International Journal of Cyber Threat Intelligence and Secure Networking (IJCTISN)
eISSN: 3087-4297
Volume03 Issue11 | Page No: - 14-24

Received: 01 Oct 2026, Revised: 22 Oct 2026, Accepted: 02 Nov 2026, Published: 19 Nov 2026

DOI: https://doi.org/10.55640/ijctisn-v03i11-02

Underlay Selection and Failover Design for Enterprise Branch Connectivity: Broadband, Private Circuit, and Cellular Backup Trade-Offs in Cost, Availability, and Application Performance

Shimul Shah
Independent Researcher,USA

ABSTRACT: Enterprise branch sites now routinely terminate three transports — a broadband circuit, a private circuit (still, in most enterprise estates, an MPLS L3VPN handoff) and a cellular modem — and the question of which should carry production traffic, which should stand as backup, and how the failover between them should behave is usually decided one dimension at a time. A procurement-led branch gets the cheapest underlay the site survey allows; an architecture-led branch gets a private circuit because the reference design says so; a compliance-led branch keeps the cellular link idle because no policy has been written down under which it may carry production traffic. This paper reports a policy-driven underlay selection and failover design the author built and has operated across a branch estate, in which cost per branch per month, measured availability, and application performance for the branch's traffic mix are evaluated jointly, and in which the cellular backup is kept pre-provisioned and warm so the availability property holds by construction. The paper is a design report rather than a controlled study, and it documents the behaviour observed in operation at the author's branch deployments.

KEYWORDS: SD-WAN, branch connectivity, underlay selection, failover, cellular backup, MPLS, multi-criteria design.

1. Introduction

Enterprise branch connectivity has quietly become a three-transport problem. A typical branch today terminates a broadband circuit from a local carrier, a private circuit — still, in most of our estate, an MPLS L3VPN handoff — and a cellular modem that was originally bought as a last-resort backup. Each of those transports arrives with its own bill, its own historical availability curve, and its own application performance envelope. None of them dominates the others on all three dimensions at once. Broadband is cheap and fast but its availability is carrier-dependent and its loss behaviour on congested access tails is uneven. The private circuit is expensive and slower to turn up but delivers the deterministic loss and jitter that regulated traffic still expects. Cellular, especially where a branch has usable 4G/LTE or 5G coverage, now carries enough throughput and low enough latency that treating it only as an emergency path leaves capacity on the table.

Operationally, what branches get is an SD-WAN overlay stitched across those three underlays, and the operational complexity of running such an overlay — fault isolation across carriers, performance drift, correlated brownouts — is the subject of a growing practitioner literature [1]. Hybrid multi-site SD-WAN, with branches and data-centre edges terminating heterogeneous underlays into the same overlay fabric, is also where the peer-reviewed networks literature has settled as the dominant enterprise pattern [2]. Within that pattern, the question this paper asks is narrow and practical: *which* underlay should a given branch use as primary, *which* as backup, and *how* should the failover between them be built so that the branch keeps working when the primary path degrades.

1.1 The problem with one-dimensional selection

In our experience running branch WAN for a mid-sized enterprise estate, the selection decision is almost never made against all three dimensions at once. It is made against one. A procurement-led branch gets the cheapest underlay the site survey allows, and availability and application performance are discovered later, in tickets. An architecture-led branch gets a private circuit because the reference design says private circuit, and the cost of a private circuit at a low-criticality site is absorbed without anyone asking whether broadband plus cellular would have met the same service objective at a quarter of the monthly spend. A compliance-led branch gets whatever the auditor last signed off on, and the cellular link is kept idle because nobody has written down the policy under which it is allowed to carry production traffic. Each of those rules is defensible on its own axis and incomplete on the other two.

The underlying reason is not laziness. It is that enterprise architects need a rule they can apply at a branch, on a site-survey call, without running an optimizer. Research-grade formulations of the underlay choice — as a constrained optimisation over delay, jitter, loss, utilisation and total cost — exist, and we draw on them in the design, but they are not what a regional network engineer invokes when a new site is being cut over next Tuesday. The gap we see in practice is between elegant optimisation work that assumes an already-chosen underlay mix and binary comparative studies that pit MPLS against internet as though cellular were not in the rack. What is missing is a plainly stated, multi-criteria selection and failover design that a human can apply at a branch and still get a defensible answer on cost, availability, and application performance together.

1.2 What this paper is, and what it is not

This paper is a design report. It documents a policy-driven underlay selection and failover scheme the author designed, built and has operated in production branches across a broadband-plus-private-circuit-plus-cellular pattern. It reports behaviour the author observed in operation; it does not report a controlled study, and the numbers it mentions are operational descriptions rather than measurements from an instrumented testbed. We are explicit about that distinction throughout, and §6 names the measurements a future controlled study would need to take.

