

Keeping the Class Online: A Thematic Review of Network Performance Optimization Trends for Reliable E-Learning, 2022 to 2026

Joshua C. Cervantes

2025-00739-CM-0 · Section MSIT_A

Master of Science in Information Technology

Polytechnic University of the Philippines, Open University System / CCIS

Manila, Philippines

joshuacruz cervantes@gmail.com

ABSTRACT

Synchronous online learning has become a permanent fixture of higher education, yet its delivery depends on residential and campus networks that were never engineered for sustained real-time media. This paper reviews fifteen peer-reviewed studies published between 2022 and 2026 to establish how contemporary operating system and network technologies address the reliability gap in e-learning. The review is organized thematically rather than by author, across six themes: measurement of quality-of-service and quality-of-experience degradation in campus and distance networks; software-defined networking for quality-of-service and bandwidth management; edge computing for latency reduction; adaptive bitrate and quality-of-experience-driven video delivery; security as a performance determinant; and the digital divide as a constraint on all of the above. Three findings emerge. First, the literature has shifted decisively from static provisioning toward learned, predictive control, with machine learning now embedded in bitrate selection, flow classification, server load balancing, and endpoint steering. Second, prioritization consistently outperforms capacity expansion, because the flow that fails in an online class is low-bandwidth and latency-sensitive rather than bandwidth-hungry. Third, a measurement gap persists: the optimization literature is validated in simulation and controlled testbeds, while the education literature measures learner outcomes, and almost nothing connects a specific network intervention to a specific learning outcome in a low-resource setting. The paper closes by mapping these trends onto a diagnosed Philippine household failure case and recommending an intervention order for school network administrators.

CCS CONCEPTS

- Networks → Network performance evaluation; Network measurement; Network management
- Applied computing → Distance learning; E-learning
- Information systems → Multimedia streaming

KEYWORDS

network performance; quality of service; quality of experience; online learning; e-learning; software-defined networking; edge computing; adaptive bitrate; latency; jitter; packet loss; digital divide

ACM Reference Format:

Joshua C. Cervantes. 2026. Keeping the Class Online: A Thematic Review of Network Performance Optimization Trends for Reliable E-Learning, 2022 to 2026. MSIT 643 Advanced Operating Systems and Networking, Final Requirement. Polytechnic University of the Philippines, Open University System, Manila, Philippines, 14 pages.

I. INTRODUCTION

A. Background of the Study

Operating systems and computer networks are the substrate on which all modern distributed computing rests. The operating system abstracts hardware, schedules contended resources, and mediates access to the network stack; the network moves the resulting traffic across administrative and physical boundaries. In cloud computing, distributed systems, and enterprise information technology infrastructure, these two layers determine whether an application is usable, and their behaviour under contention determines whether it is usable when it matters.

Online education has become one of the most demanding consumers of that substrate, and one of the least well provisioned. The displacement of instruction to synchronous video platforms during and after 2020 did not reverse; hybrid and fully online delivery are now standard in higher education, including in graduate programs delivered through open university systems. What changed permanently was not the pedagogy but the dependency: instruction now traverses consumer broadband, shared residential wireless, and campus networks whose capacity planning predates the requirement.

The mismatch is structural rather than incidental. Video conferencing imposes a continuous, bidirectional, latency-bounded load. A media frame that arrives after its playback deadline is not a slow frame but a discarded one. This distinguishes it sharply from the traffic these networks were provisioned for, where a delayed byte produces only a slower page. A household that comfortably supports several streaming devices can nonetheless fail to carry a single class session, because streaming buffers seconds ahead and video conferencing cannot buffer at all.

B. Statement of the Problem

The reliability of online learning is gated by network performance, specifically by latency, jitter, packet loss, and contended bandwidth, yet the infrastructure and the quality-of-service mechanisms available to schools and households lag the demand placed on them. The problem is sharpest in developing regions, where household bandwidth is constrained, typically shared among several simultaneous users, and delivered over wireless links whose quality varies with position inside a dwelling.

Two gaps motivate this review. The first is technical: a large and rapidly moving literature proposes optimization mechanisms, but these mechanisms are distributed across distinct research communities, namely software-defined networking, edge computing, multimedia streaming, and network security, and are rarely synthesized with respect to a single application domain. The second is evidential: the network optimization literature validates its contributions in simulation and controlled testbeds, while the education literature measures learner engagement and attainment. The two bodies of work seldom meet, and as a result a school administrator deciding where to spend a constrained budget has no evidence

base that connects a specific intervention to a specific learning outcome.

C. Objectives of the Study

1. To synthesize peer-reviewed research published between 2022 and 2026 on network performance optimization relevant to synchronous and asynchronous e-learning delivery.
2. To organize that research thematically and identify the trends, agreements, and contradictions within and across themes.
3. To map each theme onto a concrete, documented failure of an online learning session, in order to test whether the literature explains observed practice.
4. To identify the research gaps that remain, with particular attention to low-resource and developing-region contexts.
5. To derive a prioritized set of recommendations for school and household network administrators, and to indicate directions for future research.

D. Scope and Delimitation

This study is a review of related literature and is limited to secondary sources. It does not involve experimental implementation, and no new measurements were collected. Included works are peer-reviewed journal or conference publications, written in English, published between 2022 and 2026, and relevant to network or operating system performance in a learning or directly transferable multimedia delivery context. Non-peer-reviewed material, including preprints, patents, vendor documentation, and technical blogs, is excluded. Pre-2022 works are excluded except where cited as established background rather than as evidence of current trends.

The review addresses the network and operating system performance layer. It does not address instructional design, assessment validity, or learner psychology except where the literature connects those outcomes to a measured network condition. Where the paper applies the reviewed trends to a specific case, that case is a documented single-household diagnosis and is used illustratively rather than as an independent empirical contribution.

