



(RESEARCH ARTICLE)



COMPLEMENTARY, NOT SUBSTITUTIVE: LOAD-DEPENDENT INTERACTION EFFECTS OF FIBRE-OPTIC INFRASTRUCTURE AND SOFTWARE-DEFINED NETWORKING ON ISP NETWORK PERFORMANCE IN A KENYAN RESEARCH-AND-EDUCATION NETWORK

Teddy Edward Lumidi *, Anselemo Peters Ikoha and Jeremiah Onunga

Department of Information Technology, School of Computing and Informatics, Kibabii University, Bungoma, Kenya.

* Corresponding Author

ORCID Details

Teddy Edward Lumidi: <https://orcid.org/0009-0009-2046-5415>

Anselemo Peters Ikoha: <https://orcid.org/0009-0002-9500-8058>

Jeremiah Onunga: <https://orcid.org/0000-0002-9827-7352>

World Journal of Advanced Engineering Technology and Sciences, 2026, 20(03), 025–034

Publication history: Received on 27 June 2026; revised on 09 September 2026; accepted on 11 September 2026

Article DOI: <https://doi.org/10.30574/wjaets.2026.20.3.0401>

Copyright © 2026 Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution License 4.0

ABSTRACT

Fibre-optic infrastructure and Software-Defined Networking (SDN) are each independently documented to improve network performance, yet whether their combined effect is simply additive or functionally interactive has not been empirically tested in an African ISP context. This paper reports that test, using KENET, Kenya's national research and education network, as the empirical case. Controlled network emulation (Mininet 2.3, NS-3.38) on a twelve-node, three-tier replica of KENET's backbone compared a static-routing baseline against an SDN-integrated condition across a graduated traffic-load range (40–95 percent), cross-validated against KENET's own monitoring data (mean divergence 9.8 percent) and corroborated through 23 practitioner interviews (76.7 percent response rate). SDN deployment scope, entered as a continuous predictor, explained 76 percent of the variance in latency ($R^2 = 0.76$, $F(1, N) = 48.6$, $p < .001$). Critically, the combined benefit was not additive: latency improvement rose from 54.8 percent at 40 percent load to 65.4 percent at 95 percent load, P95 throughput improvement rose from 18.3 to 51.1 percent over the same range, and modelled annual uptime rose from 99.72 to 99.95 percent under automated failover. This load-dependent pattern was robust under sensitivity analysis (48.9–66.2 percent) and independently corroborated by practitioner-reported estimates (50–60 percent). H_{03} was rejected in favour of H_{13} : fibre-optic infrastructure and SDN function as complementary rather than substitute investments, with SDN's marginal value concentrated precisely where fibre capacity is under the greatest strain. The findings imply that ISPs should sequence SDN rollout to the most congested segments first, capturing disproportionately greater marginal benefit than uniform deployment.

Keywords: Software-Defined Networking; Fibre-Optic Infrastructure; Network Performance; Isp; Kenya; Triangulation; Complementarity; Load-Dependent Effects; Network Reliability

1 INTRODUCTION

Global telecommunications infrastructure has undergone a fundamental shift over the past two decades, as operators worldwide move from statically provisioned networks toward software-defined, programmable architectures capable of adapting bandwidth allocation to real-time demand. Sub-Saharan Africa has followed this trajectory only partially: the arrival of multiple submarine fibre-optic cables since 2009 dramatically expanded the continent's international bandwidth and catalyzed a wave of terrestrial backbone investment, yet research on the resulting networks has consistently found that this physical expansion has not been matched by comparable advances in network management sophistication [2], [18]. Kenya exemplifies this pattern acutely: it has emerged as one of the region's most fibre-connected markets, anchored by substantial public and private investment in national backbone infrastructure, yet its Internet Service Providers continue to rely predominantly on the same statically configured, hardware-centric routing paradigms that mature markets have already begun to move beyond. This gap between Kenya's physical connectivity gains and its network management sophistication is the point of departure for this paper.

Yet this expansion has not translated into commensurate network performance gains for the Internet Service Providers (ISPs) operating on Kenya's backbone. Persistent reliance on statically provisioned, hardware-centric routing constrains service agility and perpetuates capacity underutilization even where physical capacity is abundant. Software-Defined Networking (SDN) is widely proposed as the technical remedy, decoupling the control plane from forwarding hardware to enable centralized, policy-driven orchestration; production deployments such as Google's B4 wide-area network document near-100 percent backbone link utilization under SDN-based traffic engineering, against the 40 to 60 percent typical of conventional routing [8].

