Evaluating FEC removal in XGS PON

What the simulations establish and a credible path to vendor adoption

Prepared for Audeamus Labs
Technical review and research proposal
8 November 2026

Abstract

Removing forward error correction can increase the fraction of an optical link available for user data when the uncoded channel is sufficiently reliable. The Audeamus Labs simulation correctly illustrates that trade. Its matched clean-channel case increases simulated payload delivery from 7.969792 to 9.081856 Gb/s in each direction, a 13.9535% improvement. Independent reruns of the supplied code reproduce those results exactly.

The result is conditional on a substantial protocol change: downstream parity bytes become payload bytes. Current XGS-PON requires downstream FEC. An implementation that transmits extra payload in those positions would therefore need a defined extension, compatible transmitters and receivers, and a standards path before it could claim interoperable XGS-PON operation. Disabling a decoder alone does not create capacity.

The simulation also assumes uncalibrated receiver error curves, ideal retransmission, fixed service allocations, and whole-packet packing at 125 µs intervals. Its clean uncoded case drops 5.397% of offered packets at the source queue, despite recording no optical corruption. It demonstrates a simulated service rate under overload, not a lossless 9.6 Gb/s service or measured TCP performance. The existing paper discloses many of these limitations; the principal concern is how strongly its headline can be interpreted.

A defensible next step is to simulate an explicit experimental protocol, benchmark it against current compliant operation, and then test its framing, optical reliability, transport behavior, and interoperability on hardware. Upstream-only FEC control provides an earlier standards-supported investigation. Downstream byte reuse should remain an experimental extension until its additional capacity survives realistic errors, multi-ONU operation, and transition failures.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

The simulated optical topology

Figure 1 reproduces the original study topology in its coded reference state. The title within the image uses “measured” for simulator counters, not measurements from optical hardware.

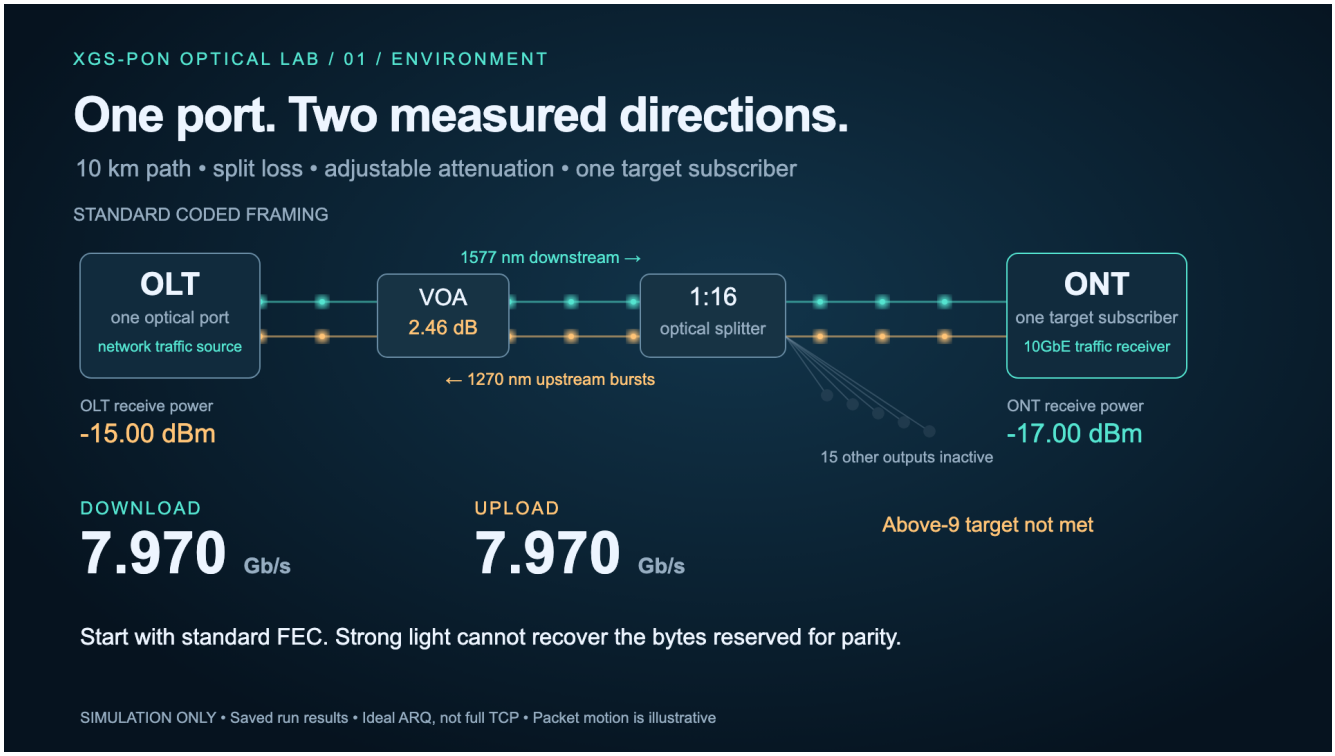


Figure 1. Original study topology, shown at frame 90 of the OpticalLab composition. Rates are saved simulation results; packet motion is illustrative. This is not a physical measurement or evidence of hardware support.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

The question worth answering

The useful research question is whether optical margin can be exchanged for more delivered payload while maintaining an agreed service objective. That objective must include packet loss, delay, availability, and service to other subscribers, not only a peak rate.

