoms, and create organizational capability to avoid repeating the same mistakes on all technology investments, not just on individual procurement mistakes.



Fig -19: Recognizing When SD-WAN is Not Appropriate

20.1 The Dual Competency Requirement for Modern IT Leadership

Over the past few years, it has been increasingly recognized that a knowledge of only one aspect of the business IT is not enough to achieve success as a CIO. CIOs must have both types of knowledge to manage IT resources as well as to manage and plan such as policy and practice development, planning, budgeting, resourcing and training. The combination of these two competencies is very challenging as most IT leaders have built their career focusing on one competency and ignoring the other. Technical professionals who are promoted to engineering management positions don't have the financial acumen to make informed business decisions, and business-oriented leaders who are recruited to fill IT positions may not have the technical expertise to make informed technology decisions.

The most common pitfall of new CIO candidates is that they want to be too technical and not business-oriented enough. Without that bigger picture it's difficult to make the leap into a CIO position with success. On the other hand, too often, other C-Suite executives believe that they can have a strong business leader with a weak technical background as a CIO, and just leave the technical expertise to other tech experts within the company. A CIO who does not have a solid understanding of technology whether it is IT automation, DevOps or cloud technologies will find it difficult to perform the role.

The SDWAN miscalculation patterns mentioned above are simply an example of this two-fold failure. To make a cloud dependency miscalculation, you need not only technical knowledge about traffic patterns, but also financial knowledge about the cost implications. When vendor manipulation is successful, leaders are either not technically savvy enough to determine if there is a problem or not, or they are not financially sophisticated enough to understand that economic projections are incorrect. Contract term mistakes are a result of poor financial analysis, and capacity planning deception is due to a lack of technical knowledge.



Technical and business only leaders do not have the holistic view to effectively safeguard the interests of the organization.

CIO needs to have business knowledge and technical knowledge. For the past five years or so CEOs and boards were looking to hire CIOs with an equal mix of both. Without this balance, organisations are systematically exposed to the risks of the kind of errors and manipulations that turn technology investments into money-losing ventures. The SDWAN context is simply an example of the general trends that are repeated in all cloud migrations, digital transformations, cybersecurity investments and all major technology investments that involve both technical and financial considerations.

20.2 The Financial Literacy Crisis in IT Leadership

Historically, financial literacy in IT has been low and a series of obstacles have contributed to this. This lack of financial literacy makes IT leaders vulnerable when they are considering vendors proposals, looking at the return on investment, or even trying to explain the value of technology to business stakeholders. A representative sample of U.S. managers from C-level executives to supervisors took a basic financial-literacy test a test that any C-level executive or junior finance person should score an A on and on average, they scored 38%. Most were not able to differentiate between profit and cash. Many were not familiar with the difference between an income statement and a balance sheet.

This alarming base level is applicable to all management functions, and IT leadership has other challenges. The financial literacy is decreasing among all leaders and managers in each department and each industry. This increasing disconnect between knowing how to communicate and understand the business language is a major barrier to IT effectiveness and the career trajectory of anyone who wants to join or stay in the management ranks.

There are direct impacts of financial illiteracy, such as SDWAN evaluation failures and other technology investment errors. IT leaders who are not financially savvy are unable to independently validate vendor savings claims, find out what costs were not mentioned, or see contract term arbitrage that makes savings look better than they actually are. They take case studies for granted without realizing the importance of the baseline costs, traffic profiles, and organizational characteristics to understand the percentage savings. They don't ask for traffic-volume based assessments because they don't know why application counts are not accurate for financial projections.

IT investment is often driven by gut feel, rather than facts and figures. The system doesn't provide enough flexibility to adjust and redirect IT dollars in response to shifting business needs. IT remains a cost center, not a strategic partner and still hasn't been able to effectively articulate the value of its work to its business partners in terms they can relate to. This is an instinct based decision making, which is indicative of inadequate financial framework for systematic evaluation of technology investments. Without being able to express the value of investments in terms that business stakeholders appreciate, IT leaders lose credibility, which is needed to counter the vendor's dubious claims or to defuse pressure to move forward with a proposal that has not been thoroughly analyzed.

Many times, CIOs are only thinking about their budget and may only have a vague idea of what's driving changes in the company's quarterly results. Your FP&A teams will likely be most aware of what is going on 'below the waterline' and this can impact company performance in a positive or negative way. This limited spending perspective means that IT leaders don't get a clear picture of what impact technology investments have on the rest of the organization's financial results. They don't have frameworks for thinking outside of



departmental expenses, so they can't relate SDWAN implementation with enterprise cost structures, working capital requirements or competitive positioning.

Too few discussions are taking place on tech financials. That's because IT management lacks experience and expertise in finance, and the business lacks understanding of what technology is, does and can do. This lack of understanding on each other's side leads to risky gaps in which neither IT nor finance has enough knowledge to make proper decisions about technology investments. It is in these gaps that vendor sales teams see an opportunity to sell technically great but economically poor proposals that neither function can be expected to see an opportunity to challenge.

20.3 Technical Knowledge Degradation in Senior IT Leadership

Financial ignorance is one important gap, but so too is technical knowledge degradation among senior IT leaders, which can also be very dangerous. In many instances CIOs will have risen through the IT ranks, and may have some IT qualifications. However, in most cases they will not have had hands on experience with technology for many years and as technology is so fast moving, this can leave them feeling rather out of touch with their technical skills and experience.

This technical obsolescence can be seen in SDWAN evaluation failures, where IT leaders can't independently assess claims of traffic patterns, validate cloud dependency calculations, or evaluate capacity planning assumptions. They have no current technical knowledge to do their own analysis, and are completely dependent on the vendor's analysis. Sales representatives know this and use it to their advantage in making technical presentations that are not designed to educate, but to dazzle, and that use terms and concepts that the outdated technical knowledge can't be used to judge.

Sometimes, a CIO will be a non-technical business leader who was promoted to CIO because he or she demonstrated an interest in technology and how it relates to the business and didn't have any technical skills at all. The most crucial aspect of a CIO's job is to be able to chart a path for technology that will empower the organization. Strategic vision is important, but if there is no technical capability at all, this means that dependency on subordinates and vendors arises, and they may have conflicting interests or incomplete vision.

In many cases this is because the consultants providing CIO services have been CIOs for many years, and have no longer been involved in the technology. Plus technology changes rapidly, their technical know-how and skills become outdated rather quickly. Technological change is ongoing and faster than it has ever been technical knowledge becomes outdated at a rapid rate. Many current IT leaders learned their technical foundation at a time before cloud architectures, software-defined networking, containerization, and other technologies that underpin many of today's investments were in use or even existed. Their mental models of network operations might be based on architectures of past decades, and not current capabilities and constraints.

