BYTETHIRST

Quarterly AI Sustainability Report

Northwind Software, Inc.

Reporting period	Q1 2026 (1 January – 31 March 2026)
Generated	19 May 2026
Methodology	v2.1 · released 7 May 2026

> **Sample report — fabricated data for demonstration purposes.**
>
> This document is a representative example of a ByteThirst Quarterly AI Sustainability Report. The team name, member identities, session counts, and all impact figures are fabricated. No real user data is reflected anywhere in this report. The methodology citations, ESG-framework alignments (see the Standards & Frameworks Alignment section below), and version pinning are accurate and reflect ByteThirst's published methodology.

Audit-ready edition · PDF Universal Accessibility (PDF/UA-1) tagged · Portable Document Format for Archive Level 3 (PDF/A-3) archival profile · Environmental, Social, and Governance (ESG)-framework aligned.

Executive Summary

For Q1 2026 (1 January – 31 March 2026), ByteThirst estimated impact for **5,640 anonymous Artificial Intelligence (AI) sessions** originating from **8 Northwind Software, Inc. team members** enrolled in ByteThirst Teams. Figures below are mid-point estimates prefixed “est.”; every number on every surface (extension popup, command-line interface (CLI), web dashboard, marketing pages, and this report) links back to the published methodology where every coefficient, source, and known limitation is disclosed.

> **Privacy posture.** ByteThirst observes anonymous traffic patterns from team members enrolled in ByteThirst Teams. Individual identities are not visible to ByteThirst, to ByteThirst Teams administrators, or to this report unless a member has explicitly opted in to identification via their personal ByteThirst account settings. The “Member A–H” labels in the per-member breakdown are persistent anonymous identifiers used for aggregation; they are not employee names. Aggregate team totals are derived from anonymous per-session estimates.

All emissions reported here align with Greenhouse Gas Protocol (GHG Protocol) **Scope 3 Category 1 — Purchased Goods & Services** (cloud-based AI inference is an upstream service). Reporters subject to International Sustainability Standards Board (ISSB) International Financial Reporting Standard S2 (IFRS S2), European Sustainability Reporting Standard E1 (ESRS E1) under the Corporate Sustainability Reporting Directive (CSRD), or the climate disclosure framework CDP (formerly Carbon Disclosure Project) should categorize accordingly.

Metric	Estimate (mid-point)	Methodology & framework
Estimated water consumption	**est. 2,140 L**	methodology v2.1 §water · Global Reporting Initiative (GRI) 303-5 · ESRS E3-4
Estimated energy consumption (upstream)	**est. 1,180 kWh**	methodology v2.1 §energy · GRI 302 (contextual)
Estimated Scope 3 Cat 1 GHG emissions	**est. 520 kg CO_2e**	methodology v2.1 §co₂ · GRI 305-3 · ESRS E1-6 · ISSB IFRS S2 §29
Team members enrolled (anonymous)	**8**	Anonymous by default; persistent labels Member A–H used in the per-member breakdown
AI sessions estimated	**5,640**	≈ 54 sessions/person/week (13-week quarter)
Top 3 platforms by usage	**ChatGPT (38%) · Claude (26%) · Copilot (18%)**	82% of total team usage

> **Reading the numbers.** Every estimate is the baseline mid-point from documented inputs and is prefixed “est.” to signal that it is a model output, not a measurement. The honesty mechanism is the methodology link, not a numeric range adjacent to the estimate — users who want to interrogate the assumptions follow the link; users who want a single readable number get one. See the Methodology section for known limitations and the full source list.

All numbers in this section are mid-point estimates pinned to ByteThirst methodology v2.1 (released 7 May 2026). Read the full methodology at bytethirst.com/methodology.

Methodology

This report is generated using **ByteThirst QueryWeight™ Calculation Methodology v2.1** (released 7 May 2026). All figures are mid-point estimates derived from per-session character-count telemetry combined with published energy and water benchmarks. The full methodology is published at bytethirst.com/methodology and is versioned, public, and open to peer review.

The methodology aligns with the **GHG Protocol Scope 3 Standard, Calculation Guidance §6.4** — specifically the *average-data method* — because per-query physical measurement is infeasible from the client side. Per-token energy and water benchmarks come from the peer-reviewed sources listed below. Methodology principles align with International Organization for Standardization (ISO) standard **14064-1** (relevance, completeness, consistency, accuracy, transparency). Numeric framing follows the United States Federal Trade Commission (FTC) Green Guides expectations for environmental claims (units everywhere, methodology link inline, no consumer-comparable analogies).