Three different propositions need separate answers:

1. **Arithmetic:** Replacing parity with useful information creates additional payload space at the same optical line rate. This is correct.
2. **Engineering:** A particular OLT and ONU can use that space reliably, including during adverse optical conditions. The supplied work does not establish this.
3. **Deployment:** The resulting mode can operate within the current standard and interoperate with deployed equipment. The proposed downstream mode does not meet the present downstream FEC requirement.

The distinction leaves a worthwhile research program. It prevents a numerical result from being presented as an available firmware setting or a field-validated service improvement.

This review examines the supplied optical-study manuscript, model code, experiment manifest, packet counters, and selected raw runs. It also reviews primary standards, experimental papers, a closely related patent disclosure, and IETF measurement guidance. Public abstracts are identified as such. No optical equipment was accessed, and no hardware performance is claimed.

Reconstructing the controlled result

The most useful comparison is run 0697 against run 0700. Both use an installed 1:16 split, one active ONU, a 10 km modeled path, received powers of −17 dBm downstream and −15 dBm upstream, identical packet formats, and the same random seed. Each measures five seconds after a one-second warmup. Both directions are offered 9.6 Gb/s of payload simultaneously.

The packet definition is a 1500-byte IPv4 MTU, with 20 bytes of IPv4 header, 20 bytes of TCP header, and 12 bytes of TCP options. That leaves 1448 payload bytes. The model charges 1528 bytes on the PON for each data packet and 1538 bytes at the Ethernet user interface. It generates a separate acknowledgment for every two successfully delivered data packets. This is TCP-like accounting, not a TCP implementation.

Metric in each direction	Coded reference	Hypothetical uncoded case
Offered packets during measurement	4,143,646	4,143,646
Successfully delivered packets	3,440,000	3,920,000
Source queue drops	703,646	223,646

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

Metric in each direction	Coded reference	Hypothetical uncoded case
Source queue drop fraction	16.981%	5.397%
Corrupt transmission attempts	0	0
Retransmissions	0	0
Delivered payload rate	7.969792 Gb/s	9.081856 Gb/s

The receiver rate follows directly from the uncoded counter:

3,920,000 packets × 1448 payload bytes × 8 / 5 seconds = 9.081856 Gb/s.

The controlled uplift is:

9.081856 / 7.969792 − 1 = 13.9535%.

This is a valid statement about this model. Its agreement with the expected parity-recovery opportunity is reassuring. It is also largely predetermined by the changed byte budget and the assumed clean channel. Repeating the calculation across similar clean scenarios does not independently prove that a physical receiver will remain clean without FEC.

Why the gain is near fourteen percent

The coded downstream frame contains 627 blocks with 216 information bytes and 32 parity bytes each. The supplied model changes the available framing-sublayer bytes from 135,432 to 155,496 per frame. The difference is 20,064 bytes every 125 µs, or 1.284096 Gb/s before higher-layer overhead. These quantities correspond to the framing described in G.9807.1 and the model's hypothetical reuse of the parity positions. ITU-T G.9807.1, Annex C.

Two percentages are easily confused. Parity consumes 32/248, or 12.903%, of a complete codeword. Recovering those positions increases information capacity relative to the original 216 information bytes by 32/216, or 14.815%. The simulated application-payload uplift is smaller because packet overhead, acknowledgments, interface limits, and discrete packing also consume capacity.

The optical line rate has not increased. The proposal changes how that line rate is allocated between redundancy and information.

The bottleneck moves to the Ethernet interface

Using the supplied model's continuous-packing calculation, with the same packets and acknowledgment ratio, produces an additional insight:

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

Continuous clean-channel bound	Downstream	Upstream
Coded PON payload bound	8.0034 Gb/s	7.9904 Gb/s
Uncoded PON payload bound	9.1892 Gb/s	9.1752 Gb/s
10GbE bound including reverse-direction ACK traffic	9.1472 Gb/s	9.1472 Gb/s

The clean uncoded case is therefore interface-limited in the continuous approximation. Its 9.081856 Gb/s frame-stepped result lies below that bound. Once the PON ceases to be the bottleneck, additional optical payload space cannot improve delivery without addressing the next limit.