Nothing can be more damaging to an executive's credibility than acting as if he or she knows what he or she doesn't know. After a few minutes, most experienced engineers, product leaders and cybersecurity professionals know the difference between curiosity and certainty. The use of jargon that leaders don't understand or the rejection of issues that haven't been discussed makes technical discussions less forthright. Teams do not raise tough issues for fear that the leadership has already decided.

This lack of credibility adds to the evaluation issues. Technical staff who know that leadership has a limited technical understanding may begin to avoid giving them honest opinions on vendors' claims, as they believe



leadership will not be able to understand their nuances technical concerns. Or, staff can take advantage of leadership's lack of technical knowledge to pursue pet projects or to escape examination of their own recommendations. Both dynamics call into question the technical validation that should prevent organisations from making bad investment decisions.

20.4 The Career Path Problem How IT Leaders Develop Blind Spots

To understand the reasons behind the creation of these competency gaps in IT leaders, it is important to look at the typical career paths that lead to systematic blind spots. An organisation may promote a superb IT Manager to a high technology position and the person will start to work as a high technology manager right away. But a new title doesn't give experience that somebody has never had the chance to get.

Perhaps, someone has been managing an IT team for years, and never had to present a technology investment case to senior leaders, or contribute to wider business planning, or own a technology strategy, or make decisions with finance or operations, or manage a big transformation programme, or be held accountable for business outcomes outside of IT. This lack of experience means that IT leaders are not ready to undertake the strategic financial assessment needed to make technology investments.

For many IT professionals, there is a frustrating dilemma in seeking to make the transition to technology leadership. They must have strategic experience in order to land the job, but they need the job to gain strategic experience. An IT Manager might be able to help shape business strategy but has not been called in to the discussions. They are familiar with technology investment but have never had an investment decision. They may have good thoughts on transformation, but they spend a lot of time in operating the current systems.

This experience paradox helps to understand why promoted IT leaders are not fully prepared to negotiate with the vendors and to evaluate the investments. Skills for operational excellence with existing infrastructure don't translate to skills for evaluating transformational investments. There is little learning in financial analysis for multi-year cost projections in successful project delivery. Technical problem solving skills are not the same as skills to spot manipulation techniques that are used to take advantage of gaps in knowledge, not technical limitations.

Technical experts tend to view the organization as a whole in terms of their role. Finance leader is a control freak. A scientific leader safeguards rigor. A deal professional is one who is focused on speed and execution. IT leaders who come up through the ranks from technical positions tend to keep this functional focus, and may consider proposals only from a technical perspective, failing to take into account financial, organizational change, or strategic implications. This working myopia helps vendors to manipulate as sales people write proposals that meet technical evaluation criteria but hide poor economics.

Typical IT manager or Director is not anywhere close to being qualified to do the kind of budgeting we are talking about and that is why you find that CIOs seldom come from the tech stream unless they have been inculcated for years. Being promoted from technical manager to IT leadership position without a conscious effort to develop the organization financially results in a leader that is structurally ill-equipped to assess investments. These leaders may know that they are not as well off as they would like to be, but may not be able to admit that there are areas where they may not be as good as they appear to be.

20.5 Organizational Factors Amplifying Leadership Weaknesses

Individual competency gaps and organizational factors exacerbate vulnerabilities and make them hard to correct. The dynamics between the head of IT and the CFO are not often a partnership, although it has



improved in recent years. This gap in relationship makes it difficult for the not so financial savvy IT leaders to get access to expertise that can fill the gap. They assess vendor proposals, without finance partnership, that would reveal hidden costs, contract term problems, or unrealistic savings estimates.

Many IT departments still report to the CFO and that's not necessarily a sign of a lack of strategic interest in increasing IT's impact through investment, but rather, a tactical lack of trust in IT's cost management capabilities. This reporting model is a symptom of the way IT leadership is seen as weak in the organization, and is a response to the symptoms, rather than to the causes, by assigning an overseer. CFO oversight can identify obvious financial mistakes, but can't deliver the integrated technical-financial perspective that is needed for good technology investments.

IT leaders aren't seen as business leaders. There's a lack of communication and understanding between IT and the business. IT's lack of knowledge of where it has business gaps, exacerbated by its continued marginalisation from business contexts and conversations. This exclusion forms self-reinforcing cycles as IT leaders are unable to learn business acumen because they are not exposed to business context and they remain excluded because they are unable to learn business acumen to make meaningful contributions.

What is intriguing is that IT employees are the most optimistic when it comes to their financial skills while financial literacy is a known gap. The perception and reality can be different. This gap between confidence and competence can be especially problematic as IT leaders may not realize that they are not competent. Leaders who think they have a good grasp of financial analysis don't need to consult others or second guess themselves. This false confidence allows vendor manipulation, as leaders believe their own estimates, and do not feel the need to have an independent verification.

One of the many reasons for building up technical debt and making suboptimal technology investments is technical incompetence at the executive level. Executives who are unable to assess technical claims on their own often rely solely on the advice of specialists who may have limited views or personal interests, or they make decisions based on non-technical considerations, which may be at odds with the technical facts.

20.6 How Competency Gaps Enable Specific SDWAN Miscalculations

These leadership competency gaps are related to the SDWAN evaluation failures mentioned above. These miscalculation patterns indicate certain areas of knowledge that better-prepared IT leaders wouldn't be making.

The miscalculation of cloud dependency based on the number of applications and not traffic volume shows a lack of technical and financial skills. Technically savvy leaders know that applications vary greatly in terms of bandwidth usage and would require traffic analysis. Financially astute leaders understand that cost forecasting should be based on volume measurements, not just on numbers. Leaders without competencies accept flawed assessments methodologically, as they do not know how to identify the flaws.

Hidden cost omission is effective due to the lack of financial intuition of leaders who are not financial. Professional financial analysts easily make comparisons between cost categories for alternatives, making sure that apples are compared to apples. IT leaders who are not trained to do this take vendor calculations at face value, and don't bother to systematically verify them, excluding the costs of professional services, training, licensing and support which have a significant impact on the economics of investment.

Capacity planning deception is based on ignorance of technical knowledge about network engineering principles. Leaders who don't know how utilization measurement methodologies are calculated and what the difference is between peak and average capacity requirements cannot challenge proposals to size circuits



for average instead of peak demands. They don't have the technical expertise to evaluate and rely on vendors' claims of the proper bandwidth tier.