E. Significance of the Study

For information technology students and practitioners, the review consolidates a fragmented literature into a single trend map, and demonstrates the operating-system and

network concepts of scheduling, queueing, contention, and prioritization operating in a familiar setting rather than an abstract one. For school network administrators, it converts a research literature into an intervention order, which is the form in which the decision is actually faced. For institutions and policymakers, it identifies where the evidence is thinnest, namely the connection between infrastructure spending and learning outcome in constrained settings, which is precisely where public investment decisions in developing regions are made. For researchers, it names a specific and addressable gap between two mature literatures that currently do not cite each other.

II. REVIEW OF RELATED LITERATURE

Fifteen studies are reviewed, organized into six themes. Themes are defined by the mechanism or constraint the work addresses rather than by publication venue or method, and several studies contribute to more than one theme; cross-references are noted where they occur.

Theme A: Measuring the Problem, Quality of Service and Quality of Experience in Campus and Distance Networks

Before optimization can be evaluated, degradation must be measured. Three studies establish what is currently measurable and what those measurements show.

The most direct measurement of the campus case is a quantitative descriptive study of the campus internet network at Universitas Islam Sumatera Utara, which measured throughput, delay, packet loss, and jitter under peak and non-peak load (Prayogi & Hulu, 2025). The finding is instructive precisely because it is undramatic: performance degraded measurably during peak hours, with lower throughput and higher delay, packet loss, and jitter, yet every parameter remained within accepted quality-of-service thresholds, leading the authors to conclude the network was adequate for academic and online learning activity. This result exposes a limitation of threshold-based assessment. A network can satisfy aggregate quality-of-service thresholds while still failing a real-time session, because the thresholds were derived for general traffic and the binding constraint for conferencing is the tail of the jitter distribution rather than its mean.

A more comprehensive framing comes from work validating a quality-of-service framework specifically for open and distance electronic

learning (Amoako et al., 2026). The proposed QoSODEL model identifies security, network bandwidth management, and conflict management as the elements determining service quality, and validates the model through survey data from 551 participants, comprising 476 students, 53 lecturers, and 18 administrators, analysed by confirmatory factor analysis using the SERVQUAL instrument, in a Ghanaian context. Two aspects matter for this review. First, the model treats bandwidth management as a first-order determinant of educational service quality rather than a technical detail, which aligns the education literature with the networking literature. Second, it places security inside the quality-of-service construct rather than beside it, anticipating Theme E: the study identifies secure, cheating-free examination as a component of quality of service, meaning that a security mechanism which degrades performance degrades educational quality directly.

The broadest evidence on outcomes is a systematic review of the impact of online learning on student performance and engagement, conducted under PRISMA guidelines across PubMed, ScienceDirect, and JSTOR, with eighteen studies retained after screening (Akpen et al., 2024). Its relevance here is a single finding: among the factors determining whether online learning succeeds, the authors identify the quality of digital tools, good internet connectivity, and student motivation. Connectivity appears as a determinant of academic outcome, on equal footing with motivation and tool quality. This is the strongest available warrant for treating network performance as an educational variable rather than an infrastructural one, and it is notable that the review reaches this conclusion without citing any networking literature.

Synthesis and gap. The three studies agree that network quality determines educational quality, but they measure at incompatible granularities. The campus study measures packets and concludes the network is adequate; the framework study measures perceptions and concludes bandwidth management is decisive; the systematic review measures outcomes and concludes connectivity matters. None of them measures a network intervention and a learning outcome in the same study. The gap is a measurement bridge, and its absence is why threshold-satisfying networks continue to fail classes.

Theme B: Software-Defined Networking for Quality of Service and Bandwidth Management

If contention is the problem, control is the proposed answer, and software-defined networking is the dominant control paradigm in the reviewed period. Four studies mark the trajectory.

The clearest statement of the mechanism is work on detecting and predicting models for quality-of-service optimization in software-defined networks (Wassie et al., 2024). The authors apply deep learning to detect so-called elephant flows, meaning the small number of large, aggressive flows that consume disproportionate capacity, and a random forest algorithm to predict them, in order to control delay, latency, bandwidth, and packet loss and thereby allocate resources to real-time applications. This is directly the online-learning problem stated in networking terms: a bulk download is an elephant flow, a class session is not, and the intervention is to identify the former in order to protect the latter.

A second study addresses contention at the server rather than the link, proposing dynamic active sensing server load management for software-defined networks (Mehmood et al., 2023). The problem addressed is that when requests concentrate on one server, end-to-end user delay rises while transfer rate and available bandwidth fall. The proposed algorithm redistributes load using throughput, response time, and round-trip time as control parameters, reducing overall delay while raising bandwidth utilization and throughput. For a school operating its own learning management system, this identifies a failure mode distinct from access-link congestion and remediable by control rather than capacity.

The trend toward learned control continues in work optimizing quality-of-service metrics for software-defined networking in federated learning settings, which combines federated learning for load balancing with genetic algorithms, explicitly preserving data privacy while distributing network load (Fallah et al., 2023). The privacy dimension is significant for education, where learner data is regulated, and it demonstrates that optimization and data protection are not necessarily in tension.

Where the preceding three studies learn, the fourth returns to analytic queue management and is the more directly actionable for that reason. Traffic engineering queue optimization models with guaranteed quality-of-service

support introduce two schemes providing guarantees either for the bandwidth allocated to a class queue or for its utilization coefficient, supporting both differentiated and integrated services (Titarenko et al., 2025). The mechanism reported is that such guarantees stabilize and control average queue length, which in turn improves average delay and packet loss probability. This is the theoretical account of why prioritization works at all, and it is the same mechanism a consumer router exposes as a quality-of-service setting.

Synthesis and gap. The theme shows a clear movement from reactive to predictive control, and from hand-tuned policy to learned policy. All four studies validate in simulation, controlled testbed, or analytic model rather than in deployment. None evaluates deployment in a resource-constrained school network, where the controller itself is an additional cost and the operating expertise is scarce. The literature has established that control beats capacity; it has not established what the minimum viable control deployment is for an institution that cannot staff a network operations centre.

Theme C: Edge Computing and Latency Reduction

Where Theme B redistributes traffic, Theme C moves computation, on the premise that physical distance to a data centre sets a latency floor no scheduling discipline can lower.