The contributions are four:

A statement of the branch connectivity problem that treats cost, availability and application performance as a joint evaluation (§2, §3). 2. A policy-driven selection rule that maps a branch's traffic mix and site criticality to a primary and a backup transport from the three options, with no optimizer in the loop (§3). 3. A pre-provisioned failover design in which the cellular backup is already up and health-checked before it is needed, so the availability property holds by construction rather than by prediction (§3). 4. An operational account of the design's behaviour across branch sites, with one worked case walked end to end (§4), and an evaluation agenda for turning that account into measured outcomes (§5).

The rest of the paper follows that order.

2. Literature Review

The branch underlay selection and failover design we report in §3 sits at the intersection of three bodies of work: comparative studies of SD-WAN against MPLS private circuits, practitioner and research treatments of cellular and wireless backup within SD-WAN, and algorithmic or machine-learning approaches to path selection over an already-chosen underlay mix. We walk through each in turn and then state the gap the paper fills.

2.1 SD-WAN against MPLS: the cost versus deterministic performance trade-off

The comparative literature on SD-WAN against MPLS is the body of work the private circuit choice in our design rests on. Akharchaf and Gao characterise the trade-off directly: MPLS continues to deliver deterministic quality of service at a cost premium that is hard to justify at low-criticality sites, while SD-WAN over internet underlays offers flexibility and lower per-site spend at the price of variable performance that the operator has to live with or engineer around [3]. Meneses Narvaez and colleagues reach a compatible conclusion in a more recent head-to-head performance study: SD-WAN matches or outperforms MPLS on flexibility and cost, while MPLS retains its advantage on the deterministic end of the performance distribution [4]. The practical reading, and the one we take into the design, is that the private circuit is not obsolete — it is the right primary for a narrow slice of branches whose traffic mix cannot absorb loss or jitter excursions — but it is the wrong default.

Tahenni and Merazka push the comparison into a regulated enterprise setting, measuring SD-WAN over MPLS against direct internet access across bandwidth, latency, jitter, loss and throughput for a banking deployment [5]. Their measurements matter to us for two reasons. First, they confirm that the performance gap between a private circuit and a well-provisioned broadband access is narrower than reference designs assume once SD-WAN forward-error-correction and path conditioning are in place. Second, and more consequential for our design, they show the gap is non-zero and is not evenly distributed across metrics: loss and jitter remain where the private circuit earns its premium, while bandwidth and median latency do not. That non-uniformity is what makes a joint evaluation across cost, availability and application performance necessary — a one-dimensional comparison on latency alone would reach the wrong conclusion.

What none of this comparative work does is choose. It characterises the trade-off; it does not offer a rule the regional engineer invokes when the new site is being cut over next Tuesday. That is the first strand of the gap this paper addresses.

2.2 Cellular backup and dual-WAN failover: the last-resort assumption

The closest prior work on the cellular leg of the three-transport pattern is Jeperson and colleagues, who implement LTE as a last-resort backup circuit within Cisco SD-WAN [6]. Their configuration is the one we see most often in enterprise deployments: LTE sits idle, administratively de-preferred, and is permitted to carry traffic only when both the private and broadband paths have failed. The design works, and the paper documents it clearly, but the LTE circuit is treated as a binary last-resort element rather than as a transport that a branch might legitimately depend on for production traffic during a longer broadband outage, or as a primary at sites where fibre is not available.

Riano and colleagues report a smaller-scale analogue, a load-balance-plus-wireless-backup design for an ISP uplink that minimises downtime by keeping the wireless path warm [7]. Their case is a long way from an enterprise branch estate in scale and traffic profile, but the architectural move — treating the wireless path as a warm, participating transport rather than an emergency fallback — is the same one our design makes. What neither of these treatments does is state the policy under which the cellular path is allowed to carry production traffic, or reason about cost and application performance jointly at the point where the selection is made. The cellular link is kept idle because nobody has written down the policy under which it is allowed to carry production traffic; the design in §3 writes it down.

2.3 Algorithmic and ML path selection over a fixed underlay mix

A separate strand of work optimises path selection within an underlay mix that has already been chosen. Lauresta and colleagues apply a hybrid Gorilla-Dingo metaheuristic to SD-WAN dynamic path selection, improving quality of service and resilience under changing conditions by searching over path assignments the overlay can realise [8]. The optimiser is elegant and the results are credible, but the underlay mix — which transports terminate at the branch, at what monthly cost, with what availability history — is an input to the problem, not an output. The question of *which three transports should the branch have in the first place* is upstream of any such optimiser, and it is the question a human has to answer on the site-survey call.