Contract term arbitrage takes advantage of the financial illiteracy in relation to contractual commitments and transition planning. Unskilled leaders might not be aware that existing MPLS contracts may have limitations on when to transition and how much it will cost. They take savings projections that are based on unconstrained optimization, but do not confirm that the projected timelines for implementation are possible under the contract.

Phantom feature attribution is a success when leaders fail to see the difference between what matters and what doesn't and think that what doesn't is impressive. Technical demonstrations make impressionistic claims of value without support, as leaders don't have the technical expertise to evaluate the relevance of features or the financial acumen to insist on quantifiable benefits.

Smart managers know that one of the most costly technology blunders start out as management blunders. Projects proceed with no well-defined goals. When reliability ought to be the top priority, teams optimize for speed. Organizations buy complex platforms without knowing whether or not current processes will help achieve the desired results. All of those decisions are not in the least bit technical. They involve decisions on priorities, governance, accountability and allocation of resources.

20.7 The Vendor Awareness Advantage

Sales people in the networking industry are aware of the competency gaps in IT and create engagement strategies to take advantage of them. Sales training focuses on knowing where the customer's lack of knowledge is and presenting accordingly. Technical demonstrations outwit leaders who have no up to date technical knowledge and financial projections take advantage of poor financial analysis skills.

When the CEO is considering an investment in AI or a cybersecurity program or a major systems overhaul, he or she is not often the most technically knowledgeable person in the room. The issue begins when the executive either turns a blind eye to the experts, or when they feigns understanding of the details when they are not clear. In both instances, the leadership team is spared of the due diligence that would be expected for a significant investment. Vendors know these trends and take advantage of both failure modes. Leaders who give away all control, vendors make sure that the technical people are given proposals that highlight the technical strengths and downplay the financial weaknesses. Leaders feign understanding and vendors give complex-sounding, but false, explanations that leaders don't have the means to judge, but they don't feel like challenging.

When IT leaders do not know as much as vendors, it is a structural advantage to manipulation. Vendors have a wealth of knowledge regarding their solutions, negotiation experience with their customers and knowledge of the competitive landscape. IT leaders tend to assess large investments in networking technologies infrequently, have a general understanding of the technologies, and don't have comparative data from their peers. This imbalance exacerbates the competency gaps and even reasonably competent leaders can be easily manipulated by sophisticated manipulation tactics.

Vendors also take advantage of the dynamics of the organization that magnifies the weaknesses of individuals. Sales people discover that there are people in the company who are interested in technology, even if it isn't beneficial to the company. They know that when IT executives are pushed to prove their innovative nature or cost savings, they are likely to take optimistic forecasts that fit a desired conclusion. They



understand that politics in the organization could stop the finance teams from questioning IT recommendations or IT teams from admitting that they are not sure about claims from vendors.

20.8 Developing Comprehensive IT Leadership Competency

To fill the leadership competency gap, it is necessary to develop both technically and financially. At least some basic understanding of accounting and finance is a must for IT leaders. But more significantly, it is essential to have a tight relationship with the CFO team, discussing costs and other financial matters. It is important to have an understanding of the business financial requirements and operations. You don't have to be a CFO, but you should know how IT's budget fits into the overall company budget. The CIO is at the nexus of technology, business strategy and people leadership it's not just about the tech, it's about helping the company to succeed.

At the same time, IT leaders need to stay technologically up-to-date while being busy with operations and keeping their hands off technology. Those leading CIOs are spending time to become intimate with the technology of AI and bleeding edge tech so they can educate their colleagues on what they need to know about those technologies and what impact they can have on the business. She has been introduced to dozens of CIOs these days, who are taking technical courses. That may be a conscious effort to seek out opportunities to influence outside of IT, understand the commercial impact of technology, own the business outcomes, delegate operational issues, don't solve all the problems yourself, and get involved earlier in business and technology decisions.

Organizations also need to put in place structural safeguards that make up for the competency gaps of individuals. To groom the next generation of technology leaders within their own ranks, organisations must provide opportunities for these individuals to work outside of the IT department. This may be as simple as providing an IT Manager with more visibility into senior decision making, or to include them in investment discussions or to lead projects with an impact on the broader organisation.

20.9 Structural Protections Against Leadership Competency Gaps

But, an individual's competency development is not the sole solution to prevent SDWAN and other investment mistakes from happening. Structural protections are for when there are gaps that development has not yet met in the individual leaders. Cross-functional evaluation teams bring together IT, finance and operations skills to make sure there are no gaps in one competency that allow for poor investment decisions. The finance team members find cost items that are not apparent to IT leaders and contract term problems. Operations representatives evaluate the feasibility of implementation and the risk of disruption to business. IT technical personnel validate traffic analysis techniques and capacity planning assumptions. The sharing of these skills makes up for weaknesses in individuals, and allows the capability of the organization to grow through the sharing of the evaluation experience.

Independent verification requirements help ensure that claims are independently verified prior to making major investment decisions. Cloud dependency calculations, traffic pattern analysis and capacity requirements are verified by third party network assessments. Independent financial review validates total cost of ownership projections/savings calculations. It's a small investment to verify to avoid significant losses due to inaccurate vendor assessments.

Standardized evaluation frameworks ensure that there is a systematic evaluation process irrespective of the expertise of the individual evaluators. Checklists that specifically call for detail on the cost categories, contract terms, capacity assumptions and feature benefit quantification eliminate the possibility of omissions that may be taken advantage of in a vendor-prepared proposal. The Framework's compliance requirements



move the focus of evaluation from a judgment to a systematic process, and away from the reliance on competencies that individual leaders may not have.

Organisational profile-specific requirements on vendor references guarantee that case studies and success claims are based on truly similar situations. Requirements for references from organizations with similar cloud dependency, geographic distribution and application portfolios ensure that there is no manipulation through irrelevant benchmark comparisons. Lack of relevant references indicates possible mismatching of vendor capabilities and organizational requirements.

Post implementation reviews, which compare the results of the investment with the pre-investment projections, create organizational learning to avoid redoing the same mistakes over and over. Evidence base for future vendor claims of projection accuracy, unexpected costs, and/or implementation problems. This institutional memory can make up for the turnover of individuals, and lessons learned from past investments will be used to guide future investments, irrespective of which individuals are doing the evaluations.

20.10 The Board and Executive Team Responsibility

In the end, it's up to the boards and the executive team to make sure that their IT leadership competency is sufficient to meet organizational needs. There must be somebody who can balance those competing priorities. This is executive work and can't be delegated just because the technology is complex.

Knowing about the gaps in competency in the IT leadership area can help in the hiring process, investing in development, and providing structural protections. When recruiting CIOs, Boards should evaluate technical and financial skills great technical skills don't make up for poor financial skills and vice versa. Executive teams should be mindful of the development needs of IT leaders and ensure that they are offered development opportunities to address those needs and not assume that promotion will automatically provide them with the necessary capabilities.