This does not invalidate the gain. It makes its dependence on the traffic format explicit. Separate bidirectional data streams, standalone acknowledgments, TCP options, IPv6, VLANs, PPPoE, packet-size distributions, or different interface speeds can change the result.

What removing FEC changes in the standard

G.9807.1 specifies that the downstream FEC flag must be set and that the OLT is statically configured to insert downstream parity. Upstream FEC can be activated or deactivated for an ONU through the applicable burst profile. These are direction-specific provisions; upstream flexibility is not permission to reuse downstream parity positions. The relevant clauses are C.10.1.1.1.3, C.10.1.3.1.2, and C.10.1.3.2.4. ITU-T G.9807.1, 2025 amendment.

Three implementation changes must be distinguished:

- **Bypassing error correction:** The receiver skips correction but still processes the existing frame layout. There is no automatic payload increase.
- **Suppressing parity while preserving its positions:** The transmitter substitutes filler or unprotected bytes that the receiver still discards. The payload budget again remains unchanged.
- **Reusing parity positions for payload:** Both ends adopt a larger information budget and compatible interpretation of those bytes. This is the change that generates the headline gain.

The last option affects more than the Reed-Solomon block. A complete design must specify frame parsing, data-path widths or rates, buffer sizing, synchronization, control-field protection, capability discovery, mode changes, counters, and fallback. Encryption and authentication must remain intact, with their framing interactions checked rather than assumed.

A vendor could build a proprietary pair for an isolated experiment. That would show implementation feasibility. It would not make the mode compliant with the current downstream requirement, establish support in existing silicon, or guarantee interoperability with another vendor.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

Shared downstream operation is the difficult case

A downstream transmission is shared by multiple receivers with different optical margins. A policy based only on the strongest subscriber can expose weaker ONUs to a format or error rate they cannot tolerate. Nor can a new format assume that legacy ONUs will safely ignore frames they cannot parse: control information, synchronization, and subsequent frame alignment may still matter.

The simplest initial experiment should therefore use a dedicated port whose OLT and every attached ONU support the experimental format. Mixed legacy and enhanced operation is a separate research problem. Any selective downstream mode would need an explicit compatibility design, protected common control, and tests proving that legacy and weak receivers remain synchronized and retain their service.

Removing FEC also removes corrected-error counts as an observation mechanism. Upstream monitoring already has a specified alternative: G.9807.1 clause C.10.1.3.2.4 uses the FS-trailer BIP-32 to estimate BER when upstream FEC is off. The experimental downstream profile would need its own defined monitoring and adaptation policy, potentially combining frame-error counters with periodic coded probes. It must include hysteresis and a reliable return to the coded mode. Otherwise a transient can damage both user traffic and the control exchange needed to recover.

The architectural idea has prior art

Denis Khotimsky's patent disclosure, published in 2015 and later granted as US10044467B2, discusses the same 135,432-to-155,496-byte transition for XG-PON1 and NG-PON2. It separates FEC status from an extra-data indicator and addresses the encoder, decoder, and data-path transition. This is strong evidence that the framing problem is concrete and already recognized. It is not evidence of a shipping implementation, measured performance, or current XGS-PON compliance. Patent disclosure US20150381313A1.

Audeamus should position its contribution around reproducible evaluation, measured operating conditions, and a carefully specified adoption path. The general concept of replacing parity with payload should not be presented as a new discovery. A commercial implementation would also warrant a separate professional intellectual-property review; no freedom-to-operate conclusion is offered here.

Where the simulation is strongest

The supplied work has several real strengths.

First, it makes byte accounting inspectable. Packet headers, acknowledgments, framing overhead, and upstream burst costs are explicit rather than hidden behind a generic percentage.

Second, it distinguishes installed split loss from active subscribers. A lightly used splitter does not become a lower-loss splitter. The matched-power experiments also avoid confusing optical power changes with capacity changes.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

Third, the saved manifests, raw counters, and model hashes make individual results reproducible. Independent reruns of the two headline cases reproduce their complete saved results. The artifact's audit reports 32 passing unit tests and reconciled counters; those are useful checks of implementation consistency, although they do not validate the physical model.

Fourth, the original manuscript already warns that downstream byte reuse is hypothetical, hardware support is unverified, receiver curves are assumed, and transport is not full TCP. It also states that external benchmarks use different layers or scenarios. A fair critique should preserve those disclosures. The problem is the risk of readers treating a conditional model result as a deployable improvement despite them.

Finally, the direction of the optical tradeoff is plausible. Fixed protection can be more conservative than some individual paths require. Research on flexible PON designs supports investigating whether that excess margin can be used productively. It does not establish the operating conditions of this particular proposal.