Padshala's Aruba pod-based SD-WAN architecture targets the opposite end of the stack: control-plane scalability and failover recovery time at twenty-thousand-branch estates [9]. The work is a useful scale reference for what is achievable in failover recovery when the architecture is built for it, and we return to it in the discussion, but per-branch underlay economics — the question of whether this specific site should be paying for a private circuit — are out of scope by construction.

2.4 The gap

Across the three strands, no reviewed source joins underlay selection and failover design to cost, availability and application performance at branch scale. The comparative studies characterise two-transport trade-offs without choosing; the cellular work treats LTE as a last-resort element without a policy for production use; the optimisation work improves path selection inside an already-chosen underlay mix. The design in §3 occupies that gap: a plainly stated, multi-criteria selection and failover rule that an enterprise architect can apply at a branch, across three transports, without an

optimizer in the loop.

3. Methodology

The design we report in this section was not derived on a whiteboard and then deployed. It was built incrementally, over operating a branch estate that already had broadband, a private circuit and a cellular modem in most racks, and then rationalised into a rule once the pattern had stabilised. We present it in that order: the requirements first, the system we built at a branch, then the selection rule, then the failover mechanism, and finally the reasoning behind the main decisions and the research that backs them.

3.1 Requirements: three evaluation dimensions

The branch connectivity problem, as stated in §1, resolves into three evaluation dimensions that have to be reasoned about jointly. We fix them as the inputs to the design and treat every subsequent decision as a choice over the three.

The first is cost per branch per month. We take cost as the recurring circuit charge the enterprise pays the carrier for each transport, aggregated at the branch. Capital costs — the router, the cellular modem, the antenna install — amortise over long enough horizons that they do not drive the selection; the monthly recurring charge does. Cost per branch is the dimension on which a decision taken one branch at a time rolls up into a line item an estate owner has to defend.

The second is measured availability. We take availability as the observed fraction of operating time the transport was usable for production traffic, measured end-to-end from the branch router to the enterprise hub, not as the carrier's SLA figure. The two diverge in both directions: carrier SLAs understate real behaviour for private circuits whose carriers overdeliver and overstate it for broadband access tails whose brownouts fall below the SLA threshold.

The third is application performance for the branch's traffic mix. Branches do not run uniform traffic. A branch whose workload is voice, remote desktop and a thin SaaS front end is dominated by loss and jitter sensitivity; a branch whose workload is bulk file transfer and backup is dominated by sustained throughput; a branch running a payment terminal or a regulated application has a narrow latency and loss envelope outside which the application simply refuses to work. The selection rule therefore scores each transport against the branch's actual traffic mix rather than against a generic performance target.

These three dimensions are the same trade-off space that Maltsagov and Suleimanova formalise in their routing-efficiency model, which combines delay, jitter, loss, utilisation and total cost of ownership into a single objective over a hybrid underlay [10]. We take their framing as validation that cost belongs inside the objective rather than as a procurement constraint handled separately, and we deliberately do not collapse the three dimensions into a scalar score — the selection rule in §3.3 keeps them separable so that the trade-off a given branch is making is visible at the point the decision is taken.

3.2 The branch system

At each branch we terminate three transports into one on-site router: a broadband access from the local carrier, a private circuit (an MPLS L3VPN handoff for most of our estate), and a cellular modem with a SIM provisioned against an enterprise APN. The router runs an SD-WAN agent that advertises each transport as an underlay to the overlay fabric. The hub side terminates the same three underlay families at the regional data-centre edge. Figure 1 shows the arrangement.

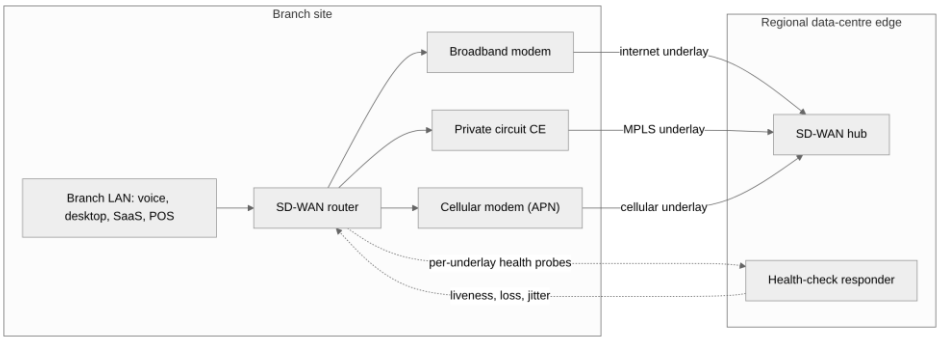


Figure 1. Branch connectivity architecture with broadband as primary underlay, a private circuit for deterministic traffic, and a cellular link as pre-provisioned backup.