A comparative review of machine-learning-based approaches for mobile edge computing evaluates several methods, including reinforcement-learning task offloading using proximal policy optimization, improved deep Q-learning with transfer learning, chaotic quantum particle swarm optimization for task allocation, personalized hierarchical federated learning, and convolutional-recurrent models for dynamic resource allocation, assessed on a large integrated fibre-to-wireless access network topology of one hundred edge nodes carrying mixed vehicular, augmented and virtual reality, and internet-of-things traffic (Marmat & Thankachan, 2025). The breadth is the contribution: it establishes that no single learned approach dominates, and that evaluation topology determines which appears best.

The mechanism transfers to education directly, since a conferencing service that terminates media at a nearby edge node rather than a distant region reduces round-trip time for every participant. A further candidate study

on fifth-generation edge computing was excluded from this review because its publication year could not be resolved unambiguously across indexes, and citing a source whose year is uncertain would violate the 2022 to 2026 delimitation this review claims to observe.

The most operationally interesting study in this theme sits at the boundary with Theme D. An adaptive video streaming architecture leveraging quality-of-experience forecasting and content steering on the edge-cloud continuum integrates a content steering service that selects a suitable endpoint dynamically during a session, continuous real-time monitoring of playback, network, and node state, and a deployed machine learning model that forecasts quality of experience over a prediction horizon in order to reconfigure the system before congestion occurs (Gama et al., 2026). The shift from reactive adaptation to forecast-driven pre-emption is the single clearest expression of where the field is moving, and content steering is notable for being a recent addition to the standardized adaptive streaming protocols rather than a research-only construct. The evaluation is also instructive for the present application: placing the cache service as close to users as possible performed best as a baseline and held service-level objective violations to a minimum up to roughly forty concurrent users, which is approximately the size of a single class section.

Synthesis and gap. Edge computing addresses a latency component that prioritization cannot, and the convergence with predictive control is evident. The gap is economic and geographic. Every study assumes edge infrastructure exists to be offloaded to. In the settings where online learning fails most often, it does not, and the literature does not address what edge placement is achievable where the nearest interconnect is already distant.

Theme D: Adaptive Bitrate and Quality-of-Experience-Driven Video Delivery

This theme addresses the mechanism that actually carries a lecture, and with four studies it is, alongside Theme B, one of the two best-populated themes in the review.

The most directly relevant study optimizes the bitrate of a video course using both network bandwidth and user bandwidth, applying a recurrent neural network with long short-term memory to adaptive bitrate selection within a collaborative e-learning platform, selecting

bitrate from network throughput, download time, current buffer level, previous bitrate, average bandwidth, rebuffering time, and next chunk size (Venkatesh Naik & Madhavi, 2024). Its significance for this review is that the application domain is not merely analogous to education but is education, and that the feature set enumerates precisely the state variables a conferencing client uses. The reported evaluation attains a quality-of-experience score of 0.975 with a buffer size of 45.654 and a selected bitrate of 90.766 bits per second, and is validated against a dataset drawn from an operating e-learning platform rather than a synthetic trace, which is uncommon in this literature and materially strengthens its transferability to the present problem.

Two studies establish the machine learning trajectory within bitrate adaptation. An architecture using convolutional and recurrent neural networks extracts features from a video stream and predicts optimal per-frame bitrate from network conditions and user preferences, reporting improved quality alongside reduced bandwidth consumption (Darwich & Bayoumi, 2024). Extending this, a deep reinforcement learning enhanced optimization algorithm addresses the sample efficiency, stability, and generalization weaknesses of prior reinforcement learning approaches, using proximal policy optimization with a dual clipping mechanism combined with long short-term memory and local attention (Zhang et al., 2025). Read together, the two mark a maturation: the field has moved past demonstrating that learning helps and is now addressing whether learned policies generalize beyond their training distribution, which is the question that determines field deployability.

A critical counterweight is provided by a study treating quality of experience as a decision criterion in heterogeneous multimedia network optimization, which surveys the benefits of quality-of-experience-driven optimization while identifying the technical obstacles, namely the difficulty of measuring quality of experience accurately and of managing heterogeneous network complexity, and identifies real-time quality-of-experience measurement and machine learning integration as open prospects (Omar et al., 2024). This is the theme's most useful contribution to a review, because it constrains the optimism of the other three: every mechanism above optimizes a proxy for quality of experience, and the field does not yet agree on how to measure the quantity being optimized.

Synthesis and gap. Adaptive delivery is the most mature area reviewed, and it converges strongly on learned, predictive control informed by buffer and bandwidth state. The measurement critique is the outstanding problem, and it compounds the Theme A gap: if quality of experience cannot be measured reliably in real time, it cannot be correlated with a learning outcome. Additionally, every study in this theme that proposes a delivery mechanism optimizes downstream video delivery. Synchronous instruction is bidirectional, and the upstream direction, which is where residential asymmetry binds, is not addressed.

Theme E: Security as a Performance Determinant

Security is conventionally treated as orthogonal to performance. The reviewed period does not support that separation.

A systematic literature review of the implementation and challenges of zero trust architecture across domains synthesizes seventy-four peer-reviewed articles published between 2016 and 2025, spanning cloud computing, the internet of things, healthcare, enterprise and remote work systems, industrial and supply chain networks, mobile networks, artificial intelligence and machine learning, blockchain, and big data and edge computing (Mushtaq et al., 2025). Two findings bear on this review. First, authentication, authorization, and access control are the most consistently implemented components, while auditing, orchestration, and environmental perception remain underexplored, meaning that deployments concentrate on the per-request verification path, which is exactly the path that adds latency. Second, the challenges identified include scalability limitations, insufficient lightweight cryptographic solutions for resource-constrained systems, and weak orchestration.

The implication for education is direct and underappreciated. Zero trust replaces perimeter trust with continuous per-request verification, so its cost is paid in added latency and added state on every access. In a well-provisioned enterprise this is negligible. On a constrained access link serving a real-time session, a verification round trip is drawn from the same latency budget as the media. The absence of lightweight cryptography for resource-constrained systems, identified by the review as a gap, is therefore a performance gap in the e-learning setting, not only a security one.