How displayed numbers are framed

ByteThirst displays mid-point estimates end-to-end. Every displayed estimate is the baseline mid-point from documented inputs, is prefixed “est.”, and links back to this methodology page where every coefficient, source, and known limitation is published. The honesty mechanism is the methodology link, not a numeric uncertainty range adjacent to the estimate. Users who want to interrogate the assumptions follow the link; users who want a single readable number get one. This rule applies on every surface (extension popup, CLI, web dashboard, marketing pages) and to every section of this report.

Known limitations

1. **Token estimation is approximate.** Character-to-token ratios are averages for English text. Actual tokenization varies by language, content type (code vs. prose), and specific model version. Errors of 10–20% in token estimation are possible.
2. **Model-tier detection is heuristic.** ByteThirst infers the active model from Document Object Model (DOM) elements on each AI platform’s interface. If a platform changes its User Interface (UI), model detection may temporarily misclassify the model tier until we update the extension.
3. **Energy-per-token varies widely.** The energy cost of inference depends on Graphics Processing Unit (GPU) type (H100 vs. A100 vs. TPUv5), batch size, quantization level, and server utilization. Our baseline assumes mid-range conditions, but actual energy consumption for any single query could be 2–3× higher or lower than the displayed mid-point.
4. **Water consumption depends on local climate and cooling technology.** A data center in Iowa using evaporative cooling will consume significantly more water per watt-hour than a data center in Finland using free air cooling. We cannot determine which data center serves any individual query, so we use an industry-average ratio.
5. **Cached and short-circuited responses are not detected.** Some queries may be served from cache or routed to smaller models, consuming far less energy than the baseline mid-point suggests. We have no way to detect this from the client side.
6. **Reasoning-model uncertainty is high.** Models like o1, o3, and o4-mini generate internal chain-of-thought tokens that are not visible to the user. The number of internal tokens can vary from 2× to 50× the visible output length. Our multiplier is a conservative midpoint, but individual queries may vary significantly.
7. **All constants are point-in-time.** The energy efficiency of AI inference is improving rapidly. Our constants are based on data available as of early 2026 and will be updated as new measurements are published.

We address this uncertainty through transparency rather than per-query bands. The methodology page enumerates every coefficient, every published source, every known limitation, and the rationale for every modeling choice. The link itself is the disclosure mechanism.

Per-coefficient citations

Each input coefficient in methodology v2.1 is sourced to a published reference. The table below maps the coefficient to its primary citation; the full source list follows.

Coefficient	Mid-point value	Primary source
Water per Wh (data-center cooling)	0.55 mL/Wh	Li et al. (2023), “Making AI Less Thirsty,” UC Riverside; cross-checked against Google (2025) arXiv:2508.15734.
Energy per token (frontier-tier inference)	0.34 Wh / 1,000 tokens	Jegham et al. (2025), arXiv:2505.09598 “How Hungry is AI?”; cross-checked against Epoch AI (Feb 2025) and Luccioni et al. (FAccT 2023).
Energy per token (efficient-tier inference)	0.12 Wh / 1,000 tokens	Google (Aug 2025), “Environmental Report: AI and Energy Use”; OpenAI (Altman, Aug 2025).
CO_2e per Wh (grid mix, US average)	0.39 g/Wh	US EPA, “eGRID Summary Tables” (2023 data).
Character-to-token ratio (English text)	4.0 chars / token	OpenAI tokenizer documentation; cross-checked against Anthropic public guidance.
Reasoning-model output multiplier	6× visible output	SemiAnalysis (2024), “Inference Cost Analysis”; Couch (2026), “Electricity use of AI coding agents.”
AI code-builder per-session energy	per-tool published estimate	Lovable.dev product documentation; StackBlitz/Bolt.new documentation.