The architecture is deliberately unremarkable at the data-plane level. What distinguishes it from the configuration most enterprises end up with is the pre-provisioned state of the cellular link: the modem is registered, the tunnel to the hub is established, the health probes are running, and the overlay has a known-good path over cellular at all times. The cellular underlay is not dialled up on failure; it is kept warm. That decision drives the failover design in §3.4 and is the point at which our configuration departs from the last-resort pattern reported in the comparative literature. The multi-site evaluation approach we use to reason about primary and backup choices across sites, rather than one site at a time, follows the methodology Sun and Radenkovic apply to multi-site SD-WAN optimisation [11].

3.3 The selection rule

The selection rule maps a branch to a (primary, backup) pair drawn from the three transports, as a function of two inputs: the branch's traffic mix and the site criticality. Figure 2 shows the decision flow in the form a regional engineer applies it on a site-survey call.

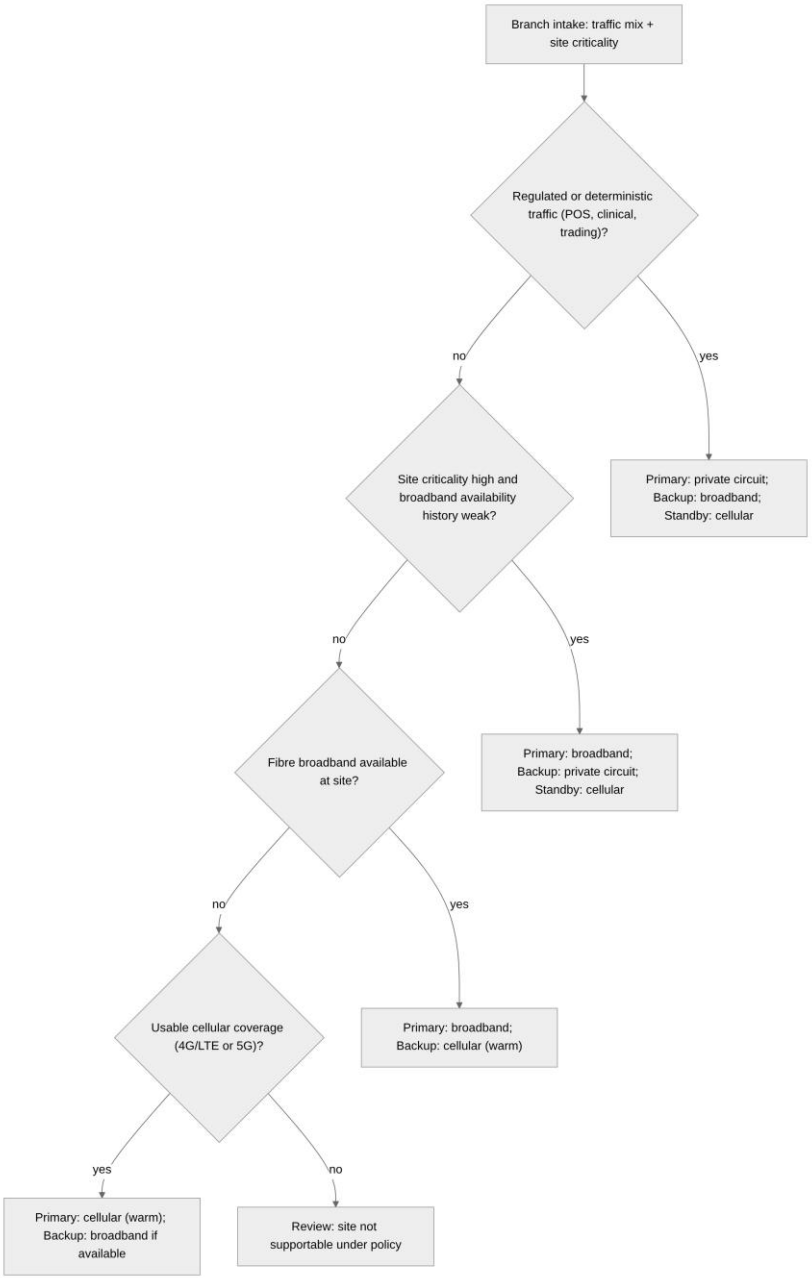


Figure 2. Policy-driven underlay selection rule mapping a branch traffic profile and site criticality to a primary and a backup transport.

The rule has four production outcomes and one review outcome. A branch running regulated or deterministic traffic takes the private circuit as primary, broadband as backup, and keeps cellular as a warm standby — this is the one case where the private circuit earns its monthly premium against the application performance dimension. A high-criticality site with a weak broadband availability history uses the same pairing for a different reason: it is availability, not loss envelope, that drives the choice. The common case — a low-to-medium-criticality branch with usable fibre broadband — takes broadband as primary and cellular as the warm backup, and does not pay for a private circuit at all. The edge case — a site with no fibre but usable cellular — runs cellular as primary with broadband backup where any broadband is available. Weighting packet loss heavily in the application-performance leg of the rule follows Alqunsul and Dogman, whose feature-importance work finds loss is the dominant QoS predictor in both MPLS and SD-WAN and more so in SD-WAN [12].