This connects back to Theme A, where secure examination was identified as a component of educational quality of service (Amoako et al., 2026): the two studies together imply that security and performance in education are a single joint optimization rather than two separable objectives.

Synthesis and gap. The literature establishes that security architecture has a performance cost and that the cost is worst where resources are scarcest. What it does not provide is any quantification of that cost for real-time educational traffic. No reviewed study measures the latency contribution of an access control decision inside a conferencing session.

Theme F: The Digital Divide as a Binding Constraint

The preceding five themes propose mechanisms. This theme establishes the conditions under which they must operate, and it is the theme where the developing-region evidence sits.

A qualitative narrative study of the effects of unstable internet connectivity on access to academic resources and collaborative learning interviewed ten tertiary students in Palimbang, a rural area of the Philippines, using in-depth open-ended interviews (Akmad & Abatayo, 2024). The reported effects are concrete: delayed task submission attributable to erratic connectivity, time-consuming access to online resources, reduced enthusiasm and reduced ability to participate, communication difficulty, and impaired timeliness of group work. Where connectivity was strong, students accessed resources faster, researched more efficiently, and submitted on time.

The methodological contrast with the rest of this review is the point. Every study in Themes B through D that proposes a mechanism reports an improvement in a network metric established by simulation, controlled testbed, or analytic model. This study reports what network failure does to a learner, in the specific setting where the failure is most common, and it does so without measuring a single packet. Neither literature can substitute for the other, and the fact that they do not cite each other is the central finding of this review.

Read against Theme A, the pattern sharpens. The campus study found thresholds satisfied and declared adequacy (Prayogi & Hulu, 2025); the systematic review found connectivity determining outcomes (Akpen et al., 2024); this study found students unable to participate. A network can be adequate by measurement and

inadequate by experience, and the discrepancy is not noise but the consequence of measuring aggregate central tendency where the binding constraint is variability at the tail.

Synthesis and gap. The constraint literature is qualitative, local, and outcome-focused; the optimization literature is quantitative, general, and mechanism-focused. The gap is the absence of any study applying a Theme B, C, or D mechanism in a Theme F setting and measuring both a network metric and a learning outcome. This is the single most consequential gap identified in this review, and it is directly addressable.

Synthesis Matrix

Table 1 consolidates the fifteen reviewed studies. The columns are chosen to expose the pattern that motivates Section IV: the mechanism proposed, how it was validated, and in what setting. Reading the validation column vertically is the fastest way to see the structural gap, because simulation and controlled testbed dominate every mechanism-oriented theme while the two studies that measure people propose no mechanism at all.

Table 1. The fifteen reviewed studies by theme, mechanism, validation method, and setting.

Theme	Study	Mechanism or constraint addressed	Validation	Setting
A	Prayogi & Hulu (2025)	Measurement of throughput, delay, loss, jitter at peak and non-peak	Field measurement	University campus network, Indonesia
A	Amoako et al. (2026)	QoS framework for open distance e-learning: security, bandwidth management, conflict management	Survey, confirmatory factor analysis, SERVQUAL	Tertiary institutions, Ghana
A	Akpen et al. (2024)	Determinants of online learning performance and engagement	Systematic review under PRISMA, 18 studies	Cross-national
B	Wassie et al. (2024)	Elephant flow detection and prediction to protect real-time flows	Deep learning and random forest, simulation	SDN testbed
B	Mehmood et al. (2023)	Dynamic active sensing server load management	Algorithm evaluation on QoS control parameters	SDN telecom testbed
B	Fallah et al. (2023)	Federated learning with genetic algorithms for privacy-preserving load balancing	Comparative testing against round-robin and five baselines	SDN, three test servers
B	Titarenko et al. (2025)	Queue guarantees on bandwidth or utilization coefficient, DiffServ and IntServ	Analytic model	Router interface
C	Marmat & Thankachan (2025)	Comparative evaluation of learned task offloading and resource allocation	Simulation on 100-node topology	Integrated fibre-wireless access network
C	Gama et al. (2026)	Content steering with QoE forecasting to pre-empt congestion	Prototype with ML forecasting, SLO violation measurement	Edge-cloud continuum, up to 40 users
D	Venkatesh Naik & Madhavi (2024)	RNN-LSTM adaptive bitrate selection from network, time, and buffer state	Evaluation on an operating e-learning platform dataset	Collaborative e-learning platform
D	Darwich & Bayoumi (2024)	CNN and RNN per-frame bitrate prediction	Experimental evaluation	Video streaming service
D	Zhang et al. (2025)	PPO reinforcement learning with dual clipping, LSTM and local attention	Simulation, sample efficiency and generalization	Adaptive bitrate streaming
D	Omar et al. (2024)	QoE as a decision criterion; critique of QoE measurability	Review and research agenda	Heterogeneous multimedia networks
E	Mushtaq et al. (2025)	Zero trust architecture components and challenges	Systematic literature review, 74 articles, 2016-2025	Ten application domains
F	Akma & Abatayo (2024)	Effects of unstable connectivity on resource access and collaboration	Qualitative narrative analysis, 10 in-depth interviews	Rural Palimbang, Philippines

Three properties of the table carry the argument. First, **nine of fifteen studies validate in simulation, analytic model, or controlled testbed**, and of the six that do not, three are reviews or syntheses that propose no mechanism of their own, two measure perceptions or outcomes, and one measures packets in the field. Second, **of the nine studies that both propose and validate a mechanism, only one is set in a learning environment**, namely Venkatesh Naik and Madhavi (2024). Five studies are situated in educational settings, but the other four, Prayogi

and Hulu (2025), Amoako et al. (2026), Akpen et al. (2024), and Akma and Abatayo (2024), measure networks or learners without proposing an intervention. Third, **the settings cluster in the global South**: Indonesia, Ghana, Burkina Faso, Iran, Ethiopia, India, and the Philippines are all represented, which is where the constraint binds hardest and, notably, is not where the mechanism-oriented work is deployed.