Not seeing technology as part of business strategy is one of the most common pitfalls. There are many instances where organizations get involved in purchasing a new platform, automating or utilizing AI without a clear definition of the problem they are trying to solve or how to measure success. Technical teams can perform well, but the business still doesn't, because leadership never focused on aligning the investment with strategic priorities. The best CEOs prevent that from happening by starting every big technology decision with the business objective, not the technology.

The SDWAN evaluation failures detailed in this analysis are all symptomatic of underpinning leadership competency failures. When procurement failures occur, and they are not addressed in a way that corrects the capability gap, it means that all technology investments remain vulnerable. Companies that want to avoid being taken for a ride by vendors either or both making incorrect investments need to invest in building a broad IT leadership competency, both technically and financially. This capability is the only one that allows for the sound judgment that is necessary to safeguard the interests of the organization in more and more complex technology environments.

21. FUTURE CONSIDERATIONS

There are a number of new trends that could impact SDWAN economics in the future. Technology markets change all the time, as new innovations, consolidation and competitiveness reshapes the market that organizations are navigating when they make infrastructure investments. Today's decisions will be made under conditions that could be very different from what occurs in the future. As new trends emerge that

impact the SDWAN economics, financial leaders must keep an eye on them and change their evaluation frameworks and the timing of investments accordingly. Although it is impossible to predict what the market will look like in the future, knowing the direction of the market will help make better decisions that will consider the likely evolution of the market. The following are some considerations that reflect important trends that should be continually monitored when organizations are assessing and running SDWAN investments.




Fig -20: Future Considerations For SD-WAN Economics

21.1 SASE Convergence

Secure Access Service Edge frameworks are SDWAN frameworks that incorporate security services. This convergence could alter the cost dynamics by combining previously separate costs. Financial evaluation frameworks need to be adjusted for measuring the economics of bundled services.

SASE is an architectural model that integrates network connectivity and security services into all-in-one cloud-based solutions. Traditional methods involved using SDWAN for network optimization and security solutions such as firewalls, secure web gateway and cloud access security brokers separately. These functions can be combined in SASE convergence, which may streamline operations and alter the procurement and cost allocation landscape. As organisations consider SDWAN, they are now forced to look at alternatives that combine security services with connectivity management SASE.

SASE bundling adds a lot of complexity when it comes to financial evaluation. Standalone SDWAN assessment looks at network costs in a relatively easy way, looking at the cost of circuits, the platform and the operational overheads compared to the baseline network connectivity spend. Evaluating SASE, however, demands a holistic view of network and security economics, and a comparison of bundled services with the existing spend on the various services. Organizations need to shift security dollars from various cost centers, integrate



security evaluation into the network and security teams with different priorities, and determine if bundled security capabilities are suitable for what was previously done with security point solutions.

The economics of bundling are very different among organisational profiles. SASE solutions can help organizations with existing complex security tool stacks realize consolidation benefits by eliminating overlapping functionality and making vendor management easier. SASE may not be as financially appealing to organizations that have recently implemented comprehensive security architectures, since they already have to pay for some of the capabilities offered by SASE. However, before assuming that SASE bundling offers favorable economics, it's important to take stock of existing security capabilities, their costs and their effectiveness.

Increasingly, vendor's strategies are focused on SASE positioning, which may be reducing SDWAN availability over time. Organizations that opt for unbundled solutions may also have fewer options as vendors shift portfolios to integrated solutions. It is important to track vendor roadmaps and market direction in order to predict changes in availability that impact future procurement opportunities. Evaluation frameworks may have to be modified, depending on preference, as market evolution limits alternatives.

21.2 5G Transport Options

Cellular transport can offer extra connectivity options which can alter SDWAN economics. Assess the impact of 5G deployment in your sites on new transport configurations, and how it can enhance financial results. 5G cellular networks provide bandwidth and latency that, in covered areas, can be similar to wired broadband networks. This capability allows to take into account cellular transport as primary or secondary connectivity for branch locations, which may even eliminate wireline circuit requirement at some branch locations. Today, SDWAN platforms are more and more able to support cellular integration and dynamically route traffic based on availability, performance, and cost, across both wired and wireless transport options.

The financial aspects of 5G transport options should be taken into account. 5G could be costly relative to terrestrial broadband, so enterprises will be judicious in their use of 5G. This is where SD-WAN can help, by allowing enterprises to define policies to route latency-sensitive applications through 5G, and other applications through business broadband. This smart traffic control enables organisations to maximise cost efficiency and benefit from the 5G performance attributes where they are needed. Geographic factors play a key role on the viability of 5G as primary transport.

The coverage is expanding, albeit most of it is in larger metropolitan areas at this time. Depending on SDWAN platform capabilities, organizations with rural locations, international sites or facilities in areas without good cellular infrastructure cannot use 5G as their primary connectivity. Before adding cellular transport to the financial projections, assess 5G availability at each location. The development of recent days shows the increasing support of vendors to integrate 5G.

Solutions from vendors, such as Versa, have been announced for 5G Extender solutions, which solve persistent problems in branch connectivity, such as signal strength in wiring closets where network appliances are installed. These solutions allow businesses to deploy 5G radio equipment where it will get the best signal by windows, building edges or rooftops and control traffic via existing SDWAN infrastructure. These innovations help to lower the hurdles towards 5G adoption for enterprise connectivity applications.

21.3 Application Modernization



As more and more industries continue to migrate to the cloud, over time, the percentage of cloud dependency will increase. Organizations that fall below current thresholds can attain SDWAN viability via natural modernization of applications without having to migrate.

Even those organizations that are not at viability levels are industry-wide on a modernization path that will lead to SDWAN. Enterprise software vendors are increasingly offering functionality only in the cloud, and they don't have on-premises deployment options for new functionality. As companies continue to manage legacy applications, they are under increasing pressure from both the approaching end of vendor support and the accumulation of security vulnerabilities, as well as being unable to leverage new features offered by modern applications. These pressures result in eventual migration, irrespective of infrastructure-planning preferences.

Natural migration progression allows patient organisations to migrate to SDWAN viability without forcing migrations. In three to five years, a 20% cloud dependent organization could be 50% cloud dependent due to natural application lifecycle management, even if they don't try to accelerate the migration process. By waiting for the natural threshold to be achieved, transition costs, change management burden and operational risks of forced migration are avoided and SDWAN benefits are still realized. When assessing investment timing, financial projections should include realistic migration timelines and not aspirational timelines.