Where the evidence is insufficient

The optical conclusion is assumed rather than measured

The receiver model uses a logarithmic BER curve anchored at a reference sensitivity. Its slope is swept through 0.5, 1.0, and 1.5 decades per dB, with an optional 2 dB penalty and a floor of 10^{-15}. These are declared assumptions, not fitted receiver measurements.

At the matched reference powers, the selected curve produces a pre-FEC BER of 10^{-14} in both directions. For a 1518-byte Ethernet frame, the model then predicts an uncoded packet-error probability near 1.2144×10^{-10}. Across 3.92 million attempts, the expected number of corrupt attempts is only about 0.00048. A zero-error run is the overwhelmingly likely consequence of that input.

An optical power reading cannot establish that BER. Receiver noise, extinction ratio, dispersion, reflections, temperature, aging, interference, and burst acquisition can produce different error behavior at the same average received power. A reference sensitivity point does not determine the slope or error floor away from it.

The immediate question for a vendor is therefore not whether −17 dBm “works” in this model. It is what measured error distribution its receiver exhibits across power, temperature, channel impairments, and time, and whether that distribution satisfies an uncoded service objective.

Error location matters as well as average BER

The coded model calculates an analytical RS codeword-failure probability under independent errors and an ideal bounded-distance decoder, using more than 16 erroneous byte symbols as the failure condition. It assumes a fixed worst-case nine intersected codewords for each 1528-byte PON packet, converts that probability to a packet-erasure probability, and then samples packet erasures independently. It does not simulate individual codewords, shared-codeword correlations between packets, synchronization failures, miscorrection, correlated bursts, or decoder implementation behavior.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

Brandonisio and colleagues experimentally showed why this distinction matters. Their 10 Gb/s burst-mode PON experiment measured pre-FEC BER, post-FEC BER, frame loss, and error locations. Shorter preambles produced localized errors and frame-loss degradation that average BER alone did not adequately describe. Their receiver and long-reach setup are not a universal XGS-PON specification, so their numerical thresholds should not be transferred directly. The relevant lesson is to calibrate error distribution and acquisition behavior, not only mean BER. Brandonisio et al., OFC 2017.

This objection cuts both ways. An ideal independent-error model can overstate coded protection when errors cluster. But that does not vindicate uncoded operation: a cluster or synchronization failure can also remove the uncoded mode's small payload advantage. Both modes need the same measured channel traces.

Source queue drops are not optical losses

The 223,646 drops in the uncoded reference are admission drops at a finite source queue, caused by offering more data than the modeled service can deliver. They are not corrupt optical packets, and subtracting them again from the reported receiver goodput would double-count the undelivered traffic.

They still matter. The case does not demonstrate delivery of all offered traffic. It delivers roughly 9.082 Gb/s while discarding approximately 0.518 Gb/s of offered payload. The coded case discards more, so the comparison remains useful as a saturated service-rate test.

A lossless-throughput claim requires a separate offered-load search and an appropriate loss objective. RFC 2544 distinguishes throughput and frame-loss tests, and recommends final throughput trials of at least 60 seconds. Those laboratory methods should inform hardware evaluation rather than be represented as a certification of the present simulation. RFC 2544, sections 26.1 and 26.3.

The queues persist across warmup and are not drained at the end of a run. The clean matched counters reconcile without ambiguity, but error-heavy tests should explicitly report initial and final queue occupancy, outstanding retries, and in-flight feedback. Long-run delivery claims need conservation checks across all those states.

Ideal retransmission does not reproduce TCP

The model has fixed feedback delay, reliable acknowledgments, and selective-repeat retries. It lacks a congestion window, receive window, slow start, retransmission timeout behavior, acknowledgment loss, and implementation-specific congestion control. Its retry queue is separate from the bounded new-packet queue.

Those choices are reasonable for isolating byte accounting. They are not adequate for predicting application TCP behavior when packets disappear. Real transport performance depends on the entire path, including round-trip time, bottleneck bandwidth, buffering, and loss response. RFC 6349 specifically distinguishes packet-layer measurements from end-to-end TCP performance and identifies these dependencies. RFC 6349.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

Consequently, describing 9.081856 Gb/s as measured TCP goodput or an iperf result would be incorrect. “Simulated delivered payload with TCP-sized headers and ideal ARQ” is accurate.

The nine gigabit threshold is sensitive to packing

The reference model sends 86 data packets per frame interval in the coded case and 98 in the uncoded case. One 1448-byte payload packet per 125 µs interval represents **92.672 Mb/s**.

The uncoded result exceeds 9 Gb/s by **81.856 Mb/s**, less than that single-packet step. This does not prove that the result is wrong or that the true rate must be lower. It shows that a threshold claim is sensitive to whole-packet packing, residual byte carryover, and the treatment of fragmentation and acknowledgments. Continuous serialization could recover some discarded space; realistic control or encapsulation costs could consume it.