Sources cited

1. Google, “Environmental Report: AI and Energy Use” (August 2025).
2. Altman, S., “AI and Energy” blog post, OpenAI (August 2025).
3. Epoch AI, “Estimating the energy consumption of LLM inference” (February 2025).
4. Jegham, N. et al., “Energy Consumption of Large Language Models: A Systematic Benchmark,” arXiv (May 2025).
5. Luccioni, A. et al., “Power Hungry Processing: Watts Driving the Cost of AI Deployment?” FAccT (2023).
6. US EPA, “eGRID Summary Tables” (2023 data).
7. Li, P. et al., “Making AI Less Thirsty,” UC Riverside (2023).
8. SemiAnalysis, “Inference Cost Analysis” (2024).
9. Couch, S.P. (2026), “Electricity use of AI coding agents” — per-token energy rates for agentic AI coding sessions.
10. Lovable.dev product documentation — AI code builder architecture and consumption models.
11. StackBlitz/Bolt.new documentation — WebContainers architecture; Claude 3.5 Sonnet integration.
12. Google (2025), “Measuring the Environmental Impact of Delivering AI at Google Scale” (arXiv:2508.15734).
13. Jegham, N. et al. (2025), “How Hungry is AI? Benchmarking Energy, Water, and Carbon Footprint of LLM Inference” (arXiv:2505.09598).

Version pinning

All numbers in this report were generated under methodology v2.1; the methodology may evolve, but this report is locked to v2.1. Future ByteThirst reports may use a newer methodology version, but historical reports always remain pinned to the version active at generation time. This prevents retroactive changes to numbers a stakeholder has already seen, and supports the audit-trail requirements of GHG Protocol §3.5 (consistency) and ESRS E1 §AR-7 (traceability).

All numbers in this section are mid-point estimates pinned to ByteThirst methodology v2.1 (released 7 May 2026). Read the full methodology at bytethirst.com/methodology.

Standards & Frameworks Alignment

The table below maps each report section to the global ESG framework requirement it satisfies, partially satisfies, or explicitly does not address. This is provided to help sustainability, legal, and assurance teams integrate ByteThirst data into an existing disclosure workflow.

Frameworks covered: **GHG Protocol** (Corporate Standard + Scope 3 Standard); **GRI Standards** (302 Energy, 303 Water, 305 Emissions); **ISSB IFRS S2** Climate-related Disclosures; **ESRS** E1 (Climate), E3 (Water & marine resources), E5 (Resource use) under EU CSRD; **CDP** Climate Change disclosure; **IPCC AR6** reporting conventions; and **ISO 14064-1** GHG quantification principles.

Framework	Requirement / clause	How this report aligns	Sections
GHG Protocol	Scope 3 Cat 1 — Purchased Goods & Services	AI inference is a purchased cloud service from upstream providers (OpenAI, Anthropic, Google).	§2 Executive Summary (CO_2e total) · §5 Per-Platform · §7 Boundary
GHG Protocol	Scope 3 calculation method — average-data method	Per-query physical measurement is infeasible; methodology uses published per-token energy/water benchmarks. Aligns with GHG Protocol Scope 3 Calculation Guidance §6.4.	§3 Methodology · §3 Sources cited
GRI	GRI 302-1 — Energy consumption within the organization	Note: AI inference energy is upstream (Scope 3); for primary GRI 302-1 disclosure, this is contextual data, not the headline number.	§2 (energy total) · §5 (per-platform energy)
GRI	GRI 303-3 / 303-5 — Water withdrawal / consumption	Methodology reports water consumption (the consumed-and-not-returned portion). Aligns with GRI 303-5 conceptual framing.	§2 (water total) · §5 (per-platform water) · §7 Boundary
GRI	GRI 305-3 — Other indirect (Scope 3) GHG emissions	CO_2e total maps to GRI 305-3 reporting. Methodology disclosure aligns with GRI 305-1/2/3 calculation-method disclosure expectations.	§2 · §5 · §3 Methodology
ISSB / IFRS S2	§29 — Cross-industry metrics (Scope 3 GHG)	Provides the underlying Scope 3 Cat 1 figure with methodology disclosure suitable for ISSB-aligned reporting.	§2 · §3 · §7
ISSB / IFRS S2	§13–15 — Governance / Strategy / Risk Management	This report does not address governance; it is an input dataset to the customer's own governance disclosure.	Not in scope of this report
ESRS (CSRD)	ESRS E1-6 — Gross Scope 3 GHG emissions	CO_2e total + per-platform breakdown supports ESRS E1-6 disclosure for organizations subject to EU CSRD.	§2 · §5
ESRS (CSRD)	ESRS E3-4 — Total water consumption	Water-consumption total supports ESRS E3-4 disclosure with explicit methodology and boundary documentation.	§2 · §5
ESRS (CSRD)	ESRS E5 — Resource use	Per-platform breakdown supports a value-chain-resource-intensity narrative for ESRS E5.	§5 · §7
CDP	C6.5 — Scope 3 emissions calculation methodology disclosure	Methodology, sources, and known limitations are explicitly disclosed.	§3 Methodology
CDP	C6.10 — Scope 3 emissions intensity	Per-member breakdown enables intensity calculation (kg CO_2e per FTE-quarter).	§4 Per-Member Breakdown
IPCC AR6	Working Group III reporting conventions	Methodology, boundary, and known limitations follow the WG III convention of disclosing modeling assumptions and source provenance for each input coefficient.	§3 Methodology · §7 Boundary
ISO 14064-1	Quantification principles (relevance, completeness, consistency, accuracy, transparency)	Methodology, sources, boundary, and limitations are documented to support an ISO 14064-1 inventory submission.	§3 · §7