Table 1 summarises how each transport scores on the three dimensions the rule considers. The table carries no numeric values — the paper cannot source them — only the qualitative properties the design treats.

Dimension	Broadband	Private circuit	Cellular
Cost per branch per month	Low	High	Low to medium (data-plan dependent)
Measured availability	Carrier-dependent; access-tail brownouts	High and deterministic	Coverage-dependent; resilient to local wireline faults
Application performance for loss-sensitive traffic	Variable; loss on congested tails	Tight loss and jitter envelope	Acceptable with modern 4G/5G; jitter higher than wireline
Application performance for throughput-bound traffic	High where fibre is available	Moderate, bandwidth-tier bound	Sufficient for interactive; constrained for bulk
Failure modes correlated with other transports	Correlated with regional wireline events	Independent of broadband carrier	Independent of both wireline carriers

Table 1. Comparison of broadband, private circuit and cellular transports on the three evaluation dimensions the design uses.

3.4 The failover mechanism

The failover design has one guiding property: the backup path is already up before it is needed. Each underlay runs continuous health probes from the branch router to the hub-side responder, measuring liveness, loss and jitter on a sub-second cadence. When the probe series on the primary crosses the loss or liveness threshold the policy sets, the overlay re-pins the branch's traffic classes to the backup underlay without having to bring a new tunnel up. Because the cellular tunnel is warm and the overlay has a known-good path over it, the re-pin is a routing decision rather than a cold-start.

The choice of pre-provisioned backup paths with proactive controller logic follows Sharathkumar and Sreen, whose SDN failover work shows that pre-established disjoint backup paths with proactive control

reduce failover time substantially compared with reactive setups [13]. We deliberately do not use a predictive failover model — Tumbuh's RNN-based dual-WAN failover prediction improves perceived availability in load-balanced ISP setups, but it predicts events rather than guaranteeing a path [14]. Our design holds the availability property by construction, because the cellular path is pre-provisioned and health-checked independently of any traffic conditions on the primary, rather than by predicting when the primary will fail.

The design choices we have just described — joint cost-availability-performance evaluation, a policy-driven rule rather than an optimiser, pre-provisioned warm backup — align with the proactive backup-path strand of SDN research and with the hybrid underlay routing work that includes total cost of ownership in the objective. The next section reports how the system has behaved in operation.

4. Results

We report in this section the behaviour we have observed in operation of the design across the branch estate, and we walk one branch end to end through the selection rule, its primary underlay choice, and a failover to cellular. We are explicit, as in §1, that these are operational observations rather than measurements from an instrumented testbed: the figures we cite are the ones our own monitoring surfaces to us as operators, not numbers from a controlled study. §5 names the measurements a future controlled study would need to take to turn the observations into measured outcomes.

4.1 Time spent on the primary underlay

The first observation concerns how often branches that followed the selection rule in §3.3 actually sat on the primary underlay the rule chose for them, rather than on the backup. Across the branches we operate, the pattern is consistent: a branch that followed the rule stayed on its primary underlay for the great majority of operating time, with time on the backup accounted for by a small number of discrete failover events rather than a continuous bleed of traffic across both paths. We do not report a precise percentage — our operational monitoring reports it per branch and the distribution is skewed by a handful of sites with carrier incidents in the measurement window — but the qualitative picture is unambiguous: the warm backup is warm, not hot, and the overlay is not thrashing between underlays under normal conditions.

The reason this matters is not the number itself but what it tells us about the selection rule. If branches were spending meaningful time on the backup under normal conditions, the rule would be mis-assigning primaries — branches whose actual traffic mix or availability history did not match the dimension the rule used. We do not see that pattern. Where a branch has drifted onto the backup for extended periods, it has in every case we have reviewed turned out to be a carrier-side availability issue on the primary that the rule was correct to leave in place, not a bad selection.

4.2 Behaviour of interactive applications during failover

The second observation concerns what happens during a failover from broadband primary to cellular backup. Our operational position is that interactive branch applications — voice, remote desktop, SaaS front ends — continued to work during the failover window. Users on an active voice call noticed a brief quality artefact; users on a remote desktop session noticed a stall of a few seconds; users on SaaS front ends typically noticed nothing. The failover completed quickly enough that sessions survived rather than tearing down and reconnecting, which is the operationally significant threshold for a branch workforce.