A byte-accurate serializer with fragmentation and cross-frame carryover is needed before attaching significance to the distinction between 8.99 and 9.08 Gb/s. Extra decimal places in deterministic counters are not measurement accuracy.

Many runs do not replace physical validation

The package contains 1089 candidate records, of which 843 were executed and 246 were excluded by modeled reference-range checks. The 22,680 link-budget rows and 18,600 continuous-screening rows are calculations, not additional frame experiments. Executed measurement time totals 4710 virtual seconds across separate scenarios. Most runs last five seconds; nine longer runs last 60 seconds.

This is a broad parameter exploration. It is not 843 independent observations of deployed optical links. The same assumed receiver family and simplified protocol govern every run, and similar clean runs have little stochastic variation because their input error probabilities are tiny.

For scale, even an actual five-second error-free bit test at 9.95328 Gb/s would only put an approximate one-sided 95% upper bound of 6.02×10^{-11} on BER under independent, stationary errors. Demonstrating an upper bound of 10^{-12} with zero observed errors would require about 301 seconds; 10^{-14} would require about 8.36 hours. These are elementary statistical exposure calculations, not sufficient qualification plans. Environmental and burst effects require additional testing. A synthetic test driven by an assumed BER cannot establish that physical BER at any duration.

Low subscriber activity is doing substantial work

The headline is a one-active-ONU result. The supplied 1:16 case with 10 Mb/s reserved for each of fifteen peers yields approximately 8.989184 Gb/s downstream and 8.711168 Gb/s upstream for the target. The stricter “above 9 Gb/s in both directions” objective already fails under that modest modeled sharing.

This is not evidence that sharing is malfunctioning. Capacity is being allocated to peers, and upstream bursts add overhead. It demonstrates why a peak subscriber rate must be reported

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

with the active-ONU count, peer demand, burst profile, and fairness policy. Fixed reservations do not reproduce a vendor's dynamic bandwidth allocation algorithm or its transient behavior.

Quantifying the reliability trade

A simple calculation makes the risk visible. Assume independent bit errors, a 1518-byte frame, and packet erasure if any of its 12,144 bits is wrong. Then the packet-error probability is:

p = 1 − (1 − BER)^12144.

Uncoded BER	Calculated packet-error probability
10^{-8}	0.0121%
10^{-7}	0.1214%
10^{-6}	1.2071%
10^{-5}	11.4356%
10^{-4}	70.3130%

These are illustrative calculations, not receiver predictions. They omit errors in other framing and control fields, burst dependence, and actual transport behavior.

If the uncoded mode's attempt capacity were fixed at its clean value, every failed attempt simply consumed capacity, and the coded reference remained clean, the rough throughput break-even condition would be:

9.081856 × (1 − p) > 7.969792.

This permits p below approximately 12.245%, corresponding to BER around 1.08×10^{-5} under the independent-error assumption. Such a loss rate would be unacceptable for many services. It is a mathematical throughput comparison, not an operational threshold.

The “above 9 Gb/s” objective is much tighter. Under the same deliberately simplified calculation, the failure fraction must remain below 0.9013%, corresponding to BER around 7.46×10^{-7}. Real ACK coupling, recovery delays, TCP behavior, and control errors can change both estimates. A production reliability requirement may be orders of magnitude stricter than either.

The practical conclusion is that an uncoded policy needs two gates: it must improve useful delivery, and it must independently satisfy the service's reliability and delay requirements. Passing one does not imply passing the other.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

What the wider literature supports

Upstream optimization is already an active research subject

The 2025 ECOC paper *Optimization of Upstream XGS-PON Throughput by Adjusting Burst Preamble Length and Enabling Forward Error Correction* directly addresses the relationship among optical budget, fiber length, preamble length, and FEC selection. Its topic supports pursuing a direction-specific optimization rather than assuming downstream and upstream have identical implementation constraints. This review does not use inaccessible full-paper results as quantitative validation. ECOC 2025 publication.

A 2026 OFC abstract by Simon and colleagues reports estimates using field data from 68,000 ONUs, with gains as high as 70% from tuning upstream header and FEC parameters. This is evidence that a large field population can inform the optimization question. It is an estimate with a particular baseline and combined adjustments, not a 70% physical line-rate increase, a downstream result, or confirmation of Audeamus's 14% comparison. Without the full methods and distributions, the headline cannot be transferred to this model. Simon et al., OFC 2026.

Flexible coding can be demonstrated on hardware

The FLCS-PON work reports a flexible 100 Gbit/s prototype and operator-trial research, adjusting modulation, coding, and probabilistic shaping to available conditions. It establishes a broader engineering precedent for exploiting heterogeneous optical paths with an explicitly designed physical layer. Its technology and protocol differ from XGS-PON parity-byte reuse, so it cannot supply a numerical gain or compatibility claim for this proposal. Borkowski et al., Journal of Optical Communications and Networking, 2022.