But there are opportunity costs associated with passive waiting, which should be taken into account. The SDWAN market is expected to grow from \$9.5 billion in 2026 to \$44.3 billion in 2033 at a CAGR of 24.7% over the forecast period. Businesses that wait may not be able to benefit from the competitive edge that can come from better application performance, security, and operational flexibility, which early adopters enjoy. Consider optimal timing for achievement of thresholds in the context of strategic value of early capability acquisition.

21.4 Vendor Consolidation

Further consolidation in the market could lead to a decrease in competition, which could impact pricing. Keep track of vendor market position and leverage through multiple vendor evaluation processes. There has been a lot of consolidation in the SDWAN market, with larger technology firms acquiring innovative startups and competitors consolidating to get to scale. Cisco Systems, VMware, Fortinet, and Huawei Technologies are among the key players in the market, which offer a wide range of products designed to meet the different needs of enterprises. This consolidation can lead to fewer vendors and less competitive pressure, which in the past helped to spur feature innovation and lower prices.

In 2024, SDWAN has become commonplace, with 87–90% of enterprises deploying or actively deploying it. This wide adoption is a major paradigm change from growth based customer acquisition to retaining and expanding existing accounts. As the market penetration matures and competitive displacement becomes less frequent, vendors might be tempted to focus on profitability over being the most competitive.

In 2026, 60% of new SDWAN purchases will be as part of a single-vendor SASE offering compared to 15% in 2023. This trend of convergence also further focuses buying power on integrated platform vendors who can provide integrated networking and security solutions. The market is consolidating towards comprehensive platforms, and the choice of best-of-breed solutions may be reduced for organizations looking for them.

As the market consolidates, it takes work to maintain negotiation leverage. There need to be real multi-vendor evaluations, rather than the processes merely rubberstamping decisions that have been made. Showing a desire to choose other vendors puts pressure on the vendor that can't be duplicated in single-source negotiations. Carry out multiple supplier relationships during the contract, with alternative supplier



relationships available if incumbent supplier relationships break down. Look for contract terms that allow for the possibility of switching vendors at reasonable costs and avoid those that require locking you into a vendor that lacks competitive leverage at contract renewal.

22. CONCLUSION

SDWAN technology can be a cost effective solution in certain scenarios, but not all. Cloud dependency is the central financial truth that dictates SDWAN value. SDWAN can deliver significant cost savings on the network for organizations that are highly dependent on the cloud. For organizations that are less cloud dependent, there will likely be an increase in costs, not savings.

Unfortunately, many IT leaders don't understand the relationship, and opt for SDWAN investments based on incomplete analysis or vendor-influenced projections. When sales people are under pressure they have an incentive to sell you a rosy picture, even if it is not the truth. Financial teams will need to build up their ability to independently verify or not take technical recommendations at face value.

The bottom line for financial practitioners is simple. Evaluate SDWAN proposals after performing demand traffic volume based cloud dependency assessment. Use the normal capital allocation criteria such as the required rate of return and payback period. Make sure that the cost modeling is complete and all categories are included. Conduct sensitivity analysis to gain insight into the robustness of projections. And set clear limits of approval based on measurable conditions not on the IT department's advocacy, or vendor promotion cycles.

But, organizations that put a hard financial analysis behind SDWAN decision making will invest in SDWAN solutions that deliver the returns promised. Those that take up the rosy forecasts without due diligence will keep seeing the value drain that makes SDWAN one of the most often disappointing IT investments, across industries. The only distinction between these results is whether the financial teams are willing to pose the tough questions and require tangible responses.

ACKNOWLEDGEMENT

This article came about because of firsthand experience in enterprise network transformation projects, where there were large gaps between expectations and actual results in financial outcomes, and there was a need for deeper analysis. It emerged from several large scale network implementations that there were some common elements to the financial forecasts that were being made at the time of approval but weren't coming in at the operational phase. These findings, along with conversations with finance and tech practitioners facing similar issues, highlighted the need for simple and understandable frameworks for evaluating investments, enabling financial teams to truly engage in technical investment ideas. Thanks to colleagues in Finance and Information Technology for feedback given freely. Their guidance and practical experience combined with my recent work experience on SD-WAN projects and the problems faced in the evaluation of proposals helped drive the practical advice and **how to** attitude of this guide.

REFERENCES

- [1] Administrator. (n.d.). Complete Guide to SD-WAN. Technology Benefits, SD-WAN Security, Management, Mobility, VPNs, Architecture & Comparison with Traditional WANs. SD-WAN Providers Feature Checklist.