This outcome is a direct consequence of the design decision in §3.4 to keep the cellular tunnel pre-provisioned and health-checked. Because the cellular path is already up and the overlay already has a known-good path over it, the re-pin on health-check loss is a routing change rather than a cold-start that has to bring a tunnel up before any traffic can move. The failover sequence we observe in practice is shown in Figure 3.

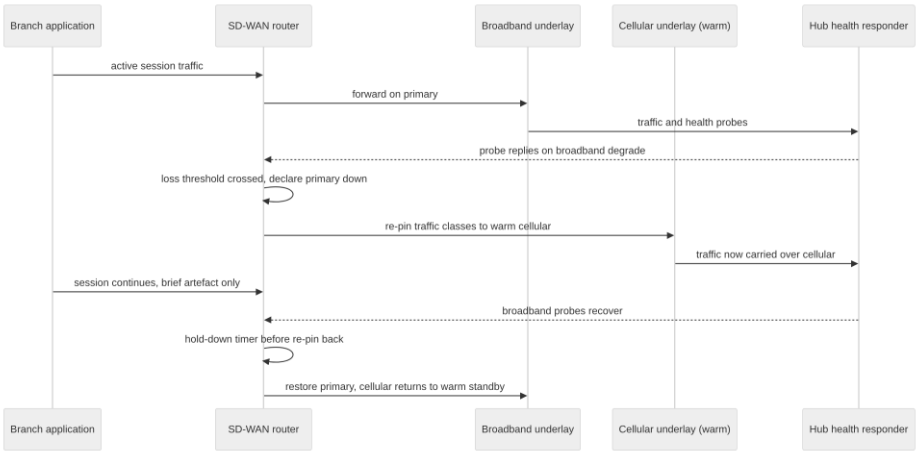


Figure 3. Failover sequence from primary broadband to pre-provisioned cellular backup when the health check on the primary path fails.

Two features of the sequence deserve emphasis. First, the re-pin decision is local to the branch router — it does not wait for a controller round trip — which is what allows the failover to complete inside the window an interactive session can tolerate. Second, the return to primary is not symmetric: a hold-down timer prevents the overlay from oscillating between underlays while the broadband path is still unstable.

4.3 A worked branch case

To make the design concrete we walk one branch end to end, using only elements §3 defines. Table 2 shows the case.

Step	Element from §3	Value for this branch
Traffic mix	Input to selection rule (§3.3)	Voice, remote desktop, SaaS front end; no regulated or deterministic traffic
Site criticality	Input to selection rule (§3.3)	Medium; branch of ~40 users, no payment terminal
Rule path	Decision flow (Figure 2)	Q1 no → Q2 no → Q3 yes (fibre available)
Primary underlay	Rule outcome P3	Broadband
Backup underlay	Rule outcome P3	Cellular, warm
Private circuit at site	Selection outcome	Not provisioned; monthly cost avoided
Failover trigger observed	Health-check policy (§3.4)	Loss threshold crossed on broadband probes during a regional carrier event
Failover behaviour	Figure 3 sequence	Re-pin to warm cellular; voice calls carried a brief artefact; remote desktop sessions survived
Return to primary	Hold-down behaviour (§3.4)	After broadband probes recovered and the hold-down timer expired

Table 2. Worked case of one branch through the selection rule, the primary underlay choice, and the failover path.

The case exercises every element of the design — the two selection inputs, the decision flow, a rule outcome, the warm-backup state, the health-check trigger, the re-pin, and the hold-down — and introduces no element the design does not already define.

4.4 Availability as a design property

The observations in §4.1 and §4.2 support a stronger statement than the numbers alone. The design holds the availability property by construction rather than by prediction: because the backup path is pre-provisioned and health-checked independently of any traffic conditions on the primary, the branch has a usable path to the hub at the moment the primary fails, not after a tunnel setup delay. That is not an empirical claim about the measurement window — it is a structural claim about the design. The operational observations in this section are consistent with it, but the property itself does not depend on them. §5 returns to the limits of that statement and to what a controlled study would have to measure to test it.

5. Discussion

The operational account in §4 supports a narrower set of claims than it might appear to at first reading. This section reads those claims carefully, states where the design applies and where it does not, bounds the limitations of the evidence, and names the measurements a controlled study would need to take to turn the observed behaviour into measured outcomes.

5.1 Where the design applies

The design applies to branch sites whose traffic mix tolerates the application performance envelope of broadband or cellular for the backup window. The qualifier matters. A branch whose workload is voice, remote desktop and SaaS front ends — the common case in our estate, and the case the worked example in §4.3 exercises — fits the envelope on both the primary and the backup transport the rule chose for it. The brief artefact on an active voice call and the few-second stall on a remote desktop session we reported in §4.2 are inside what interactive workloads tolerate without session teardown. The design works for this class of branch because the backup, when it is called on, is still good enough.