III. METHODOLOGY

A. Research Design

This study employs a qualitative descriptive design executed as a review of related literature. It analyses existing scholarly works only; no primary data were collected, no system was implemented, and no experiment was conducted. The design is appropriate because the research objective is to establish the state and direction of a distributed literature rather than to test a hypothesis about a single mechanism.

B. Sources of Data

All data are secondary and consist of peer-reviewed journal and conference publications. The fifteen retained records were drawn from Wiley Online Library (four), MDPI (three), SpringerLink (two), ScienceDirect (two), and one each from AIP Publishing, Open Praxis, Philippine E-Journals, and the Indonesian Journal of Informatics and Computer Science. IEEE Xplore, the ACM Digital Library, PubMed Central, and Google Scholar were additionally searched as coverage checks and yielded no records that survived screening.

C. Selection Criteria

Inclusion criteria. A record was included if it was (a) relevant to network or operating system performance in a learning context, or in a multimedia delivery context whose mechanism transfers directly to synchronous instruction; (b) published in a peer-reviewed journal or peer-reviewed conference proceedings; (c) published between 2022 and 2026; and (d) written in English.

Exclusion criteria. A record was excluded if it was (a) non-peer-reviewed, including preprints, patents, vendor white papers, technical blogs, and repository postings without a peer-reviewed venue; (b) published before 2022, except where cited as established background rather than as evidence of a current trend; (c) not relevant to the network or operating system performance layer; or (d) unavailable in a form permitting verification of authorship, venue, and year.

D. Data Collection Procedure

Searches were conducted using keyword combinations drawn from the six themes, including quality of service in online learning, campus network quality-of-service measurement, software-defined networking bandwidth management, edge computing latency reduction, adaptive bitrate quality of experience, zero trust performance, and

internet connectivity in Philippine education. Titles and abstracts were screened for relevance against the inclusion criteria. Records surviving screening were retrieved and read in full where open access permitted, and to the extent of the abstract, stated findings, and reported method where it did not.

A verification step was applied to every record before citation. Author list, publication year, journal or proceedings title, volume and issue, and digital object identifier were confirmed against the publisher record or, where indexed, the Crossref metadata registry. The step was not incidental, and it was iterative. An initial pass corrected five venue, attribution, or year errors in the working source pool and excluded one candidate record whose publication year could not be resolved unambiguously across indexes. A second, later pass against Crossref corrected three reference entries whose author lists were incomplete or whose lead author was misattributed, and supplied three digital object identifiers absent from the first pass. Both passes are reported here rather than only their outcome, because the second pass demonstrates that a single verification round was insufficient. The final retained set is fifteen studies, each with its full author list confirmed against a publisher or registry record.

E. Data Analysis Technique

Thematic synthesis was applied. Each retained study was coded for the mechanism or constraint it addressed, the method of validation used, the setting of that validation, and the findings reported. Codes were grouped into the six themes presented in Section II, with themes defined by mechanism rather than venue. Within each theme, studies were compared for agreement, contradiction, and complementarity, and a research gap was derived from what the theme collectively did not establish. Cross-theme analysis then examined convergence across themes and identified the gaps recurring in more than one, on the reasoning that a gap appearing independently in several themes is structural rather than incidental to a single literature.

F. Ethical Considerations

This study uses only published secondary sources. All sources are attributed in American Psychological Association seventh edition format. No data were fabricated, altered, or selectively reported to support a conclusion, and findings inconsistent with the review's

overall direction are reported where they occur, most notably the campus measurement study whose conclusion of adequacy contradicts the review's general position. No human subjects were involved. Where a source's stated findings could be verified only at the level of the abstract and reported method, the claim drawn from it is confined to what that material supports.

IV. RESULTS AND DISCUSSION

A. Three Cross-Theme Trends

Trend one: from static provisioning to learned, predictive control. The most consistent movement across the reviewed period is the displacement of fixed policy by learned policy, and then of reactive adaptation by forecasting. Flow handling has moved from static classification to deep-learning detection of elephant flows (Wassie et al., 2024). Server load balancing has moved from round-robin to federated learning with genetic optimization (Fallah et al., 2023). Bitrate selection has moved from buffer-threshold heuristics to recurrent networks (Venkatesh Naik & Madhavi, 2024; Darwich & Bayoumi, 2024) and then to reinforcement learning explicitly addressing generalization (Zhang et al., 2025). Endpoint selection has moved from static assignment to forecast-driven content steering that reconfigures before congestion rather than after (Gama et al., 2026). The direction is unambiguous and it is toward self-optimizing networks.

Trend two: prioritization outperforms capacity. The reviewed mechanisms that report the largest improvements do not add bandwidth; they decide which traffic gets it. Queue guarantees stabilize queue length and thereby improve delay and loss probability (Titarenko et al., 2025). Elephant flow detection protects real-time flows by identifying the flows that crowd them out (Wassie et al., 2024). Server load management reduces delay without additional server capacity (Mehmood et al., 2023). This convergence has a specific cause. The flow that fails in an online class is not bandwidth-hungry; a conferencing session requires roughly one to three megabits per second in each direction, far less than a single high-definition stream. It fails because it cannot defer, and because a first-in-first-out queue serves an aggressive bulk flow in proportion to how many packets it offers. Capacity expansion raises the ceiling all flows compete under; prioritization changes who loses when the ceiling is reached. For the sensitive, low-

bandwidth flow, the second intervention dominates.

Trend three: security has entered the performance budget. Zero trust concentrates its implemented components on authentication, authorization, and access control, while auditing and orchestration remain underdeveloped, and lightweight cryptography for constrained systems is identified as an unmet need (Mushtaq et al., 2025). Independently, secure examination appears as a constituent of educational quality of service (Amoako et al., 2026). Together these establish that in education the security and performance objectives are jointly constrained, and that the joint constraint binds hardest on the least-provisioned links.