- <https://www.firewall.cx/security/sase-and-sd-wan/sd-wan-networks-benefits-management-security-architecture.html>
- [2] Afzalzada, A. (2026, October 1). Data Breach Statistics 2026: Cost, Frequency, Causes and Timelines. Stingrai. <https://www.stingrai.io/blog/data-breach-statistics-2026>
- [3] Breaking down the Real-World benefits of SD-WAN. (n.d.). Palo Alto Networks. <https://www.paloaltonetworks.com/cyberpedia/benefits-of-sd-wan>
- [4] Business Broadband Hub. (2026, September 3). SD-WAN costs: What businesses pay in 2026 for SD-WAN. <https://www.businessbroadbandhub.co.uk/sd-wan-costs/>
- [5] Carter, R. (2025, June 25). Sase sd wan: Top 5 Powerful Benefits & Key Differences 2025. Modern Theme. <https://netsharx.com/sase-sd-wan/>
- [6] Cd, & Cd. (2026, August 24). Extending the conversation on Return-On-Investment (ROI) analysis. Health Care Financing Initiative. <https://hcfi.georgetown.edu/whats-new/extending-the-conversation-on-return-on-investment-roi-analysis/>
- [7] Checking your browser - reCAPTCHA. (n.d.). <https://pmc.ncbi.nlm.nih.gov/articles/PMC12567460/>
- [8] Cindi. (2026, January 19). SD-WAN's Impact on Business Networking in 2026. kinectcommunications.com. <https://kinectcommunications.com/why-sd-wan-is-transforming-business-networking-in-2026/>
- [9] CIOPages. (2026, June 25). Buyer's Guide: Cloud networking & SD-WAN. CIOPages. <https://www.ciopages.com/buyer-guides/cloud-networking-sd-wan>
- [10] Cisco vs Fortinet vs Palo Alto vs HPE Aruba Edgeconnect: An SD-WAN Comparison Guide for 2026. (2026, September 23). <https://technologymatch.com/blog/cisco-vs-fortinet-vs-palo-alto-vs-hpe-aruba-edgeconnect-an-sd-wan-comparison-guide-for-2026>
- [11] Cost of Data Breach 2026: The human Risk factor | Kymatio. (n.d.). <https://kymatio.com/blog/cost-of-data-breach-2026-human-risk>
- [12] Data Breach statistics for 2026: The numbers IT and security leaders need to know | SWIF. (n.d.). Swif.ai. <https://www.swif.ai/blog/data-breach-statistics>
- [13] Dell'Oro Group. (2026, February 3). 5-Year SASE forecast reaches \$97B as spending nearly triples, according to Dell'Oro Group - Dell's Oro Group. <https://www.delloro.com/news/5-year-sase-forecast-reaches-97b-as-spending-nearly-triples/>
- [14] Devasia, A. (2026, February 15). Cost of Data breach Statistics: All you need to know in 2026. The Network Installers. <https://thenetworkinstallers.com/blog/cost-of-data-breach-statistics/>
- [15] George, D. (2024). 5G-Enabled Digital Transformation: Mapping the landscape of possibilities and problems. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.11583365>
- [16] Edwards, J. (2024, July 1). Well-Intentioned things IT leaders do that hurt team productivity. Information Week. <https://www.informationweek.com/it-leadership/well-intentioned-things-it-leaders-do-that-hurt-team-productivity>
- [17] George, D. (2026a). Security Service Edge (SSE) and SASE: A complete guide to Cloud-Native Zero Trust architecture for enterprise security. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.19974566>
- [18] Edwards, J. (2025, July 23). Want to become a CIO? Here's what you need to know. Information Week. <https://www.informationweek.com/it-leadership/want-to-become-a-cio-here-s-what-you-need-to-know>
- [19] George, D. (2026c). AI-Assisted Incompetence and Fraud in IT-Separating Real Talent from Artificial Credibility in Organizations. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.21314189>
- [20] Fortinet. (n.d.). Secure SD-WAN Buyer's Guide: Key considerations to help you choose. In Fortinet. <https://www.fortinet.com/content/dam/fortinet/assets/ebook/buyer-guide-sd-wan.pdf>
- [21] George, D. (2026b). The Convergence of Intelligence and Infrastructure: A review of Emerging Technology Trends Reshaping Modern Operational Technology (OT) ecosystems. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.20353282>
- [22] FortiSASE | Fortinet. (n.d.). Fortinet. <https://www.fortinet.com/products/sase>
- [23] George, D. (2026d). Generative Artificial Intelligence and Expertise Development among Generation Z: A Critical Review of Cognitive Offloading, Critical Thinking, and AI Literacy. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.21971812>
- [24] Globalgig. (2026, June 18). Networking for the future: What happens when 5G meets SD-WAN? Globalgig. <https://globalgig.com/networking-for-the-future-what-happens-when-5g-meets-sd-wan/>



- [25] George, D., George, A., & Dr.T.Baskar. (2023). SD-WAN Security Threats, Bandwidth Issues, SLA, and Flaws: An In-Depth Analysis of FTTH, 4G, 5G, and Broadband technologies. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.8057014>
- [26] Gray, K. (2026a, March 16). How gaps in IT leadership quietly shrink your C-Suite's bandwidth. Warren Averett CPAs & Advisors. <https://warrenaverett.com/insights/it-leadership/>
- [27] George, D., Dr.T.Baskar, & Srikanth, P. B. (2026). Biometric Authentication Contemporary Trends and Systemic Failures-A Review. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.22055700>
- [28] GTT. (2025, September 15). SD-WAN Architecture explained. GTT. <https://www.gtt.net/resources/blog/sd-wan-architecture-explained/>
- [29] George, D. (2026e). Large language model firewalls runtime safeguards for securing generative systems in enterprise deployment. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.21725826>
- [30] Hancock, S. (2026a, January 22). SASE vs SD-WAN 2026: Network Security Convergence Guide | Socium IT. Socium IT. <https://www.sociumit.com/resources/blog/sase-vs-sd-wan-2026>
- [31] Hancock, S. (2026b, January 22). SD-WAN Costs 2026: "00-\$500/Site + Save 50-84% vs MPLS. Socium IT. <https://www.sociumit.com/resources/blog/sd-wan-cost-guide-2026>
- [32] Hartley, A. (2026, May 26). 3 things Leaders miss when Evaluating an MSP | Hi-Tek Data. Hi-Tek Data MSP. <https://www.hitekdata.com/3-things-leadership-teams-miss-when-evaluating-managed-it-providers/>
- [33] House, N. (2026, October 2). Cyber Security Breach Statistics [2026]: Key Facts & Data. StationX. <https://app.stationx.net/articles/cyber-security-breach-statistics>
- [34] How do 5G and SD-WAN work together? (n.d.). Palo Alto Networks. <https://www.paloaltonetworks.com/cyberpedia/sdwan-and-5g>
- [35] How Much Does SD-WAN Cost?, Zpedia. (n.d.). How much does SD-WAN cost? | ZPedia. <https://www.zscaler.com/zpedia/how-much-does-sd-wan-cost>
- [36] How SD-WAN helps Today's IoT. (n.d.). Palo Alto Networks. <https://www.paloaltonetworks.com/cyberpedia/how-sd-wan-helps-iot>
- [37] Info-Tech Research Group. (2026a, April 13). Develop your IT leadership team's financial literacy. <https://www.infotech.com/research/ss/develop-your-it-leadership-team-s-financial-literacy>
- [38] Info-Tech Research Group. (2026b, April 13). Take a holistic approach to IT financial management. <https://www.infotech.com/research/ss/take-a-holistic-approach-to-it-financial-management>
- [39] INTEGRATED SASE - ATOMNIA. (2025, December 9). ATOMNIA. <https://www.sdwan-solutions.global/omnia/security/>
- [40] Kessler, L. (2026, August 20). Top 10 best Fully Managed SD WAN Services of 2026. Gitnux. <https://gitnux.org/service/fully-managed-sd-wan/>
- [41] Khalil, M. (2026a, July 28). Data Breach Statistics 2025-2026 Global Trends and Insights. DeepStrike. <https://deepstrike.io/blog/data-breach-statistics-2025>
- [42] Khalil, M. (2026b, September 13). 150 Cybersecurity statistics for 2026. DeepStrike. <https://deepstrike.io/blog/cybersecurity-statistics>
- [43] Konzept Conference. (n.d.). Konzept Conference 2026 Phuket | Tech and Innovation. <https://www.konzeptconference.com/knowledge-center/the-tech-debt-is-a-leadership-problem-the-smart-management-way-of-not-failing-in-the-long-term>
- [44] Madawi, M. (2025, October 3). The hidden cost of running Cloud-Hosted SD-WAN for IAAS. MegaPort. <https://www.megaPort.com/blog/the-hidden-cost-of-running-cloud-hosted-sd-wan-for-iaas/>
- [45] Mahajan, G. (2026, September 16). Enterprise Network Infrastructure Market Companies, Size & Trends 2026-2034. <https://www.precedenceresearch.com/enterprise-network-infrastructure-market>
- [46] Mali, S. (n.d.). SD-WAN Market Size, Share, and Growth Forecast, 2026 – 2033 – Persistence Market Research. Persistence Market Research. <https://www.persificencemarketresearch.com/market-research/sd-wan-market.asp>
- [47] Marketplace, C. (2026, September 6). Data Breach Cost Calculator 2026 | AI Threat Edition. BreachCalculator. <https://www.databreachcostcalculator.com/>
- [48] Markets, R. A. (2026a, April 22). \$93.39 Bn Enterprise Network Equipment Market Trends and Growth Forecasts, 2026-2031: Surging Bandwidth Demand from IoT and Edge Computing and Rapid Enterprise-wide Upgrades to Wi-Fi 7 Campus LANs. Yahoo Finance. <https://finance.yahoo.com/sectors/technology/articles/93-39-bn-enterprise-network-084400077.html>