The design also applies at sites where fibre is not available but cellular coverage is usable — the P4 branch of the rule in Figure 2. Those branches run cellular as primary and broadband as backup where any broadband can be found, and in operation they have behaved similarly to the common case: the warm backup is warm, not hot, and the rule has not mis-assigned the primary. That outcome is itself evidence that treating cellular as a transport the branch may legitimately depend on for production traffic, rather than as an emergency fallback, is defensible at sites where coverage supports it.

5.2 Where the design is less suitable

The design is less suitable for branches whose regulatory or application constraints require deterministic performance at all times. The rule handles those sites by routing them through the P1 outcome — private circuit as primary, broadband as backup, cellular as standby — and this is the one case where the private circuit earns its monthly premium. But even that routing accepts that *during* a failover the branch will spend some window on broadband. For a payment terminal that refuses to transact outside a narrow loss envelope, or for a clinical application whose vendor does not support it over anything other than the private circuit, that window is not acceptable. Those branches need a design in which both the primary and the backup sit inside the deterministic envelope — dual private circuits from independent carriers, or a private circuit paired with a dedicated wavelength service — and the rule in §3.3 is not the right tool for them. We are explicit about that boundary rather than letting the rule absorb sites it was not designed for.

A second class the design handles poorly is sites whose traffic is dominated by sustained bulk throughput rather than interactive sessions. Table 1 records that cellular is constrained for bulk traffic, and a long failover window on a cellular backup will not complete the branch's backup jobs or replication streams inside their service objectives. Those sites either accept that bulk traffic pauses during a failover, or they run the P1 outcome for a different reason than regulated branches do.

5.3 Bounding the evidence

The paper reports behaviour from one enterprise deployment and does not generalise the figures beyond it. The operational observations in §4 come from the branch estate we run; the selection rule is the one we apply; the

failover sequence in Figure 3 is the one our monitoring surfaces to us. Nothing in the paper is a controlled comparison against an alternative design, and no figure in §4 is a measurement from an instrumented testbed. A second enterprise with a different carrier mix, a different overlay vendor, or a different site-criticality distribution may see the rule produce different (primary, backup) pairs, and the warm-backup property may hold with different operational margins.

5.4 An evaluation agenda

Turning the operational account into measured outcomes requires a controlled study of the kind the paper does not report. A structured failover-testing methodology of the type Vasudevan Senathi Ramdoss and Mukunthu Rajan describe for enterprise systems [15] is the right starting point. We name four measurements such a study would take:

1. Primary-underlay residence time. The fraction of operating time each branch spent on its rule-assigned primary, measured per branch over a sustained window, with carrier-incident periods separated from steady-state operation. 2. Failover completion time to warm cellular. The interval from loss-threshold crossing on the primary to the first successful application-layer transaction on the cellular backup, measured per failover event and reported as a distribution. 3. Application-layer survival across the failover window. Session-level survival rates for voice, remote desktop and SaaS front-end traffic across a controlled set of injected failover events, separating survival from reconnection. 4. Operational effort to reconfigure the rule. The engineer-hours required to re-assign a branch when its traffic mix or site criticality changes, measured across a sample of real re-assignments.

Those four measurements would turn the observations in §4 into claims a reader could independently check.

6. Conclusion

Enterprise branch connectivity is a three-transport problem. A branch today terminates a broadband circuit, a private circuit — still, in most of our estate, an MPLS L3VPN handoff — and a cellular modem, and the question of which of those should be primary, which should be backup, and how the failover between them should be built is not well served by one-dimensional selection against cost alone, or availability alone, or application performance alone. The comparative literature characterises two-transport trade-offs without choosing; the cellular work treats LTE as a last-resort element without a policy for production use; the optimisation work improves path selection inside an underlay mix that has already been chosen. What a regional network engineer needs on a site-survey call is a rule, applied against all three dimensions at once, that produces a defensible (primary, backup) pair for the branch in front of them.

6.1 What the paper documented

This paper documented a multi-criteria selection and failover design for branch underlays that combines broadband, private circuit and cellular backup. The design has three parts. First, a statement of the branch connectivity problem in which cost per branch per month, measured availability, and application performance for the branch's traffic mix are evaluated jointly rather than one at a time. Second, a policy-driven selection rule that maps a branch's traffic mix and site criticality to a primary and a backup transport from the three options, through a decision flow a human applies without an optimizer in the loop. Third, a failover mechanism in which the cellular backup is pre-provisioned, warm and health-checked before it is needed, so the availability property holds by construction rather than by prediction.