B. Real-World Application: A Diagnosed Household Failure

To test whether the reviewed literature explains observed practice, the trends are applied to a documented diagnosis of a single failing session. A graduate student attended a synchronous class over shared residential wireless. The session was initially stable and degraded after approximately ten minutes, producing pixelated and freezing video, delayed and distorted audio, repeated reconnection notices, and intermittent disconnection. Concurrently a television streamed high-definition video, a second household member played an online multiplayer game, and a smartphone performed background downloads. The student was positioned at the far end of the dwelling with several walls between the client and the access point.

Layered analysis identified two independent, simultaneous causes. The first was saturation of the shared uplink at the network and transport layers, producing queueing delay and discard. The second was physical-layer signal attenuation from distance and intervening walls, reducing the negotiated data rate, increasing frame retransmission, and consuming shared airtime disproportionately. Neither alone accounts for the full symptom set; together they account for all of it.

The reviewed literature accounts for each element of this case, and the mapping is worth stating explicitly.

Table 2 identifies the protocols carrying each traffic class in the session, because the symptom asymmetry cannot be explained without them. Table 3 then assigns each observed symptom to the layer at which it originates rather than the layer at which it is

perceived, which is the distinction the reviewed literature repeatedly shows practitioners collapsing.

Table 2. Protocols active in the failing session and their role.

Protocol	Layer	Role in the session
DNS	Application	Resolves the conferencing service hostname at join
HTTPS and TLS	Application, Presentation	Login, meeting join, encryption of signalling and media
WebSocket over TCP	Application, Transport	Persistent signalling, chat, participant list
RTP	Application over Transport	Carries audio and video with sequence numbers and timestamps
RTCP	Application over Transport	Reports loss, jitter, round-trip time; drives adaptive bitrate
UDP	Transport	Preferred carrier for media: no retransmission, no head-of-line blocking
TCP	Transport	Signalling, chat, and media fallback where UDP is blocked
STUN, TURN, ICE	Application	NAT traversal and candidate path selection
IP	Network	Addressing and routing between client and server
DHCP	Application over UDP	Address and gateway assignment
ARP	Data Link	Resolves the gateway MAC address on the local segment
IEEE 802.11	Data Link, Physical	Wireless medium access, acknowledgment, rate adaptation

The division explains the asymmetry directly. Media travels over UDP with RTP and is not retransmitted, so a lost media packet is simply absent and surfaces as a visible or audible artefact. Chat travels over TCP and is retransmitted, so a lost chat packet is recovered transparently and surfaces only as a delay the user does not perceive. This is why the same link that renders video unusable leaves text unaffected, and it is the reason capacity-based reasoning misdiagnoses the fault.

Table 3. Layer of origin versus layer of perception for each observed symptom.

Layer	Role in the session	Failure originating here
7. Application	Client, DNS, signalling, chat, screen share	Adaptive bitrate reducing quality in response to RTCP loss reports
6. Presentation	Media encode and decode, TLS	Codec renegotiating to a lower profile; resolution downgrade
5. Session	Session establishment, keep-alives, ICE management	Keep-alive timeout producing the reconnection state
4. Transport	UDP carrying RTP, TCP carrying signalling	Loss without retransmission; jitter exceeding de-jitter buffer depth
3. Network	IP routing, client to gateway to ISP to server	Queueing delay and discard at the saturated uplink
2. Data Link	802.11 medium access, acknowledgment	Airtime contention, collisions, backoff, retransmission

Layer	Role in the session	Failure originating here
	t, rate adaptation	
1. Physical	Radio propagation between client and access point	Attenuation from distance and walls; reduced modulation rate

Layers 5 through 7 exhibit symptoms but are not causes. The reconnection state, the codec downgrade, and the adaptive bitrate reduction are all correct responses to conditions created beneath them. This is the practical value of layered analysis and the reason the fifth-ranked remediation in the case, reducing video quality in the client, was ranked last: it acts at the layer of perception while the cause persists two layers below.

The delayed onset is a queue-length phenomenon. Degradation appeared after ten minutes rather than immediately because competing flows had not reached steady-state throughput at the outset. A download's congestion window expands until it encounters loss; a streaming session's sustained rate stabilizes only after initial buffering. The interval is the time for those flows to collectively exceed capacity. The queue optimization work supplies the mechanism: it is average queue length that governs delay and packet loss probability (Titarenko et al., 2025), so degradation tracks the filling of the queue rather than the arrival of any single flow.

The competing download is an elephant flow. The traffic that displaced the class session is exactly what the software-defined networking literature targets: a small number of large aggressive flows consuming disproportionate capacity, detectable and predictable by learned classification, with resource allocation to real-time applications as the stated objective (Wassie et al., 2024). The household case is the residential instance of a problem the literature addresses at carrier scale.

The codec downgrade is adaptive bitrate operating correctly. The observed progressive quality reduction was not a fault but the adaptation mechanism responding to loss reports, selecting a lower bitrate from buffer level, throughput, and rebuffering history exactly as the reviewed adaptive delivery work describes (Venkatesh Naik & Madhavi, 2024; Darwich & Bayoumi, 2024). The client degraded quality to preserve continuity, which is the designed behaviour.

The remediation order the case derived is the order the literature predicts. The diagnosis ranked quality-of-service

prioritization first and wired connection second, ahead of capacity upgrade. This ordering was derived from the symptom analysis alone, and it coincides with Trend two. The theoretical warrant is the queue guarantee mechanism (Titarenko et al., 2025); the operational warrant is the flow protection objective (Wassie et al., 2024).

Two failures the literature does not address. The case also exposes the limits of the reviewed work. First, the binding direction was upstream. Residential connections are asymmetric, and the student's outgoing media competed with acknowledgment traffic generated by every downstream flow in the household. Every adaptive delivery study reviewed optimizes downstream delivery to a consuming client. Synchronous instruction is bidirectional, and the uplink is where it breaks. Second, the physical-layer cause is untouched by any reviewed mechanism. No amount of prioritization, edge placement, or bitrate adaptation compensates for a client that has negotiated a low modulation rate because of distance and wall attenuation, and that consequently occupies shared airtime disproportionately. The optimization literature operates above the layer where this household's problem originated.