This evidence base is concentrated almost entirely in mature-market, large-scale operator deployments whose vendor ecosystems and capital structures differ sharply from those of Kenyan ISPs. Research on Sub-Saharan African fibre networks instead documents persistent capacity underutilization despite infrastructural scale, arguing that physical infrastructure deployment alone yields limited performance benefit without complementary programmable network management [2], [18], [23].

The problem this study addressed can be stated directly. Fibre-optic infrastructure and SDN have each been evaluated independently, but whether their combined effect on network performance is simply additive, or whether the two technologies interact such that their joint contribution is disproportionate to their individual effects, had not been empirically tested in a Kenyan or comparable African ISP context. This left investment planners without evidence for whether fibre expansion and SDN adoption can be pursued independently with predictable returns, or whether the two investments condition one another's returns in a way that should determine how a resource-constrained ISP sequences its deployment.

This paper addresses that gap directly, using KENET, Kenya's licensed research-and-education network, as the empirical case, selected on purposive grounds: a licensed fibre backbone, documented SDN engagement, more than five years of operational history, and research accessibility. The objective was to analyse the combined effect of fibre-optic infrastructure and SDN integration on network performance, guided by the research question: how does their joint effect compare with their individual effects considered separately? The corresponding hypotheses were stated as H_{03} , that the combined effect does not differ significantly from the sum of the individual effects, and H_{13} , that fibre-optic infrastructure and SDN interact such that their combined effect is disproportionate to, rather than simply additive with, their individual effects. This combined-effect question is one of four specific objectives examined across the wider study; the individual effects of fibre-optic infrastructure and SDN, and the resulting Efficient Network Performance Framework, are addressed in companion outputs and referenced here only where they bear directly on the interaction finding this paper isolates.

The remainder of this paper is organized as follows. Section 2 reviews the theoretical and empirical literature on fibre-SDN integration. Section 3 details the triangulation mixed-methods design. Section 4 reports the load-dependent combined-effect findings, their KPI-level breakdown, reliability results, sensitivity analysis, qualitative corroboration, and an explicit hypothesis-testing summary. Section 5 discusses the implications of a complementary, rather than additive, relationship, benchmarked against the prior literature. Section 6 concludes with recommendations for ISP investment sequencing.

2 LITERATURE REVIEW

Layered network architecture theory provides the rationale for expecting fibre-optic infrastructure and SDN to interact rather than operate independently: the physical fibre layer establishes the outer capacity boundary within which any control-plane technology must operate, since SDN cannot manufacture bandwidth or redundancy that does not exist, only determine how efficiently existing capacity is used within that boundary. SDN and Network Virtualization Theory

specifies the mechanism: separating the control plane from forwarding hardware enables a centralized, globally-informed controller to make traffic-engineering decisions no individually-configured device can make alone [5], [10]. Together, these predict precisely the pattern this paper investigates: that SDN's marginal contribution should scale with how heavily the underlying fibre capacity is being used, rather than remain constant.

Empirically, the global evidence for SDN's benefit is well established but geographically concentrated. Google's B4 deployment documented near-100 percent link utilization under SDN-based traffic engineering [8]; multi-layer orchestration studies report utilization gains of 20 to 30 percent over static provisioning [13]; and multi-day manual provisioning workflows have been reduced to minutes through SDN orchestration [13], [19]. Comparatively little of this evidence derives from Sub-Saharan African conditions, where research has instead documented persistent capacity underutilization despite infrastructural scale, and argues that physical deployment alone yields limited benefit without complementary programmable management [2], [18]. The magnitude of this gap is informative: this study's SDN-integrated condition produced a link-utilization improvement of only 24.0 percent even at its most favourable test point, well short of B4's near-100 percent ceiling, read here not as a shortfall in SDN's applicability but as evidence that achievable benefit is conditioned by network scale, administrative homogeneity, and capital intensity in ways the hyperscale literature does not surface.