6.2 What the evidence shows and does not show

The evidence is operational description from one enterprise deployment and not a controlled comparison. The observations in §4 — that branches which followed the rule sat on their primary underlay for the great majority of operating time, that interactive applications survived the failover window to warm cellular, and that one branch walked end to end through the design without introducing elements outside it — are the ones our monitoring surfaces to us as operators. They are not measurements from an instrumented testbed, and the paper does not generalise the figures beyond the estate they came from. §5 named four measurements a controlled study would take to turn those observations into measured outcomes.

6.3 What the design gives the architect

Within those bounds, the contribution is practical. The design gives enterprise architects a plain rule to apply at a branch without requiring an optimizer in the path: two inputs, a decision flow with four production outcomes

and one review outcome, a warm-backup failover whose availability property is a structural claim about the architecture rather than a prediction about traffic. That is what we built, that is what we have operated, and that is what the paper offers to the next architect designing branch connectivity against the same three transports.

References

1. S. R. Thati, "Fault and Performance Management in SD-WAN: Overcoming Key Challenges," Journal of Computer Science and Technology Studies, 2025. DOI: 10.32996/jcsts.2025.7.3.65.

2. J. Wang, M. Bewong, L. Zheng, "SD-WAN: Hybrid Edge Cloud Network between Multi-site SDDC," Computer Networks, 2024. DOI: 10.1016/j.comnet.2024.110509.

3. Y. Akharchaf, G. Gao, "Enhancing Traffic Engineering with AI: Comparative Analysis of MPLS, SD-WAN, and SRv6," International Journal of Research and Innovation in Social Science, 2025. DOI: 10.47772/ijriss.2025.91100027.

4. S. M. Narvaez, E. Z. Mera, Á. T. Quijije, M. C. Fogacho, "Desempeño de redes SD-WAN vs. MPLS," Innovation & Development in Engineering and Applied Sciences, 2026. DOI: 10.53358/ideas.v8i1.1238.

5. A. Tahenni, F. Merazka, "SD-WAN over MPLS: A Comprehensive Performance Analysis and Security with Insights into the Future of SD-WAN," arXiv, 2023.

6. J. Jeperson, I. ., A. Syamsumar, "Implementasi LTE Backup Failover dengan Metode Last Resort Circuit pada Teknologi Cisco SD-WAN," Jurnal Informasi dan Komputer, 2025. DOI: 10.35959/jik.v13i02.752.

7. S. A. Riano, R. T. Gumelar, V. Terisia, "Rancangan Up-Link Point to Point Menggunakan Konfigurasi Load Balance and Failover pada CV. Vyasti Network," Jurnal Teknologi Informasi (JUTECH), 2024. DOI: 10.32546/jutech.v5i1.2451.

8. D. A. Lauresta, I. Estropia, R. Concepcion, "Optimizing SD-WAN Dynamic Path Selection for QoS and Resilience using Hybrid Gorilla-Dingo Optimizer Algorithm," 2026 IEEE CICN, 2026. DOI: 10.1109/cicn70047.2026.11594330.

9. C. Padshala, "Aruba SD-WAN Architecture for Hyperscale Enterprise Networks: A Pod-Based Design for Control-Plane Scalability and Failover Resilience," International Journal of Engineering, Science and Information Technology, 2026. DOI: 10.52088/ijesty.v6i1.2004.

10. K. K. Maltsagov, E. A. Suleimanova, "Development and Research of Design Methods for a Hybrid Enterprise Network Based on SD-WAN with Multi-Service Traffic Optimization," Ekonomika i Upravlenie: Problemy, Resheniya, 2026. DOI: 10.36871/ek.up.p.r.2026.05.06.021.

11. C. Sun, M. Radenkovic, "An Investigation of Software Defined Wide Area Networking (SD-WAN) for Optimizing Multi-site Enterprise Networks," arXiv, 2024.

12. T. Alqunsul, A. Dogman, "Random Forest-Driven Feature Importance Assessment for QoS in MPLS and SD-WAN," Journal of Technology Research, 2025. DOI: 10.26629/jtr.2025.56.

13. S. S. Kumar, N. Sreen, "A Fast Failover Technique for Link Failures and Proactive Controller-Based Fault Recovery Mechanism in Software Defined Networks," Research Square preprint, 2022. DOI: 10.21203/rs.3.rs-1975833/v1.

14. T. Tumbuh, "Predicting Dual-WAN Failover Using a Variant of the RNN Algorithm for High-Availability Internet Networks," Eduvest - Journal of Universal Studies, 2026. DOI: 10.59188/eduvest.v6i7.53222.

15. V. S. Ramdoss, P. D. M. Rajan, "Strengthening Enterprise Systems: Failover Testing Techniques and Best Practices," International Journal of Science and Research Archive, 2025. DOI: 10.30574/ijsra.2025.15.3.1899.