C. Implications for Schools, Industry, and Policy

For schools. The evidence supports an intervention order rather than a shopping list. Prioritization mechanisms should be exhausted before capacity is purchased, because prioritization protects the sensitive flow regardless of competing load while capacity raises a ceiling that aggressive flows will reoccupy. Where an institution operates its own learning management system, server-side load distribution is a distinct and separately remediable failure mode (Mehmood et al., 2023). The measurement caution from Theme A should govern procurement: an aggregate quality-of-service report showing thresholds satisfied is not evidence that synchronous sessions are usable, and administrators should require jitter and loss distributions under peak load rather than mean values.

For industry. Content steering and forecast-driven reconfiguration are entering standardized adaptive streaming protocols rather than remaining research constructs (Gama et al., 2026), which implies that educational platform providers will be able to consume these capabilities without bespoke

engineering. The unaddressed commercial opportunity is upstream-aware adaptation for bidirectional instructional traffic.

For policy. The digital divide literature and the optimization literature do not intersect (Akmad & Abatayo, 2024), and public investment decisions in developing regions are made precisely at that intersection. A policy instrument that funds capacity without funding the configuration competence to prioritize it will underperform, because the reviewed evidence indicates that control, not capacity, is the higher-yield intervention. The gap in lightweight cryptography for constrained systems (Mushtaq et al., 2025) has a policy dimension as well, since regulatory requirements for secure assessment impose a latency cost that falls hardest on the least-provisioned learners.

D. The Structural Gap

One gap recurred in three of the six themes, Themes A, D, and F, and is therefore treated as structural rather than local. The optimization literature validates in simulation and controlled testbeds, reporting improvements in throughput, delay, loss, and quality-of-experience proxies. The education literature measures engagement, participation, and attainment. No reviewed study measures a network intervention and a learning outcome together, and the two literatures do not cite one another.

The consequence is practical. An administrator can be shown that queue guarantees reduce packet loss probability and, separately, that connectivity affects attainment, but cannot be shown that enabling prioritization improves attainment. Two further gaps compound it: quality of experience cannot yet be measured reliably in real time (Omar et al., 2024), which removes the natural intermediate variable that would bridge the two literatures, and no study quantifies the latency contribution of access control decisions inside a real-time session (Mushtaq et al., 2025).

This gap is addressable with modest resources, which is the reason for naming it precisely. A single institution operating a constrained network could enable prioritization on one of two comparable sections, instrument jitter and loss at the client, and compare participation and assessment completion across sections. No study in this review occupies that intersection, and such a study would not require infrastructure that the settings in question lack.

E. Limitations of This Review

Four limitations bound the claims made here.

Sample size and selection. Fifteen studies satisfy the stated requirement but do not constitute an exhaustive survey of a literature this active. Selection proceeded by thematic keyword search rather than by an exhaustive database enumeration with reported yield at each screening stage, so this review is a thematic synthesis and not a systematic review in the PRISMA sense. The distinction matters because publication bias cannot be assessed from a non-exhaustive sample, and mechanisms that were evaluated and found ineffective are systematically less likely to appear.

Depth of access. Several sources were available in full text under open access, while others were read at the level of abstract, stated method, and reported findings. Claims drawn from the latter are confined to what that material supports, and no claim in Section II attributes to a study a finding its abstract and method do not state. This is a real constraint on the granularity of the synthesis: methodological weaknesses visible only in a full results section could not be assessed uniformly.

Uneven theme population. Themes B and D carry four studies each, while Themes E and F carry one each. The security and digital-divide conclusions therefore rest on a single study apiece, and while both are substantial, a systematic review of seventy-four articles in the case of Theme E and a ten-participant qualitative study in the case of Theme F, neither can be cross-validated against a second source within this review. The Theme E and Theme F findings should be read as well-supported single-source observations rather than as convergent findings.

The application case is a single diagnosis. The household analysed in Section IV is one case, diagnosed by layered inspection rather than instrumented measurement, and it is used to test whether the literature explains observed practice rather than as independent evidence. That it reproduced the literature's intervention ordering independently is suggestive and no stronger than suggestive. A single case cannot establish generality, and the two failures it exposed, the uplink direction and the physical layer, are identified as gaps in the literature on the strength of that literature's own stated scope rather than on the strength of one household.

V. CONCLUSION AND RECOMMENDATIONS

A. Conclusion

Fifteen peer-reviewed studies published between 2022 and 2026 were reviewed to establish how contemporary operating system and network technologies address the reliability of online learning. Three trends were identified. Network control is moving decisively from static provisioning toward learned and increasingly predictive policy, with machine learning now embedded in flow classification, server load distribution, bitrate selection, and endpoint steering, and with forecast-driven pre-emption emerging as the frontier. Prioritization consistently outperforms capacity expansion for this application, because the flow that fails in an online class is low-bandwidth and latency-bound rather than bandwidth-hungry, and because first-in-first-out service systematically disadvantages exactly such flows. Security has entered the performance budget, since zero trust concentrates its mature components on the per-request verification path while lightweight cryptography for constrained systems remains unmet, making security and performance a joint constraint in education rather than separable objectives.

Applied to a documented household failure, the literature explained the delayed onset as a queue-length phenomenon, the displacing traffic as an elephant flow, and the observed quality reduction as adaptive bitrate operating as designed. The remediation order the case derived from symptom analysis coincides with the ordering the reviewed literature converges on, which is a suggestive convergence between a single diagnosis and a body of simulation-validated research. Two limits were exposed: the reviewed adaptive delivery work optimizes downstream delivery while synchronous instruction fails upstream on asymmetric residential links, and no reviewed mechanism addresses the physical-layer attenuation that originated half of the case's symptoms.

The principal finding is negative and structural. The network optimization literature and the educational outcome literature do not meet. Optimization is validated on metrics in simulation; educational effect is measured on outcomes in the field; and no reviewed study connects an intervention to an outcome. Until that connection is made, decisions about how to spend constrained educational infrastructure budgets will continue to rest on inference rather than evidence.