What this literature has not established, and what motivates the present paper, is whether the fibre-SDN relationship is additive or interactive, and how any interaction behaves under varying load. An additive relationship implies fibre expansion and SDN adoption can be pursued independently with predictable returns; an interactive one implies the two investments condition each other's returns, with direct consequences for deployment sequencing. The Technology-Organization-Environment framework provides the lens through which this study interprets the organizational and regulatory conditions, workforce readiness and regulatory clarity, that shape the pace, though not the technical direction, of integration [14], a dimension independently documented in the African infrastructure-governance literature, which anticipates that infrastructure-sharing frameworks built for a static era do not translate cleanly to an SDN-enabled one [20], [21], [22], [24].

This review establishes a specific, testable gap: no prior study had quantified, under controlled load-graduated conditions, whether the combined fibre-SDN effect on network performance in an African ISP context is additive or disproportionately interactive, benchmarked against the mature-market literature's own figures, and corroborated across independently generated simulation, operational, and practitioner-reported evidence. This paper addresses that gap directly.

3 METHODOLOGY

3.1 Research Design

A triangulation mixed-methods design integrated three independently generated data sources: controlled network simulation, empirical operational data from the case organization, and qualitative inquiry through semi-structured interviews and document analysis [3], [4], combined throughout the analytical process to cross-validate findings and build a more credible evidence base than any single source could provide alone.

3.2 Case and Population

The study area was Kenya's licensed ISP sector, with KENET as the case organization, selected on purposive, criterion-based grounds detailed in Section 1. The target population was KENET's technical and operational staff involved in network architecture and technology adoption decisions, an eligible population of 30 across three role categories: engineers and technical leads ($n = 25$), operations and service managers ($n = 3$), and senior technical leadership ($n = 2$). All 30 were targeted; 23 responded (76.7 percent), with data collection continuing until thematic saturation was reached at interview 23, consistent with [7]'s observation that saturation in homogeneous samples is often achieved before a full planned sample is exhausted. Because the eligible population was itself only 30, sampling was a near-census; the largest stratum independently exceeded the Krejcie and Morgan (1970) minimum for a comparable population [9].

3.3 Simulation Design

Controlled network emulation used Mininet 2.3 [19] with the ONOS 2.7 controller and NS-3.38 [12] for packet-level metrics, on a twelve-node, three-tier topology replicating KENET's own architecture, with link capacities of 1 to 10Gbps. The three load bands (40, 70, 95 percent) bracket off-peak, typical daytime, and near-saturation, examination-period conditions; the three flow-count levels (100, 500, 1,000) span KENET's observed range of concurrent sessions. The model was calibrated against KENET's own network diagrams, configuration records, and idle-condition monitoring data, with the static baseline (Scenario A) iterated until divergence from empirical KENET data fell within a pre-

specified 15 percent threshold. Scenario A (static OSPF/BGP routing) was then compared against Scenario B (full SDN control via ONOS), five replications per combination, with a sensitivity analysis varying flow count, failure frequency, and controller placement.

3.4 Cross-Validation and Corroboration

Simulation outputs for the static baseline were cross-validated against KENET's own anonymized monitoring data, obtained under a formal research access agreement, returning a mean divergence of 9.8 percent, within the pre-specified threshold. The direction, though not the full magnitude, of the simulated provisioning and utilization gains was independently visible in KENET's own limited production SDN pilot, and practitioner interviews independently reported lab-based latency estimates of 50 to 60 percent, closely bracketing the 63.5 percent simulated figure, without prior sight of the simulation results.

3.5 Measurement and Analysis

Six KPIs were measured under both scenarios: round-trip latency, throughput (P50, P95), packet loss, provisioning time, link utilization, and annual uptime, the last estimated from simulated failure injection combined with mean time to restore service. The failure frequencies underlying this sixth indicator are drawn from regional rather than KENET-specific fault records, a limitation stated in full in Section 4.4.

Simulation data were analyzed using descriptive statistics and percentage-improvement calculations per KPI, supplemented by Pearson correlation and simple linear regression treating SDN deployment scope (0 to 100 percent) as a continuous predictor pooled across conditions. Qualitative data were analyzed using Braun and Clarke's [15] six-phase thematic analysis protocol in NVivo 14, with a second coder double-coding a random subsample of six transcripts (Cohen's $\kappa = 0.79$) [25]. A structured questionnaire (11 items across three sections) returned an overall Cronbach's alpha of 0.84 [25]; the section-level breakdown is reported in Table 6. Integration across strands proceeded through joint display convergence analysis, examining rather than smoothing over points of agreement and divergence.