What this report does not claim

- This is **not a third-party-assured disclosure**. It is input data suitable for inclusion in your own assured ESG disclosure. See “Assurance & next steps for the ESG team” below.
- This report does not address governance, strategy, or risk-management disclosures (ISSB IFRS S2 §13–§28; ESRS E1 §16–§26). Those are entity-level disclosures the reporting organization must compose.
- This report does not establish a science-based target nor a transition plan. ByteThirst data can inform such targets but the targets themselves are an organizational decision.

All numbers in this section are mid-point estimates pinned to ByteThirst methodology v2.1 (released 7 May 2026). Read the full methodology at bytethirst.com/methodology.

Per-Member Breakdown

Per-member figures are anonymous aggregations grouped under persistent pseudonymous labels (Member A–H). ByteThirst does not link these labels to identified employees unless a member has explicitly opted in to identification via their personal account settings. **All cells are mid-point estimates** — ByteThirst displays mid-point values end-to-end on every surface, and this table follows the same rule. Coefficients and known limitations are documented in the Methodology section above and on the live methodology page.

Per-member breakdown supports the CDP climate disclosure framework reference C6.10 (Scope 3 emissions intensity per Full-Time Equivalent (FTE)) and provides internal cost-attribution data for finance and people teams. It is not intended for external individual disclosure.

Member	Sessions	Water (L, mid)	Energy (kWh, mid)	CO_2e (kg, mid)	Primary platforms
Member A	1,180	est. 448	est. 247	est. 109	ChatGPT, Copilot
Member B	945	est. 359	est. 198	est. 87	Claude, ChatGPT
Member C	820	est. 311	est. 172	est. 76	ChatGPT, Claude
Member D	705	est. 267	est. 147	est. 65	Copilot, ChatGPT
Member E	640	est. 243	est. 134	est. 59	Claude, Gemini
Member F	525	est. 199	est. 110	est. 48	ChatGPT, Perplexity
Member G	455	est. 173	est. 95	est. 42	Claude, Lovable.dev
Member H	370	est. 140	est. 77	est. 34	ChatGPT, Mistral
Team total	**5,640**	**est. 2,140**	**est. 1,180**	**est. 520**	**All 14 supported platforms**

Column totals are anchored to the canonical team mid-point shown on the Executive Summary; per-member cells are allocated to sum exactly to that total via largest-remainder rounding so the team total cell matches §2 to the unit.

All numbers in this section are mid-point estimates pinned to ByteThirst methodology v2.1 (released 7 May 2026). Read the full methodology at bytethirst.com/methodology.

Per-Platform Breakdown

Estimated impact attributed to each of the 14 AI platforms ByteThirst supports. Percentage column reflects share of total estimated sessions for the team in the period; impact columns show the **mid-point estimate per platform**. ByteThirst displays mid-point values end-to-end; readers who want to interrogate the underlying coefficients follow the methodology link.