- [49] Markets, R. A. (2026b, May 20). Enterprise Network Infrastructure Market Outlook 2026–2032: Trends, growth challenges, regional opportunities & Strategies, \$73.9 billion Roadmap. Yahoo Finance. <https://uk.finance.yahoo.com/news/enterprise-network-infrastructure-market-outlook-085700203.html>
- [50] Markets, R. A. (2026c, May 20). Enterprise Network Infrastructure Market Outlook 2026–2032: Trends, growth challenges, regional opportunities & Strategies, \$73.9 billion Roadmap. Yahoo Finance. <https://finance.yahoo.com/sectors/technology/articles/enterprise-network-infrastructure-market-outlook-085700920.html>
- [51] MarketsandMarkets. (n.d.-a). Enterprise Networking Market Report 2025–2030, By Networking Devices, Geo, Tech. MarketsandMarkets. <https://www.marketsandmarkets.com/Market-Reports/enterprise-networking-market-168098457.html>
- [52] MarketsandMarkets. (n.d.-b). Software-Defined Wide Area Network (SD-WAN) Market Report 2025 – 2030, By Services, Geo, Tech. MarketsandMarkets. <https://www.marketsandmarkets.com/Market-Reports/software-defined-wan-market-53110642.html>
- [53] MarketsandMarkets Research Pvt. Ltd. (2026, August 19). Enterprise Networking Market Surges to "93.77 billion at a CAGR 9.2% by 2030 | Report by MarketsandMarkets&sup>TM. Yahoo Finance. <https://finance.yahoo.com/technology/articles/enterprise-networking-market-surges-193-150000148.html>
- [54] McFarlane, D., & McFarlane, D. (2026, August 22). Average cost of a data breach 2026 | 46 Facts from IBM & Verizon. CNiC Solutions | We're I.T. <https://cnicsolutions.com/statistics/data-breaches-research/average-cost-of-a-data-breach-statistics-2026/>
- [55] McLaughlin, L. (n.d.). CIO role 2020: Everything you need to know about today's Chief Information Officers. The Enterprisers Project. <https://enterprisersproject.com/article/2019/9/cio-chief-information-officer-role-explained>
- [56] Mendler, A. (2026, August 31). Technical leadership: What every executive should know. Adam Mendler. <https://www.adammendler.com/technical-leadership/>
- [57] Mushroom Networks, Inc. (2024, April 9). Best practices for deploying SD-WAN in distributed enterprise networks. Mushroom Networks. <https://www.mushroomnetworks.com/blog/best-practices-for-deploying-sd-wan-in-distributed-enterprise-networks/>
- [58] Mushroom Networks, Inc. (2025, September 9). How to select the best SD-WAN solution for your enterprise. Mushroom Networks. <https://www.mushroomnetworks.com/blog/best-sd-wan-solution/>
- [59] Network Kings. (2026, July 21). Future of SD-WAN: 2026 Trends, benefits, predictions. <https://www.nwking.com/future-of-sd-wan-trends-benefits-predictions>
- [60] Oliver. (2026, March 26). Why enterprise buyers are rethinking SD-WAN in 2026 – Telecom Review Americas. TelecomReview Canada. <https://www.telecomreviewamericas.com/articles/reports-and-coverage/why-enterprise-buyers-are-rethinking-sd-wan-in-2026/>
- [61] Ouamri, M. A., Alharbi, T., Singh, D., & Sylia, Z. (2024). A comprehensive survey on software-defined wide area network (SD-WAN): principles, opportunities and future challenges. The Journal of Supercomputing, 81(1). <https://doi.org/10.1007/s11227-024-06718-1>
- [62] Penny, & Penny. (2026, September 10). The Leadership Transition Gap: Why technical experts struggle at the next level | Berman Leadership Development. Berman Leadership Development. <https://bermanleadership.com/the-leadership-transition-gap-why-technical-experts-struggle-at-the-next-level/>
- [63] Researcher, T. J. C. (2026, September 23). 250+ data Breach Statistics for 2026. Bright Defense. <https://www.brightdefense.com/resources/data-breach-statistics/>
- [64] Schwegler, K. (2025, September 8). Navigating the Future: Insights from the 2025 SASE and SD-WAN AvidThink Report. Aryaka Unified SASE Solution For Secure. <https://www.aryaka.com/blog/untangling-secure-enterprise-connectivity/>
- [65] SD WAN Deployment | Models & Best practices. (n.d.). <https://nitizsharma.com/the-ultimate-guide-to-sd-wan-deployment-models-and-best-practices/>
- [66] SDWAN Solution Market 2026 to 2034. (n.d.). <https://www.intelmarketresearch.com/sd-wan-solution-market-36296>
- [67] SD-WAN vendor selection: Making the right choice. (2026, April 21). <https://sabertoothpro.com/a/feeds/blog/best-cloud-disaster-recovery-providers-enterprises>
- [68] SD-WAN vs. MPLS: Reliability, Security, Cost, and the Future. (n.d.). Palo Alto Networks. <https://www.paloaltonetworks.com/cyberpedia/sd-wan-vs-mpls>
- [69] SD-WAN vs. SASE vs. SSE: What are the differences? (n.d.). Palo Alto Networks. <https://www.paloaltonetworks.com/cyberpedia/sdwan-vs-sase-vs-sse>