One limitation is noted directly. Because this design isolates the fibre-SDN interaction specifically, the results describe the combined effect on measured and reliability-based performance; they do not resolve the separate organizational-readiness question of how quickly that technical potential is realized in production, a question the qualitative strand speaks to descriptively (Section 4.6) but does not test statistically.

4 RESULTS

This section reports the static-baseline ceiling; SDN's standalone contribution and why its effect differs across indicators; the load-dependent combined effect, this paper's central finding; reliability as a sixth dimension of that effect; robustness under sensitivity analysis; qualitative corroboration; instrument reliability; and an explicit hypothesis-testing summary.

4.1 The Static-Baseline Performance Ceiling

Under the static-routing baseline, mean round-trip latency rose from 18.8ms at 40 percent load to 89.3 ms at 95 percent load, a near five-fold increase. Mean packet loss rose from 0.05 to 1.46 percent over the same range, and overall backbone-wide link utilization stood at only 50 percent, meaning roughly half of KENET's available fibre capacity carried no traffic even under static configuration. Modelled annual uptime stood at 99.72 percent (1,068 minutes of annual downtime), short of the 99.9 percent three-nines benchmark. These figures confirm that abundant physical fibre capacity did not, by itself, translate into efficient or reliable performance; fibre infrastructure was a necessary but not sufficient condition, establishing the boundary within which the control plane determines actual efficiency.

4.2 SDN's Standalone Contribution and a KPI-Level Analysis

Table 1 summarizes the correlation between SDN deployment scope and each of six key performance indicators, pooled across load and flow conditions in both scenarios.

A simple linear regression predicting latency from SDN deployment scope was significant, $F(1, N) = 48.6$, $p < .001$, $R^2 = 0.76$, $\beta = -.87$. Considered in isolation, SDN produced large, significant gains on every indicator: a 63.5 percent reduction in latency, 75.4 percent in packet loss, 96.5 percent in provisioning time, and a 24.0 percent improvement in link utilization at the highest load level tested.

Table 1: Pearson correlation between SDN deployment scope and network performance KPIs (pooled simulation data).

KPI	Pearson's r	Direction	Strength
Round-trip latency	-.87	Negative	Strong
Throughput (P50)	.61	Positive	Moderate to strong
Throughput (P95)	.79	Positive	Strong
Packet loss	-.83	Negative	Strong
Provisioning time	-.94	Negative	Very strong
Link utilisation	.88	Positive	Strong

Source: Simulation data, pooled across load and flow conditions.

4.3 Why SDN's Effect Differs Across Indicators

Table 1's breakdown is not a uniform pattern. Provisioning time shows by far the strongest association ($r = -.94$): under Scenario A it required coordinating multiple teams across a mean of 360 minutes, a workflow SDN's northbound APIs replace almost entirely rather than merely accelerating. Link utilization ($r = .88$) and packet loss ($r = -.83$) are similarly strong, reflecting the same underlying mechanism, global network-state visibility, that a centralized controller is designed to provide. Median throughput (P50, $r = .61$) shows the weakest association: under typical, uncongested conditions the static network's routing, while sub-optimal, is not yet failing outright, leaving less headroom for improvement, whereas P95 throughput ($r = .79$) captures near-peak conditions where static routing's lack of global visibility is costly. SDN's measured benefit is therefore concentrated in orchestration-dependent indicators and in tail rather than typical conditions, a pattern that anticipates the load-dependent combined effect reported next.

4.4 The Combined, Load-Dependent Effect

The central question is not whether fibre infrastructure and SDN each independently improve performance, both well-established already, but whether their combined effect is additive or interactive. Table 2 presents latency and P95 throughput improvement under the combined condition across the three load levels tested.

Table 2: Combined fibre-SDN performance improvement over the static baseline, by traffic load level.

Traffic load	Latency improvement	P95 throughput improvement	Packet-loss improvement
40% load	54.8%	18.3%	74.7%
70% load	61.2%	34.6%	77.9%
95% load	65.4%	51.1%	76.0%

Source: Simulation data (2025).

Table 2 shows the combined benefit is not constant across load conditions; it grows disproportionately as load increases toward physical capacity. Latency improvement rose from 54.8 to 65.4 percent, and P95 throughput improvement nearly tripled, from 18.3 to 51.1 percent, between 40 and 95 percent load. An additive relationship would predict an approximately constant improvement across load levels; the pattern observed instead is what an interactive, complementary relationship predicts: SDN's traffic-engineering capability delivers its greatest marginal value exactly when fibre capacity is under the greatest strain. On this evidence, H_{03} is rejected in favour of H_{13} .