Platform	% of usage	Water (L, mid)	Energy (kWh, mid)	CO_2e (kg, mid)
ChatGPT	38.0%	est. 813	est. 448	est. 198
Claude	26.0%	est. 556	est. 307	est. 135
Copilot	18.0%	est. 385	est. 212	est. 94
Gemini	7.0%	est. 150	est. 83	est. 36
Perplexity	4.0%	est. 86	est. 47	est. 21
Lovable.dev	3.0%	est. 64	est. 35	est. 16
Mistral	2.0%	est. 43	est. 24	est. 10
Poe	1.0%	est. 21	est. 12	est. 5
You.com	0.5%	est. 11	est. 6	est. 3
HuggingChat	0.3%	est. 6	est. 4	est. 2
Figma AI	0.1%	est. 2	est. 1	est. 1
Bolt.new	0.1%	est. 2	est. 1	est. 1
NotebookLM	0.0%	est. 1	est. 0	est. 0
Google AI Studio	0.0%	est. 1	est. 0	est. 0
Total	**100.0%**	**est. 2,140**	**est. 1,180**	**est. 520**

Cell values rounded to the nearest unit; column totals are the canonical team mid-points shown in the Executive Summary. Percentages may not sum to exactly 100.0% due to rounding.

All numbers in this section are mid-point estimates pinned to ByteThirst methodology v2.1 (released 7 May 2026). Read the full methodology at bytethirst.com/methodology.

Quarter-over-Quarter Trend

Comparison of Q1 2026 (1 January – 31 March 2026) against Q4 2025 (1 October – 31 December 2025). All values are mid-point estimates per the methodology section. The chart below is a Scalable Vector Graphics (SVG) drawing with embedded title and description for screen readers; the same data is reproduced in tabular form immediately below the chart so the report is readable without seeing the chart.

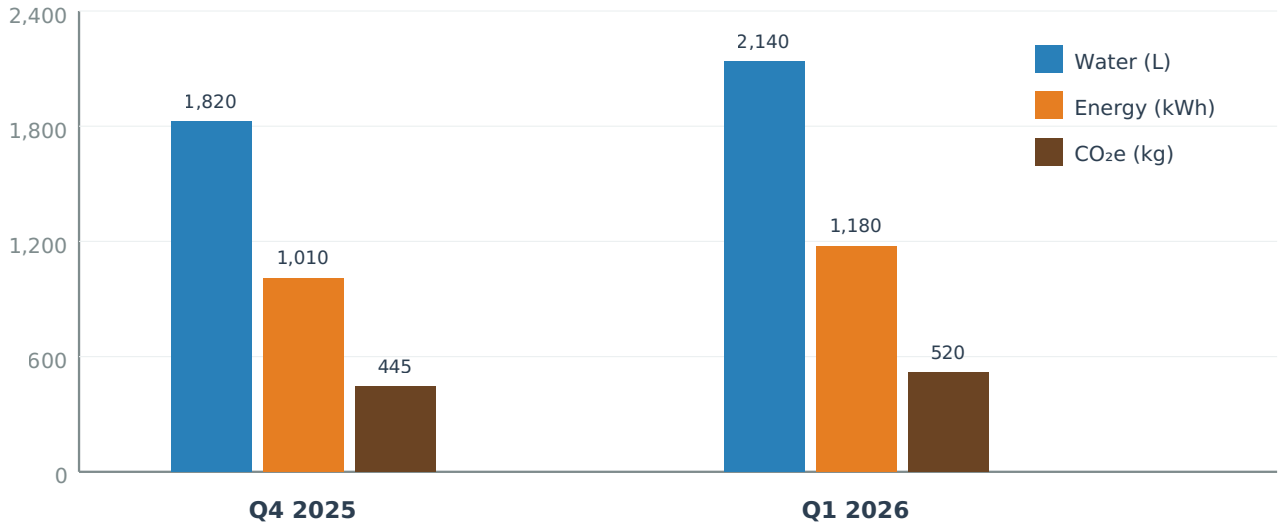

Estimated quarterly totals (mid-point) — Q4 2025 vs. Q1 2026

Metric	Q4 2025 (prior)	Q1 2026 (this report)	Δ
Water (L)	est. 1,820	est. 2,140	+320 (+18%)
Energy (kWh)	est. 1,010	est. 1,180	+170 (+17%)
CO_2e (kg)	est. 445	est. 520	+75 (+17%)

Quarter-over-quarter change is consistent with team headcount growth (7 → 8 members) and a new product-launch cycle in February. Recommendations below identify the top three workflow patterns driving the increase.

All numbers in this section are mid-point estimates pinned to ByteThirst methodology v2.1 (released 7 May 2026). Read the full methodology at bytethirst.com/methodology.