Kubo, Kariyawasam, and Kimura investigate adaptive RS decoding for power savings in a 10G-class optical access network, reporting more than 48% average decoder power saving in their evaluation. That is relevant to the idea that some links do not require the same decoding effort. Decoder energy reduction and increased payload capacity are different outcomes and should not be conflated. IEICE Communications Express, 2024.

External throughput figures need matched assumptions

ITU-T Supplement 75's illustrative XGS-PON upstream calculation gives 8.439 Gb/s at the TC layer and 8.257 Gb/s Ethernet throughput using four bursts per 125 µs interval, a 1032-byte per-burst overhead allowance including the FS header and trailer, and an average 364-byte Ethernet frame. These differ materially from the supplied large-packet, one-active-ONU, bidirectional payload case. The supplement also illustrates FEC and DBA latency contributions, reinforcing the need to evaluate the whole service path. ITU-T G Supplement 75, Tables 14 and 16.

The original paper's “10-14%” opening combines a controlled model comparison with gaps to external figures. It is not an uncertainty interval. The scientifically stronger lead is **“13.95% higher simulated payload delivery than the matched coded model, under hypothetical downstream byte reuse.”** External figures can remain as context if their measurement layer, direction, packet mix, and operating conditions remain visible.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

Why the conclusion cannot simply be extended to 50G PON

The original investigation also names 50G-PON, but the reviewed frame-stepped results are an XGS-PON model. Higher-speed PON uses a different coding and framing framework. G.9804.2 specifies LDPC coding, including LDPC(17280,14592) for the nominal 49.7664 Gb/s downstream case, and retains mandatory downstream FEC. Its upstream options are separately controlled. ITU-T G.9804.2, 2025 amendment, clause 10.3.

The corresponding code rate, decoder behavior, receiver requirements, codeword boundaries, and error floor differ from the RS model. Multiplying the XGS-PON result by five or applying the same BER threshold would not be justified.

There is relevant primary research on flexible upstream FEC for 50G-PON, which supports a separate study with the correct code family and hardware assumptions. *Flexible Upstream FEC for Higher Throughput, Efficiency, and Robustness for 50G PON*, OFC 2022. Any proposed 50G extension should have its own protocol specification, code-level model, measured channel calibration, and baseline.

A simulation program that vendors could use

Simulation is an appropriate first step. Its role is to expose failure modes, quantify conditional gains, and generate testable requirements. The next version should simulate an explicit candidate protocol rather than only increase a service-budget variable.

Preserve the existing result as a reproducibility baseline

Freeze the current source, configuration files, hashes, and raw counters. Keep the clean 7.969792-to-9.081856 Gb/s pair as a regression test. Label it as the original idealized model so later realism does not silently alter the historical comparison.

Add separate switches for downstream and upstream FEC. The optical-study environment currently couples them in a single setting. At minimum, compare coded operation, upstream-only FEC off, decoder bypass without byte reuse, and experimental downstream byte reuse. This separates a supported upstream mechanism from the nonstandard downstream proposal.

Define the experimental wire format before measuring it

Write a bit- and byte-level candidate specification covering information boundaries, parity reuse, control protection, synchronization, mode indicators, and error detection. Define what a transmitter sends and what every receiver does on an unknown, invalid, or corrupted indication.

Specify activation, capability exchange, mode-change timing, timeout behavior, and fallback. Treat a mixed legacy/enhanced population as a separate profile with an explicit compatibility

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

argument. A mode that loses the control channel needed to recover should fail the design review regardless of peak rate.

Replace the dominant abstractions in stages

1. **Serialization:** Model XGEM fragmentation, frame and grant boundaries, residual capacity, Ethernet wire time, real control traffic, and packet-size mixtures.
2. **Coding and errors:** Use a real RS encoder/decoder or a validated equivalent; compare independent errors with measured traces and burst-error processes. Include frame synchronization, header corruption, decoder failure, and miscorrection accounting.
3. **Traffic and scheduling:** Model per-ONU queues, report/grant timing, burst profiles, fairness, assured traffic, background demand, and simultaneous traffic in both directions.
4. **Transport:** Add an established TCP implementation or a validated network simulator, documenting congestion control, round-trip time, receive windows, pacing, and host limits. Retain constant-rate packet tests for isolating link capacity.
5. **Adaptation:** Model measurement noise, stale counters, hysteresis, probe overhead, temporary impairment, and recovery. Compare always-coded, always-uncoded, and adaptive policies on the same traces.