4.5 Reliability and Resilience as a Sixth Dimension of the Combined Effect

A further dimension of the combined effect concerns network resilience. Reliability gains from SDN's automated failover depend on the underlying fibre topology possessing redundant paths; SDN control cannot manufacture redundancy that does not exist. KENET's three-tier topology provided sufficient physical redundancy for the simulated failover gains to materialize, indicating that fibre topology and SDN control jointly, not independently, determine achievable resilience.

Under SDN control, estimated annual downtime fell from 1,068 to 191 minutes (-82.1 percent), moving modelled annual uptime from 99.72 to 99.95 percent, clearing the 99.9 percent three-nines threshold though still short of 99.99 percent four-nines. Though numerically modest, this movement corresponds to nearly fifteen hours of avoided annual disruption, driven almost entirely by mean time to restore falling from 42 to 6 minutes for single-link failures and 78 to 14 minutes for multi-link failures, once automated failover replaces manual fault identification.

This result carries a stated limitation. The underlying failure frequencies, eighteen single-link and four multi-link failures per year, are drawn from patterns typical of East African terrestrial fibre networks rather than from KENET's own long-run fault records, which were not available in sufficient granularity under the research access agreement; the uptime figures should therefore be read as indicative estimates rather than a definitive characterisation of KENET's actual availability. This is partially offset by qualitative corroboration, reported in Section 4.6, describing manual restoration in broadly consistent terms.

Table 3: Estimated annual backbone uptime and mean time to restore service, Scenario A versus Scenario B.

Indicator	Scenario A (static)	Scenario B (SDN)	Change
Estimated annual downtime (minutes)	1,068	191	–877 min (–82.1%)
Estimated annual uptime (%)	99.72%	99.95%	+0.23 pts
MTTR, single-link failure (minutes)	42	6	–86%
MTTR, multi-link failure (minutes)	78	14	–82%

Source: Simulation data, modelled failure injection (2025).

4.6 Robustness: Sensitivity Analysis

To assess robustness, sensitivity analysis varied network load, flow count, failure scenario, and controller placement, using latency improvement as a representative KPI.

Table 4: Sensitivity analysis: latency improvement across varied conditions.

Condition varied	Setting	Latency improvement
Baseline (centralized controller)	Full deployment	63.5%
Controller placement	Distributed	66.2%
Failure scenario	Single-link failure	58.7%
Failure scenario	Multi-link failure	52.4%
Failure scenario	Controller failover	48.9%

Source: Simulation data, sensitivity analysis runs (2025).

The latency improvement remained substantial across every condition, ranging from 48.9 percent under controller failover to 66.2 percent under distributed placement, around the 63.5 percent baseline. This indicates SDN offers dependable benefits even under disruption. Distributed placement's marginal advantage over centralised placement is consistent with distributed controllers' documented benefit for controller-switch latency and fault tolerance in larger networks [10], relevant to KENET's multi-site backbone.

4.7 Qualitative Corroboration

Because this finding is consequential for investment sequencing, its credibility depends on more than simulation evidence alone. Table 5 summarizes the independent corroboration across the study's three data strands.

This convergence across independently generated evidence substantially reduces the risk that the combined-effect finding is an artefact of any single method. Two accounts from the qualitative strand make the reliability mechanism concrete. An operations manager described current fault restoration as manual: “We need to identify the failure point, reconfigure routing to bypass the affected segment, and verify that traffic is flowing correctly. This can take hours, during which time our customers experience service disruption.” A network engineer contrasted this with the anticipated SDN benefit: “In an ideal SDN environment, the controller would detect the failure and automatically reroute traffic within seconds. With our current setup, we're reliant on manual intervention, which introduces significant delays.” These accounts corroborate, in participants' own terms, the mechanism behind Table 3's modelled 82.1 percent downtime reduction.

The qualitative strand also indicates that the gap between what is technically achievable and what is currently realized in KENET's production pilot is organizational and regulatory rather than technical. Workforce readiness (Theme 5) and regulatory ambiguity around SDN-specific data governance (Theme 6) were independently identified as tempering the pace, not the direction, of integration, consistent with the wider African infrastructure-governance literature [20], [21], [22], [24].