Boundary Disclosure

This report is a partial environmental disclosure. The boundaries below are stated explicitly so the reader can integrate this report into a broader Scope 1/2/3 disclosure under their chosen framework.

Organizational boundary

This report covers AI usage by employees of Northwind Software, Inc. enrolled in ByteThirst Teams during the reporting period. The reporting organization should apply this data to whichever consolidation approach it has adopted — equity share, financial control, or operational control — under GHG Protocol §3.4. ByteThirst does not assert any consolidation approach on the reporter's behalf.

Operational boundary & GHG Protocol Scope categorization

Cloud-based AI inference is an upstream cloud service. For the reporting organization, the emissions attributable to AI usage are **Scope 3**, primarily **Category 1 (Purchased Goods & Services)**. In some contractual arrangements (long-term enterprise commitments to a single AI provider), an argument can also be made for **Category 8 (Upstream Leased Assets)**; the reporting organization should choose the categorization most consistent with how it categorizes other cloud services. Whichever Scope 3 category is chosen, this report's CO_2e figure is the input.

Reporting boundary — what this report covers

- Estimated energy consumed by the AI model's inference computation (GPU/TPU processing on the provider's serving infrastructure).
- Estimated direct water consumed for data-centre cooling during inference (corresponds to GRI 303-5 "Water consumption" — withdrawn and not returned).
- Estimated CO_2e emitted from electricity generation powering inference hardware (using United States Environmental Protection Agency (EPA) Emissions & Generation Resource Integrated Database (eGRID) 2023 grid-mix baselines per source #6).
- Anonymous AI-platform sessions estimated from the 14 ByteThirst-supported AI platforms (see Per-Platform Breakdown), originating from the 8 team members enrolled in ByteThirst Teams during the reporting period.

What this report does not cover

- **Physical infrastructure:** laptops, monitors, peripherals, office equipment. (These fall under Scope 2 [purchased electricity] or Scope 3 Cat 11 [use of sold products] depending on framing; ByteThirst does not duplicate that work.)
- **Offline AI use:** locally-run models (Llama on a laptop, Ollama, on-device inference) are not detected by the ByteThirst extension/CLI and are not in scope.
- **Internal model training:** this report estimates inference only. Fine-tuning or training is not represented here. Where the reporter performs training, that footprint falls under Scope 2 (electricity) or Scope 3 Cat 8 (upstream leased compute).
- **Embodied carbon of devices:** manufacturing, shipping, and disposal of laptops/phones/servers are not included. Consult IT-department lifecycle assessment.
- **Network transmission:** Internet Service Providers (ISPs), Content Delivery Networks (CDNs), and routers between user and AI provider are not measured. These are typically negligible at this scale but, where material, fall under Scope 3 Cat 4.
- **Upstream water in the energy supply chain:** power-plant cooling, fuel extraction, etc. (Out of scope per GRI 303-5 conceptual boundary.)
- **Non-inference server operations at the AI provider** (load balancing, logging, storage of conversation history). The provider's published per-query benchmarks bake these in conceptually, but they are not itemized here.

Methodology version pinning

This report’s numbers were generated under **methodology v2.1 (released 7 May 2026)**. Pinning ensures that if the methodology evolves in a future version, this report’s numbers do not retroactively shift. The version active at generation time is named on the cover, in this section, in the Methodology section, and in the page footer of every section. This pinning supports GHG Protocol §3.5 (consistency) and ESRS E1 §AR-7 (traceability) requirements.

All numbers in this section are mid-point estimates pinned to ByteThirst methodology v2.1 (released 7 May 2026). Read the full methodology at bytethirst.com/methodology.

Recommendations Summary

Three highest-impact workflow patterns identified during Q1 2026 (1 January – 31 March 2026), with suggested swaps and estimated savings. Suggested swaps are **workflow-level recommendations only** — ByteThirst does not see prompt content. Estimated savings are computed under methodology v2.1 from the current vs. swap model-tier energy and water multipliers.

#1. Code review with GPT-4 (long sessions)

Observed pattern: 412 sessions averaging 8,200 characters in / 5,400 characters out; approximately 3.1× the team's median for code-review-shaped sessions.

Suggested swap: Move routine code-review sessions to GPT-4o-mini for first-pass; reserve GPT-4 for complex logic review.