Use paired channel traces and traffic inputs when comparing modes. Report sensitivity to assumptions separately from random variation between seeds. Hundreds of seeds cannot repair a systematically wrong channel model.

Measure the outcomes that matter

The primary outcome should be unique application bytes delivered at a stated loss, delay, and availability objective. Also report offered load, source drops, corrupt attempts, retransmissions, uncorrectable codewords, synchronization failures, queue occupancy, percentile latency, recovery time, and peer-service effects.

Publish the distribution across ONU conditions and traffic loads, not only the strongest subscriber. Show which fraction of measured field paths qualifies, how often the mode falls back, and whether the benefit survives stricter reliability targets. Define these tests before selecting favorable cases.

A phased path for vendors and standards work

Phase 1 Establish a standards compliant control

Start with coded downstream operation and investigate upstream FEC on/off using supported equipment controls. Verify the actual burst profiles, device counters, and transport performance. This provides useful evidence while avoiding an unsupported assumption about downstream framing. It cannot by itself reproduce the claimed symmetric uncoded result.

Exit evidence: reproducible matched measurements, an identified equipment and firmware configuration, and an upstream policy that meets the chosen service objective.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

Phase 2 Specify and simulate an experimental downstream profile

Prepare the candidate wire format, state machines, compatibility rules, error model, and reference implementation described above. Initially restrict it to an isolated all-enhanced port. Compare it with compliant operation under identical traffic and channel traces.

Exit evidence: a reproducible gain after realistic overhead and errors, no accounting violations, bounded recovery behavior, and a clear explanation of every departure from G.9807.1.

Phase 3 Demonstrate the data path on hardware

Implement the proposed mode in an FPGA or other suitably programmable OLT/ONU pair. Demonstrate that the larger information frame is transmitted, parsed, buffered, and delivered correctly. Test transitions under load, malformed indicators, resets, clock drift, and injected errors. Keep standard operation available as the reference and recovery mode.

Exit evidence: bit-exact traces and receiver measurements proving that the extra bytes are useful data, rather than an inference from a configuration flag or application speed test.

Phase 4 Establish the optical operating envelope

Measure both directions independently across received power, temperature, fiber conditions, reflections, coexistence conditions, and component variation. Exercise burst acquisition and loud-to-soft upstream transitions. Record error locations and frame losses alongside BER. Include weak ONUs and service-affecting transients.

Exit evidence: a measured envelope with confidence bounds, sufficient duration for the claimed reliability, documented exclusion conditions, and a fallback policy that works when conditions deteriorate.

Phase 5 Build the interoperability contribution

Take the evidence to the relevant standards participants as a proposed optional extension or profile, not as a claim that the existing requirement can be ignored. ITU-T SG15 Question 2 is responsible for optical systems for fibre access networks and maintains the relevant PON recommendations. ITU-T Q2/15 work programme.

The contribution should include precise proposed normative changes, receiver and transmitter requirements, capability signaling, mandatory behavior for unsupported modes, protected common control, management counters, recovery procedures, and mixed-population test cases. Whether the right vehicle is an amendment or another profile is a matter for the standards process; this review does not assume acceptance or a timetable.

Coordinate corresponding management and interoperability requirements. Broadband Forum's PON TC-layer interoperability test plan provides an existing framework to extend, but current certification should not be represented as covering an unstandardized mode. Broadband Forum technical library, TR-309.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

Exit evidence: agreement on a complete interoperable specification and tests, with independent implementations demonstrating consistent behavior. A single-vendor laboratory success is an input to that process, not a substitute.

Phase 6 Decide whether deployment is economically worthwhile

Compare the net benefit with upstream-only optimization, scheduling improvements, larger frames where appropriate, faster user interfaces, lower split ratios, and higher-rate PON upgrades. Include line-card and ONU replacement, qualification costs, operational complexity, support burden, and any licensing implications.

Exit evidence: a field-trial result showing useful aggregate or subscriber benefit at unchanged service quality, with a manageable installed-base transition. A small gain available only on unusually strong, lightly loaded ports may be technically valid but commercially unattractive.

These phases are decision gates, not a forecast or a commitment to a particular vendor. Stopping after a failed gate would be a valuable research outcome.

Conclusions for Audeamus Labs

The simulation is right about the bandwidth arithmetic and useful as a transparent, reproducible exploration. Its strongest result is a 13.95% same-model improvement under a declared hypothetical frame format and very low assumed errors.

It would be wrong to describe that result as a demonstrated standard XGS-PON upgrade, hardware measurement, lossless offered-load result, full TCP benchmark, or general field gain. The downstream standard, optical calibration, transport model, sharing assumptions, and threshold sensitivity all prevent those interpretations.

The research opportunity remains credible. The literature supports flexible protection and overhead optimization, while prior art shows that downstream byte reuse has a concrete architectural history. The next contribution should be a falsifiable protocol and validation program: define the mode, model its failures, build it, measure it, and then seek an interoperable specification.