B. Recommendations for Practitioners

6. **Exhaust prioritization before purchasing capacity.** Configure quality-of-service guarantees for instructional traffic or devices first. The reviewed evidence indicates control yields more than capacity for latency-sensitive low-bandwidth flows, and it costs nothing recurring.
7. **Eliminate the wireless variable for the delivering endpoint.** Where instruction originates from a fixed position, use wired connection. No reviewed optimization mechanism compensates for physical-layer attenuation or airtime contention.
8. **Require distributions, not means, in network assessment.** A report showing aggregate thresholds satisfied is compatible with unusable synchronous sessions. Specify jitter and loss distributions under peak load.
9. **Treat the learning platform server as a separate failure domain.** Where an institution hosts its own system, evaluate server-side load distribution independently of access-link capacity.
10. **Cost security decisions in latency, not only in risk.** On constrained links, per-request verification draws from the same budget as media. Where lightweight mechanisms are unavailable, the trade-off should be made explicitly rather than by default.

C. Directions for Future Research

11. **Bridge the measurement gap.** Conduct paired studies that apply a prioritization, edge, or adaptive delivery mechanism and measure both network metrics and learning outcomes in the same population. This is the highest-value gap identified and the most tractable.
12. **Address the uplink.** Extend adaptive delivery research to bidirectional instructional traffic on asymmetric access links, where the reviewed downstream-oriented work does not apply.
13. **Establish minimum viable control.** Determine the smallest software-defined or queue-management deployment that yields measurable benefit in an institution without dedicated network operations staff.
14. **Quantify the security latency cost.** Measure the contribution of access control decisions to end-to-end latency inside real-time educational sessions, and evaluate

lightweight cryptographic alternatives in that setting.

15. **Resolve real-time quality-of-experience measurement.** Reliable in-session measurement is the prerequisite for the intermediate variable that would connect network state to educational outcome.

VI. REFERENCES

- Akmađ, A. P., & Abatayo, A. V. L. (2024). Caught in the slow lane: Effects of unstable internet connectivity on accessing academic resources and collaborative learning. *Journal of Interdisciplinary Perspectives*, 2(9), 278-286. <https://doi.org/10.69569/jip.2024.0370>
- Akpen, C. N., Asaolu, S., Atobatele, S., Okagbue, H., & Sampson, S. (2024). Impact of online learning on student's performance and engagement: A systematic review. *Discover Education*, 3(1). <https://doi.org/10.1007/s44217-024-00253-0>
- Amoako, P. Y. O., Mnkandla, E., & Adarkwah, M. A. (2026). Validating a framework for quality of service in open distance electronic learning. *Open Praxis*, 18(1), 66-85. <https://doi.org/10.55982/openpraxis.18.1.1003>
- Darwich, M., & Bayoumi, M. (2024). Video quality adaptation using CNN and RNN models for cost-effective and scalable video streaming services. *Cluster Computing*, 27(5), 6355-6375. <https://doi.org/10.1007/s10586-024-04315-8>
- Fallah, M., Mohammadi, P., NasiriFard, M., & Salehpour, P. (2023). Optimizing QoS metrics for software-defined networking in federated learning. *Mobile Information Systems*, 2023, 1-10. <https://doi.org/10.1155/2023/3896267>
- Gama, E. S., Rodrigues-Filho, R., Madeira, E. R. M., Immich, R., & Bittencourt, L. F. (2026). Adaptive video streaming architecture leveraging QoE forecasting and content steering on the edge-cloud continuum. *Future Generation Computer Systems*, 183, Article 108521. <https://doi.org/10.1016/j.future.2026.108521>
- Marmat, A., & Thankachan, D. (2025). ML-driven latency optimization for mobile edge computing in fiber-wireless access networks. *MethodsX*, 15, Article 103594. <https://doi.org/10.1016/j.mex.2025.103594>
- Mehmood, K. T., Atiq, S., & Hussain, M. M. (2023). Enhancing QoS of telecom networks through server load management in software-defined networking (SDN). *Sensors*, 23(23),

Article 9324.
<https://doi.org/10.3390/s23239324>

Mushtaq, S., Mohsin, M., & Mushtaq, M. M. (2025). A systematic literature review on the implementation and challenges of zero trust architecture across domains. *Sensors*, 25(19), Article 6118.
<https://doi.org/10.3390/s25196118>

Omar, M. H. H., Kouraogo, J. P., Kabre, W. J. M., Tapsoba, S. D. W., & Sie, O. (2024). Using QoE metric as a decision criterion in multimedia heterogeneous network optimization: Challenges and research perspectives. *Journal of Computer Networks and Communications*, 2024, 1-18.
<https://doi.org/10.1155/2024/7864757>

Prayogi, S. Y., & Hulu, D. (2025). Quality of service (QoS) analysis of campus internet networks: A case study at Universitas Islam Sumatera Utara. *IJICS (International Journal of Informatics and Computer Science)*, 9(3), 199-202.
<https://doi.org/10.30865/ijics.v9i3.9582>

Titarenko, L., Lemeshko, O., Yeremenko, O., Savchenko, R., & Barkalov, A. (2025). Traffic engineering queue optimization models with guaranteed quality of service support. *Electronics*, 14(20), Article 4078.
<https://doi.org/10.3390/electronics14204078>

Venkatesh Naik, N., & Madhavi, K. (2024). Adaptive QoE-based video delivery by optimizing the bitrate of the video course using network bandwidth and user bandwidth. *Concurrency and Computation: Practice and Experience*, 36(16), Article e8114.
<https://doi.org/10.1002/cpe.8114>

Wassie, G., Ding, J., & Wondie, Y. (2024). Detecting and predicting models for QoS optimization in SDN. *Journal of Computer Networks and Communications*, 2024, Article 3073388.
<https://doi.org/10.1155/2024/3073388>

Zhang, J., Han, Y., Cai, Z., Feng, Y., & Zhu, L. (2025). Deep reinforcement learning enhanced optimization algorithm for adaptive bitrate video streaming. *AIP Advances*, 15(7), Article 075042. <https://doi.org/10.1063/5.0277381>