Estimated savings (mid-point): approximately 72% lower impact on this workflow → roughly est. 154 L water, est. 85 kWh energy, est. 38 kg CO_2e per quarter at the current pace.

Per ADR 0005c, the recommendation-card surface displays the mid-point only; underlying coefficients and known limitations are documented on the methodology page linked from every estimate.

#2. Long-context Claude sessions for documentation

Observed pattern: 287 sessions averaging 18,400 characters in / 12,100 characters out on Claude Opus; approximately 2.4× team median for doc-shaped sessions.

Suggested swap: Use Claude Haiku for doc drafting; promote to Opus only for final structural review.

Estimated savings (mid-point): approximately 65% lower impact on this workflow → roughly est. 98 L water, est. 54 kWh energy, est. 24 kg CO_2e per quarter at the current pace.

Per ADR 0005c, the recommendation-card surface displays the mid-point only; underlying coefficients and known limitations are documented on the methodology page linked from every estimate.

#3. Repetitive Copilot prompts on similar boilerplate

Observed pattern: 194 sessions in 7 days with the same length-bucket and code signature; high repeat-shape score per the Type 3 cache-candidate detector.

Suggested swap: Adopt a local snippet library or template repository for the recurring boilerplate; route only novel cases to Copilot.

Estimated savings (mid-point): approximately 58% lower impact on this workflow → roughly est. 72 L water, est. 40 kWh energy, est. 18 kg CO_2e per quarter at the current pace.

Per ADR 0005c, the recommendation-card surface displays the mid-point only; underlying coefficients and known limitations are documented on the methodology page linked from every estimate.

Aggregate savings if all three are adopted

est. 324 L water + est. 179 kWh energy + est. 80 kg CO_2e per quarter — equivalent to roughly a 15% reduction on water, 15% on energy, and 15% on CO_2e against this quarter's totals.

Adoption is voluntary; ByteThirst recommendations persist as user preferences and surface as toasts on the recommended platform on next visit.

Estimated savings are computed under methodology v2.1 from the current vs. swap model-tier energy and water multipliers. All numbers in this section are mid-point estimates pinned to ByteThirst methodology v2.1 (released 7 May 2026). Read the full methodology at bytethirst.com/methodology.

Assurance & Next Steps for the ESG Team

Assurance status of this report

This report is suitable as input data to a corporate ESG disclosure but is **not itself audited or third-party assured**. ByteThirst publishes the methodology, sources, and known limitations to support reasonable-assurance review by the reporter's chosen assurance provider, but neither ByteThirst nor this report itself is a substitute for that engagement.

Specifically:

- **EU CSRD (Article 8):** sustainability information published under CSRD requires *independent limited assurance* initially, transitioning to *reasonable assurance* over time. This report's data, methodology, and version pinning are designed to support that engagement, but the assurance opinion is a separate workstream the reporter must commission.
- **ISSB IFRS S2:** assurance requirements vary by adopting jurisdiction. The reporter should consult its local regulator and chosen assurance provider on materiality and audit readiness.
- **United States Securities and Exchange Commission (SEC) Climate Disclosure Rules & California Senate Bills 253 and 261 (SB 253/261):** requirements continue to evolve. Reporters subject to these regimes should consult counsel and assurance providers on the appropriate level of disclosure and documentation.
- **CDP submission:** the data and methodology disclosure here support CDP Climate Change submission Sections C5 (emissions methodology), C6 (Scope 3), and C8 (energy).

Next steps for the reporter's ESG team

1. **Categorize the CO_2e figure** as Scope 3 Category 1 (Purchased Goods & Services) — or Category 8 (Upstream Leased Assets) for long-term enterprise AI commitments — under your chosen GHG Protocol consolidation approach.
2. **Pin the methodology version** in your audit trail. This report uses methodology v2.1; a future report under v2.1 or later will use updated constants but will not retroactively change this report's numbers.
3. **Annual rollup:** aggregate the four quarterly reports for your fiscal year. Quarterly granularity is provided to support quarterly internal reporting and the GRI 305 / ESRS E1 requirement to disclose annual totals.
4. **Cross-reference with your assurance provider** on the suitability of the methodology and boundary statements for your chosen disclosure framework. Provide them with this report and the live methodology page (bytethirst.com/methodology).
5. **If subject to a regulated disclosure regime** (CSRD, SEC Climate Rules, California SB 253/261, the United Kingdom (UK) Sustainability Disclosure Requirements (SDR), Singapore Exchange Securities Trading (SGX-ST) listing rules, etc.), engage your assurance provider before incorporating this data into the regulated filing. The methodology, sources, and limitations published with ByteThirst are designed to support such review but cannot substitute for it.