Recommended public claim: Audeamus Labs reproduced a 13.95% increase in simulated delivered payload by allowing hypothetical downstream FEC-byte reuse and upstream FEC removal under idealized clean-channel conditions. Current XGS-PON does not permit that downstream mode. A staged simulation, hardware, and standards program is proposed to determine whether any practical deployment can retain the benefit while meeting reliability and interoperability requirements.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

Reproducibility and evidence notes

- **Source materials reviewed:** the supplied *XGS-PON Optical Environment White Paper*, capacity.py, optical_environment.py, simulate_pon.py, experiment manifest, data audit, CSV experiment index, and raw JSON runs. The source manuscript already labels the work as simulation and documents its major assumptions.
- **Independent checks:** code inspection, recalculation of payload rates and queue-drop fractions, continuous-packing bounds, and exact reruns of saved cases 0697 and 0700. This review did not independently rerun all 843 executed cases.
- **Source identity:** capacity.py SHA-256 b799eae42b7218d9fac96f0ad75968f5da5029e247c6c5945500cd707f0897f4; optical_environment.py SHA-256 c9680ebf4d3ae226978cf17b6ebec312fa0c5c31698e793250e8dd58baf79484; simulate_pon.py SHA-256 836fb2d7fd95ab1a610ff649cc1d6db2e9defa65e4e1e8f073ea9c01189b274f.
- **Literature access:** standards and the cited OFC 2017 manuscript were inspected in full or in relevant primary-source sections. The 2025 upstream-optimization and 2026 field-data discussions are limited to accessible publication metadata and abstracts. Flexible-PON and adaptive-decoder references establish related research directions, not validation of Audeamus's exact scenario.
- **Interpretation:** all new break-even, loss-probability, interface-bound, and statistical-exposure calculations in this review are analytical illustrations with their assumptions stated. They are not new optical measurements.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.

References

1. ITU-T. G.9807.1 Amendment 1, May 2025. XGS-PON physical and transmission-convergence requirements, especially Annex C.
2. ITU-T. G Supplement 75, December 2021. Illustrative capacity and latency comparisons, Tables 14 and 16.
3. Khotimsky, D. A. Method and apparatus of downstream forward error correction on-off control in XG-PON1 and NG-PON2 TWDM-PON systems. US20150381313A1, published 2015; subsequently US10044467B2. Prior technical disclosure, not performance validation or legal advice.
4. Brandonisio, N., Porto, S., Carey, D., Ossieur, P., Talli, G., Parsons, N., and Townsend, P. D. Forward Error Correction Analysis for 10Gb/s Burst-Mode Transmission in TDM-DWDM PONs. OFC 2017, Th2A.28. DOI: 10.1364/OFC.2017.Th2A.28.
5. Optimization of Upstream XGS-PON Throughput by Adjusting Burst Preamble Length and Enabling Forward Error Correction. ECOC 2025. Discussion limited to accessible publication information.
6. Simon, G., Chanclou, P., Le Huérou, S., Potet, J., Torillec, B., and Saliou, F. PON Bitrate Optimization Adapting Flexible Header and FEC According to Field Data. OFC 2026, Th1K.6. Abstract-level evidence.
7. Borkowski, R., et al. FLCS-PON an opportunistic 100 Gbit/s flexible PON prototype with probabilistic shaping and soft-input FEC. Journal of Optical Communications and Networking 14, C82, 2022. Related architecture and operator-trial research.
8. Kubo, K., Kariyawasam, A., and Kimura, S. Power saving control method with adaptive FEC decoder in optical access network. IEICE Communications Express 13(10), 401-404, 2024.
9. Bradner, S., and McQuaid, J. RFC 2544 Benchmarking Methodology for Network Interconnect Devices. 1999.
10. Constantine, B., Forget, G., Geib, R., and Schrage, R. RFC 6349 Framework for TCP Throughput Testing. 2011.
11. ITU-T. G.9804.2 Amendment 3, November 2025. Higher-speed PON common transmission-convergence specification.
12. Flexible Upstream FEC for Higher Throughput, Efficiency, and Robustness for 50G PON. OFC 2022, M3G.5. Related research, not a validation of the XGS-PON model.
13. ITU-T. Question 2/15 work programme for optical systems for fibre access networks. Standards responsibility and scope.
14. Broadband Forum. Technical library and PON TC-layer interoperability test plans. Consult the current TR-309 edition when developing an interoperability contribution.

© 2026 Audeamus Group, LLC. All rights reserved. This report is the property of Audeamus Group, LLC and reflects solely its opinions. It is for informational purposes only and must not be relied upon as legal or investment advice. Third-party materials remain the property of their respective owners.