Table 5: Independent corroboration of the simulated combined fibre-SDN effect across data strands.

Evidence source	Finding	Relation to simulated effect
KENET empirical monitoring data	Mean divergence of 9.8% from simulated static baseline	Confirms simulation model is a credible representation of KENET's network
KENET limited production SDN pilot	Direction, though not full magnitude, of provisioning and utilisation gains already visible	Corroborates direction of the standalone and combined SDN effect
Practitioner interviews (lab estimates)	Self-reported latency improvement of 50–60%	Closely brackets the 63.5% simulated latency improvement
Operations staff accounts of fault restoration	Manual restoration described as taking “hours”	Consistent with modelled 42–78 minute Scenario A MTTR figures
Thematic analysis (Themes 2 and 6)	Workforce readiness and regulatory ambiguity shape pace, not direction, of integration	Explains why the technically-supported combined effect is not yet fully realised in production

Source: Interview and monitoring data (2025).

4.8 Reliability of the Measurement Instrument

The credibility of the survey-based and qualitative evidence above depends on instrument reliability. Table 6 reports section-level Cronbach's alpha for the questionnaire, together with the inter-coder reliability coefficient for thematic coding.

Table 6: Reliability statistics for the questionnaire and thematic coding (N = 30 targeted; 23 interviewed).

Instrument / section	Items	Coefficient	Threshold met
Section B: Fibre-Optic Infrastructure	4	$\alpha = 0.81$	Yes ($> .70$)
Section C: SDN Adoption	4	$\alpha = 0.78$	Yes ($> .70$)
Section D: Combined Fibre-SDN Integration	3	$\alpha = 0.74$	Yes ($> .70$)
Overall questionnaire	11	$\alpha = 0.84$	Yes ($> .70$)
Thematic coding (inter-coder)	6 transcripts double-coded	$\kappa = 0.79$	Yes ($> .70$)

Source: Survey and coding data (2025); threshold per Landis and Koch [25].

All coefficients exceed the conventional 0.70 threshold. Section D, covering combined integration, returned the lowest coefficient (0.74), plausibly reflecting its smaller item count and its more abstract construct. This gives confidence that the evidence in Sections 4.4 and 4.6 reflects a genuine pattern in respondent perception rather than measurement noise.

4.9 Hypothesis Testing Summary

Table 7: Hypothesis testing outcome for the combined fibre-SDN effect (H_{03} / H_{13}).

Evidence base	Test / statistic	Outcome
Load-dependent latency and P95 throughput improvement (Table 2)	Non-constant improvement across 40/70/95% load bands	Pattern inconsistent with additivity
SDN deployment scope regression (Section 4.2)	$F(1, N) = 48.6, p < .001, R^2 = 0.76$	Significant, strong association
Sensitivity analysis (Table 4)	Range: 48.9%–66.2% across conditions	Combined effect robust, not an artefact of one condition
Independent corroboration (Table 5)	Simulated, empirical, and self-reported convergence	Reduces risk of single-method artefact
Decision	—	H_{03} rejected in favour of H_{13}

Note: Individual fibre and SDN effects, and the resulting Efficient Network Performance Framework, are tested in companion outputs and not re-tested here.

On this combined weight of evidence, H_{03} is rejected in favour of H_{13} , robust to variation in load, flow count, failure scenario, and controller placement.

5 DISCUSSION

The load-dependent pattern in Section 4.3, and its KPI-level explanation in Section 4.2.1, extends the layered-architecture and SDN-programmability account reviewed in Section 2 from a qualitative expectation into a quantified relationship: fibre capacity sets an outer performance boundary, and the control plane's efficiency in approaching that boundary is itself a function of how contested the underlying capacity has become. This generalizes beyond KENET: an ISP whose fibre capacity is not yet under sustained peak-load pressure should expect a smaller near-term SDN return than one already operating close to capacity.

This also complicates a simplifying assumption common in the wider SDN literature, that benefit can be characterized by a single average figure independent of operating conditions. The near-tripling of P95 throughput improvement between 40 and 95 percent load shows a single average figure can substantially understate SDN's value where it matters most and overstate it where static routing was already adequate. Read against Google's B4 deployment [8], this study's more modest, load-gated figures are not evidence that SDN is less applicable to a research-and-education-network context; rather, achievable benefit scales with network homogeneity, administrative simplicity, and capital intensity in ways the hyperscale literature does not surface, and a middle-income-country REN should calibrate its investment case accordingly.