All numbers in this report are mid-point estimates pinned to ByteThirst methodology v2.1 (released 7 May 2026). Read the full methodology at bytethirst.com/methodology.

Acronyms Used in This Report

For reader convenience, every acronym used in the body of this report is spelled out at first occurrence. This appendix consolidates them as a glossary for cross-reference.

Acronym	Spelled out	Notes
ADR	**Architecture Decision Record**	Internal ByteThirst document type; ADR 0007 governs this report's format
AI	**Artificial Intelligence**	Refers throughout to the AI platforms ByteThirst supports
AR-7	**Application Requirement 7**	ESRS-internal clause numbering
CDP	**CDP (formerly Carbon Disclosure Project)**	Voluntary climate disclosure platform
CDN	**Content Delivery Network**	Cited under out-of-scope network transmission
CLI	**Command-Line Interface**	One of the ByteThirst surfaces where mid-point estimates are displayed
CO_2e	**Carbon dioxide equivalent**	Standard unit for aggregating greenhouse gases by global warming potential
CSRD	**Corporate Sustainability Reporting Directive**	European Union directive (2022/2464); implemented via ESRS
DOM	**Document Object Model**	Web-page structure ByteThirst's extension reads for model-tier inference
eGRID	**Emissions & Generation Resource Integrated Database**	EPA dataset of US power-plant emissions (2023 data used here)
EPA	**Environmental Protection Agency**	United States federal agency
ESG	**Environmental, Social, and Governance**	Umbrella term for non-financial corporate disclosures
ESRS	**European Sustainability Reporting Standards**	Mandatory standards under the EU CSRD; this report aligns with E1 (Climate), E3 (Water), E5 (Resource use)
EU	**European Union**	
FTC	**Federal Trade Commission**	United States agency; cited for Green Guides framing of environmental claims
FTE	**Full-Time Equivalent**	Used as the denominator for emissions intensity per worker
GHG	**Greenhouse Gas**	The GHG Protocol is the foundational corporate-emissions accounting framework
GPU	**Graphics Processing Unit**	Primary inference hardware for most AI providers
GRI	**Global Reporting Initiative**	This report aligns with GRI 302 (Energy), 303 (Water), 305 (Emissions)
IFRS	**International Financial Reporting Standards**	This report aligns with IFRS S2 (Climate-related Disclosures) issued by the ISSB
IPCC	**Intergovernmental Panel on Climate Change**	Sixth Assessment Report (AR6) cited for working-group reporting conventions
ISO	**International Organization for Standardization**	Standards 14064-1 (GHG quantification) and 14289-1 (PDF/UA-1) referenced here
ISP	**Internet Service Provider**	Cited under out-of-scope network transmission
ISSB	**International Sustainability Standards Board**	Issues IFRS Sustainability Disclosure Standards
PDF/A-3	**Portable Document Format for Archive Level 3**	ISO 19005-3 archival profile for long-term records retention
PDF/UA-1	**PDF Universal Accessibility — version 1**	ISO 14289-1; tagged accessible-PDF standard
SB 253/261	**California Senate Bills 253 and 261**	California's climate disclosure laws (Climate Corporate Data Accountability Act and Climate-Related Financial Risk Act)
SDR	**Sustainability Disclosure Requirements**	United Kingdom regime
SEC	**Securities and Exchange Commission**	United States regulator; Climate Disclosure Rules adopted then partially stayed
SGX-ST	**Singapore Exchange Securities Trading**	Singapore-listed-issuer disclosure regime
SVG	**Scalable Vector Graphics**	Vector chart format used for the quarter-over-quarter chart

Acronym	Spelled out	Notes
TPU	**Tensor Processing Unit**	Google’s custom AI inference hardware
UI	**User Interface**	
UK	**United Kingdom**	
US	**United States**	

All numbers in this report are mid-point estimates pinned to ByteThirst methodology v2.1 (released 7 May 2026). Read the full methodology at bytethirst.com/methodology.