- [70] SD-WAN vs SASE: Which Is Better for Your Organization's Security? | Fortinet. (n.d.). Fortinet. <https://www.fortinet.com/resources/cyberglossary/sd-wan-vs-sase>
- [71] St-Jean, E., Bernstein, C., & McLaughlin, E. (2025, September 10). CIO job description: Chief information officer responsibilities. IT Strategy. <https://www.techtarget.com/it-strategy/reference/CIO-job-description-Chief-information-officer-responsibilities>
- [72] The benefits and challenges of SD-WAN architecture - LogicMonitor. (2026, September 2). LogicMonitor. <https://www.logicmonitor.com/deep-dive/network-admin-guide/sd-wan-benefits>
- [73] The top 10 criteria for selecting an SD WAN vendor. (2019, December 10). Verizon Enterprise. <https://www.verizon.com/business/resources/articles/the-top-10-criteria-for-selecting-an-sd-wan-vendor/>
- [74] ThunderCat Technology. (2025, December 18). SD-WAN Costs & pricing in 2023: The complete guide. ThunderCat Technology - a Service-Disabled, Veteran-Owned Small Business. https://www.thundercattech.com/tcat_blog/sd-wan-cost-pricing/
- [75] Top 7 SD-WAN Challenges: SD-WAN Risks, issues, & Solutions. (n.d.). Palo Alto Networks. <https://www.paloaltonetworks.com/cyberpedia/sd-wan-challenges>
- [76] Understanding the top benefits of SD-WAN - The Comprehensive Guide | Fortinet. (n.d.). Fortinet. <https://www.fortinet.com/resources/cyberglossary/benefits-of-sd-wan>
- [77] UpUply. (2025, December 19). sd wan cost: Comprehensive Analysis of Costs, Models, ROI and Optimization. UpUply. <https://www.upuply.com/blog/sd-wan-cost>
- [78] Venter, L. (2026, October 1). When being a great IT manager becomes a barrier to leadership. Bristow Holland. <https://www.bristowholland.com/insights/industry-insights/great-it-manager-becomes-a-barrier-to-leadership/>
- [79] Versa Networks. (2026, April 28). Versa frees enterprise 5G from the wiring closet with new SD-WAN-Native extender. <https://versa-networks.com/news/2026/versa-frees-enterprise-5g-from-the-wiring-closet-with-new-sd-wan-native-extender/>
- [80] Violino, B., & Violino, B. (2022, March 14). 5 SD-WAN gotchas to avoid. Network World. <https://www.networkworld.com/article/970766/5-sd-wan-gotchas-to-avoid.html>
- [81] Vliegen, H. (2026, August 4). Is your SD-WAN ready for AI-powered operations? Cisco Blogs. <https://blogs.cisco.com/security/ai-sdwan-readiness>
- [82] Weinberg, N., & Weinberg, N. (2021, July 20). SD-WAN buyer's guide: Key questions to ask vendors (and yourself). Network World. <https://www.networkworld.com/article/970038/sd-wan-buyers-guide-key-questions-to-ask-vendors-and-yourself.html>
- [83] What is SASE (Secure Access Service Edge)? | A starter guide. (n.d.). Palo Alto Networks. <https://www.paloaltonetworks.com/cyberpedia/what-is-sase>
- [84] What is SD-WAN? Definition & Solutions. (n.d.). Fortinet. <https://www.fortinet.com/products/sd-wan>
- [85] What is SD-WAN? Software-Defined Wide Area Network. (n.d.). Palo Alto Networks. <https://www.paloaltonetworks.com/cyberpedia/what-is-sd-wan>
- [86] Wikipedia contributors. (2026a, August 8). Chief information officer - Wikipedia. https://en.wikipedia.org/wiki/Chief_information_officer
- [87] Wikipedia contributors. (2026b, September 25). SD-WAN. Wikipedia. <https://en.wikipedia.org/wiki/SD-WAN>
- [88] Yelland, H. (2026a, March 11). SD-WAN disadvantages & problems: What IT teams should know. Netify. <https://netify.co.uk/insights/problems-with-sd-wan/>
- [89] Yelland, H. (2026b, April 17). SD-WAN RFP Questions: Complete Guide with Copy-Paste Examples. Netify. <https://netify.co.uk/insights/sd-wan-rfp-template-complete-guide-and-copy-paste-example/>
- [90]
- [91] Alchemy (2018). SD WAN: A pragmatic perspective. [online] Alchemy Tech Group. Available at: <https://alchemytechgroup.com/blog/sd-wan-a-pragmatic-perspective>
- [92] Anusha Dhasarathy, Gill, I. and Khan, N. (2020). The CIO challenge: Modern business needs a new kind of tech leader. [online] McKinsey & Company. Available at: <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-cio-challenge-modern-business-needs-a-new-kind-of-tech-leader>
- [93] Author, N. (2021). 7 Business skills every IT leader needs to succeed. [online] CIO. Available at: <https://www.cio.com/article/191559/7-business-skills-every-it-leader-needs-to-succeed.html>
- [94] Cisco. (2025). Benefits of SD-WAN - Cisco SD-WAN. [online] Available at: <https://www.cisco.com/site/us/en/learn/topics/networking/what-is-sd-wan.html>



- [95] Cisco. (n.d.). Design zone for branch/WAN - Cisco SD-WAN design guide. [online] Available at: <https://www.cisco.com/c/en/us/td/docs/solutions/CVD/SDWAN/cisco-sdwan-design-guide.html>
- [96] Lightyear. (2025). Lightyear. [online] Available at: <https://lightyear.ai/blogs/sd-wan-cost>
- [97] Open Systems. (2019). Open systems. [online] Available at: <https://open-systems.com/blog/sd-wan-will-drive-enterprise-adoption-of-5g-services/>.
- [98] Rigdon, S. (2024). 8 Business benefits of using SD-WAN | vonage. [online] Vonage.com. Available at: <https://www.vonage.com/resources/articles/benefits-of-sd-wan/>.
- [99] Shinde, P. (2024). Essential decision making frameworks for technical leaders. [online] Priyanka Shinde. Available at: <https://www.thepriyankashinde.com/post/essential-decision-making-frameworks-for-technical-leaders>.
- [100] Tiwari, V., Singh, N. and Canales, C. (2024). Forecast: Enterprise network equipment, worldwide, 2024-2030, 2Q26. [online] Gartner. Available at: <https://www.gartner.com/en/documents/>
- [101] www.catonetworks.com. (n.d.). Secure access service edge (SASE). [online] Available at: <https://www.catonetworks.com/sase/>
- [102] www.fortunebusinessinsights.com. (n.d.). Enterprise networking market size, industry share | forecast, 2032. [online] Available at: <https://www.fortunebusinessinsights.com/enterprise-networking-market-105887>