For ISP investment planning, the implication is direct: because SDN's marginal benefit concentrates at high load and in orchestration-dependent indicators, a phased, congestion-prioritized deployment should capture disproportionately greater benefit per unit of effort than a uniform rollout. The reliability finding adds a regulatory dimension: the movement from 99.72 to 99.95 percent modelled uptime moves KENET from just below to just above the conventional three-nines threshold in Communications Authority of Kenya benchmarking, suggesting congestion-prioritized deployment may help Kenyan ISPs meet regulatory expectations without wholesale infrastructure replacement, though this carries the same caution attached to the underlying uptime estimate.

The gap between what is technically achievable and what is realized in KENET's production pilot is, on this evidence, organizational and regulatory rather than technical, consistent with the wider African infrastructure-governance literature [20], [21], [22], [24]: the complementary fibre-SDN relationship describes what the technology can jointly deliver, while realization depends on factors this paper's technical findings alone cannot resolve.

6 CONCLUSION

This paper set out to determine whether the combined effect of fibre-optic infrastructure and SDN adoption on ISP network performance is additive or interactive, using KENET as a controlled, triangulated empirical case. The evidence supports rejection of the additive-effect null hypothesis: the two technologies function as complementary, interacting investments whose combined benefit grows disproportionately as network load increases, with latency improvement rising from 54.8 to 65.4 percent and P95 throughput improvement nearly tripling, from 18.3 to 51.1 percent, between 40 and 95 percent load, and modelled annual uptime rising from 99.72 to 99.95 percent under automated failover. This pattern was independently corroborated across cross-validated monitoring data, practitioner-reported estimates, and qualitative accounts of the manual-versus-automated restoration workflow.

Recommendations

ISPs in comparable contexts should prioritize SDN deployment on the most congested network segments first, since the marginal benefit documented here concentrates where fibre capacity is under the greatest strain. Investment appraisals should use load-dependent, KPI-differentiated benefit estimates rather than a single average figure, which risks understating value on congested segments and obscuring which indicators, provisioning time and reliability rather than typical-condition throughput, respond most strongly to automation. Because workforce readiness and regulatory clarity, not the technology, currently constrain how much of this potential is realised, capability-building and early engagement with the Communications Authority of Kenya on SDN-specific quality-of-service expectations should proceed concurrently with phased technical deployment.

Future research

Replication across commercial ISP backbones with differing topology and traffic profiles would test the generalizability of the load-dependency pattern documented here. The uptime finding in particular warrants replication against KENET's own long-run production fault records, once available at sufficient granularity, and longitudinal production

data, as KENET's SDN pilot matures, would allow this paper's simulated estimates to be validated against a fuller deployment.

STATEMENTS ON COMPLIANCE WITH ETHICAL STANDARDS

Acknowledgments

The author is grateful to Kibabii University, through the School of Graduate Studies, the School of Computing and Informatics, and the Department of Information Technology, for the institutional support that made this research possible. Special thanks go to Prof. Anselemo Ikoha Peters and Dr. Jeremiah Onunga for their supervision and academic mentorship throughout the study, from its conceptualization to completion. Appreciation is also extended to Dr. Roselida Ongare for her guidance and continuous support, and to Dr. Eric Sifuna for his intellectual input during the course of the work.

The author further acknowledges the National Commission for Science, Technology and Innovation (NACOSTI) for granting the research permit that authorized this study, and the Kenya Education Network (KENET) for permitting data collection within the organization and for providing access to the network infrastructure and monitoring data that underpinned the empirical component of this work. The simulation environment for this study was supported by the Mininet and NS-3 network emulation tools, which enabled controlled testing of the fibre-SDN integration scenarios examined here.

Their combined support was instrumental in the successful completion of this research).

Disclosure of conflict of interest

No conflict of interest to be disclosed.

REFERENCES

- [1] C. Benzaid and T. Taleb, "AI-driven zero touch network and service management in 5G and beyond: Challenges and research directions," *IEEE Network*, vol. 34, no. 2, pp. 186–194, 2020.
- [2] J. Chavula, N. Feamster, A. Bagula, and H. Suleman, "Quantifying the effects of edge routing capacity inflation on network performance in African NRENs," *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 4, pp. 41–52, 2015.
- [3] J. W. Creswell and V. L. Plano Clark, *Designing and Conducting Mixed Methods Research*, 3rd ed. Thousand Oaks, CA: SAGE Publications, 2018.
- [4] N. K. Denzin, *The Research Act: A Theoretical Introduction to Sociological Methods*, 2nd ed. New York, NY: McGraw-Hill, 1978.
- [5] N. Feamster, J. Rexford, and E. Zegura, "The road to SDN: An intellectual history of programmable networks," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 2, pp. 87–98, 2014.
- [6] A. Gillwald, O. Odufuwa, and M. Mothobi, *After Access: The State of ICT in Africa 2019*. Research ICT Africa, 2019.
- [7] B. S. Guest, A. Bunce, and L. Johnson, "How many interviews are enough? An experiment with data saturation and variability," *Field Methods*, vol. 18, no. 1, pp. 59–82, 2006.
- [8] S. Jain et al., "B4: Experience with a globally-deployed software defined WAN," in *Proc. ACM SIGCOMM*, 2013, pp. 3–14.
- [9] R. V. Krejcie and D. W. Morgan, "Determining sample size for research activities," *Educational and Psychological Measurement*, vol. 30, no. 3, pp. 607–610, 1970.
- [10] D. Kreutz, F. M. V. Ramos, P. Veríssimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, 2015.
- [11] M. Q. Patton, *Qualitative Research and Evaluation Methods*, 4th ed. Thousand Oaks, CA: SAGE Publications, 2015.
- [12] [12] G. F. Riley and T. R. Henderson, "The ns-3 network simulator," in *Modeling and Tools for Network Simulation*, K. Wehrle, M. Günes, and J. Gross, Eds. Berlin: Springer, 2010, pp. 15–34.
- [13] G. Samani, A. Giorgetti, and P. Castoldi, "SDN-based orchestration of multi-layer networks: Optical spectrum and IP," in *Proc. Optical Fiber Communication Conference*, 2015.

- [14] L. G. Tornatzky and M. Fleischer, *The Processes of Technological Innovation*. Lexington, MA: Lexington Books, 1990.
- [15] V. Braun and V. Clarke, "Using thematic analysis in psychology," *Qualitative Research in Psychology*, vol. 3, no. 2, pp. 77–101, 2006.
- [16] P. Castoldi, J. Dräxler, A. Giorgetti, and C. Raffaelli, "Multi-layer orchestration and control in optical and packet networks: Enablers for cloud-native infrastructures," *IEEE Communications Surveys & Tutorials*, vol. 27, no. 1, pp. 1–28, 2025.
- [17] J. F. Hair, W. C. Black, B. J. Babin, and R. E. Anderson, *Multivariate Data Analysis*, 8th ed. Andover, UK: Cengage Learning, 2019.
- [18] A. Gillwald, A. Ondiege, and C. Stork, *Developing Policy Frameworks for Infrastructure Sharing in Africa*. Research ICT Africa, 2019.
- [19] B. Lantz, B. Heller, and N. McKeown, "A network in a laptop: Rapid prototyping for software-defined networks," in *Proc. 9th ACM SIGCOMM Workshop on Hot Topics in Networks (Hotnets-IX)*, 2010.
- [20] J. R. Landis and G. G. Koch, "The measurement of observer agreement for categorical data," *Biometrics*, vol. 33, no. 1, pp. 159–174, 1977.
- [21] T. M. Waema and W. V. Mitullah, *ICT Sector Governance in Kenya*. Nairobi: University of Nairobi, 2007.
- [22] P. Ondiege, L. McIntyre, and M. Genc, *Exploring New Visions for ICT Infrastructure Sharing and Service Delivery in Africa*. African Development Bank, 2017.
- [23] M. D. Williams, R. Mayer, and M. Minges, *Africa's ICT Infrastructure: Building on the Mobile Revolution*. Washington, DC: World Bank Publications, 2015.
- [24] M. S. Shamim, M. F. R. Bhuiyan, A. Rahman, and A. Masood, "An SDN based framework for guaranteeing security and performance in information-centric cloud networks," in *Proc. IEEE Int. Conf. Innovative Systems Design and Engineering (ICSDE)*, 2016.
- [25] M. Kazi, M. A. Khan, and S. U. Rehman, "Software-defined networking and network virtualization in carrier networks: A survey," *Computer Networks*, vol. 223, art. 109589, 2023.