

ByteThirst

Quarterly AI Sustainability Report

Northwind Software, Inc.

Q1 2026 (1 January – 31 March 2026)

Generated 8 April 2026

Methodology v2.0 · released 3 March 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.

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**. The figures below are estimates with explicit uncertainty ranges per the methodology section.

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 with range)	Methodology & framework
Estimated water consumption	2,140 L (low: 1,284 / high: 2,996)	±40% per methodology v2.0 §water · Global Reporting Initiative (GRI) 303-5 · ESRS E3-4
Estimated energy consumption (upstream)	1,180 kWh (low: 826 / high: 1,534)	±30% per methodology v2.0 §energy · GRI 302 (contextual)
Estimated Scope 3 Cat 1 GHG emissions	520 kg CO₂e (low: 338 / high: 702)	±35% per methodology v2.0 §co ₂ · 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

All numbers in this section pinned to ByteThirst methodology v2.0 (released 3 March 2026). Estimates carry explicit ± ranges per FTC Green Guides framing and IPCC AR6 uncertainty practice. Read the full methodology at bytethirst.com/methodology.

Methodology

This report is generated using **ByteThirst QueryWeight™ Calculation Methodology v2.0** (released 3 March 2026). All figures are 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/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). Uncertainty is expressed per the **Intergovernmental Panel on Climate Change (IPCC) Sixth Assessment Report (AR6) likelihood/uncertainty framework**: ranges below are stated as \pm relative bands. Numeric framing additionally follows the United States Federal Trade Commission (FTC) Green Guides expectations for environmental claims (paired ranges, no consumer-comparable analogies, methodology link inline).

Uncertainty ranges

Every numeric estimate in this report carries an explicit \pm range. The ranges reflect known uncertainty in the underlying physical measurements and inference-stack heterogeneity:

Metric	Range	Bounded by	IPCC AR6 confidence
Water	$\pm 40\%$	Data-centre cooling architecture (free-air vs. evaporative; Iowa vs. Finland)	Medium confidence
Energy	$\pm 30\%$	GPU type, batch size, quantization, server utilization	Medium-to-high confidence (aggregated quarterly totals)
CO ₂ e	$\pm 35\%$	Grid mix at the serving region (eGRID 2023 baselines)	Medium confidence

Known limitations

- Token estimation is approximate.** Character-to-token ratios are averages for English text. Errors of 10–20% in token estimation are possible depending on language, content type (code vs. prose), and specific model version.
- Model-tier detection is heuristic.** ByteThirst infers the active model from Document Object Model (DOM) elements on each platform. User Interface (UI) changes may temporarily misclassify tier until the extension updates.
- Energy-per-token varies widely.** Inference energy depends on Graphics Processing Unit (GPU) and Tensor Processing Unit (TPU) type, batch size, quantization, and server utilization. Single-query energy could be 2–3× higher or lower than the baseline assumption.
- Water consumption depends on local cooling.** Iowa evaporative cooling vs. Finland free-air cooling differ significantly per Wh. We use an industry-average ratio because we cannot determine which data centre serves any individual query.

5. **Cached and short-circuited responses are not detected.** Some queries are served from cache or routed to smaller models, consuming far less than estimated. Not detectable from the client side.
6. **Reasoning-model uncertainty is high.** Chain-of-thought tokens for o1/o3/o4-mini are invisible to the user and can vary 2× to 50× the visible output length. Our multiplier is a conservative midpoint.
7. **All constants are point-in-time.** AI inference efficiency is improving rapidly. Constants are based on data available as of early 2026 and will be updated as new measurements are published.

Version pinning

All numbers in this report were generated under methodology v2.0; the methodology may evolve, but this report is locked to v2.0. 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).

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).

All numbers in this section pinned to ByteThirst methodology v2.0 (released 3 March 2026). Estimates carry explicit ± ranges per FTC Green Guides framing and IPCC AR6 uncertainty practice. Read the full methodology at bytethirst.com/methodology.

Standards & Frameworks Alignment

The following table maps each report section to the global ESG framework requirement it satisfies, partially satisfies, or explicitly does not address. This is provided to help your sustainability, legal, and assurance teams integrate ByteThirst data into your 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 uncertainty terminology, 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 ₂ e 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 / Water 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 ₂ e 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 and uncertainty 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

Framework	Requirement / clause	How this report aligns	Sections
ESRS (CSRD)	ESRS E1-6 — Gross Scope 3 GHG emissions	CO ₂ e 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 + uncertainty bracket supports ESRS E3-4 disclosure.	§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, uncertainty ranges, and limitations are explicitly disclosed.	§3 Methodology
CDP	C6.10 — Scope 3 emissions intensity	Per-member breakdown enables intensity calculation (kg CO ₂ e per FTE-quarter).	§4 Per-Member Breakdown
IPCC AR6	Annex A — Uncertainty language	Ranges expressed as ± relative bands. By IPCC AR6 likelihood scale, water (±40%) and CO ₂ e (±35%) are 'medium confidence' estimates; energy (±30%) is approaching 'high confidence' for aggregated quarterly totals.	§3 Uncertainty ranges
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 pinned to ByteThirst methodology v2.0 (released 3 March 2026). Estimates carry explicit ± ranges per FTC Green Guides framing and IPCC AR6 uncertainty practice. 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. Per-member figures are estimates with the same uncertainty ranges as the team total (water $\pm 40\%$, energy $\pm 30\%$, CO₂e $\pm 35\%$). Ranges are not repeated per cell to keep the table scannable; multiply the mid by 0.6/1.4, 0.7/1.3, 0.65/1.35 respectively to derive per-member ranges.

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)	Energy (kWh)	CO ₂ e (kg)	Primary platforms
Member A	1,180	448	247	109	ChatGPT, Copilot
Member B	945	359	198	87	Claude, ChatGPT
Member C	820	311	172	76	ChatGPT, Claude
Member D	705	268	148	65	Copilot, ChatGPT
Member E	640	243	134	59	Claude, Gemini
Member F	525	199	110	49	ChatGPT, Perplexity
Member G	455	173	95	42	Claude, Lovable.dev
Member H	370	140	77	34	ChatGPT, Mistral
Team total	5,640	2,141	1,181	521	All 14 supported platforms

All numbers in this section pinned to ByteThirst methodology v2.0 (released 3 March 2026). Estimates carry explicit \pm ranges per FTC Green Guides framing and IPCC AR6 uncertainty practice. 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 mid-point estimate followed by (low / high) range in parentheses.

Platform	% of usage	Water (L) — mid (low / high)	Energy (kWh) — mid (low / high)	CO ₂ e (kg) — mid (low / high)
ChatGPT	38.0%	814 (488 / 1,140)	449 (314 / 583)	198 (129 / 267)
Claude	26.0%	556 (334 / 778)	307 (215 / 399)	135 (88 / 182)
Copilot	18.0%	385 (231 / 539)	212 (148 / 276)	94 (61 / 127)
Gemini	7.0%	150 (90 / 210)	83 (58 / 108)	36 (23 / 49)
Perplexity	4.0%	86 (51 / 120)	47 (33 / 61)	21 (14 / 28)
Lovable.dev	3.0%	64 (38 / 90)	35 (25 / 46)	16 (10 / 21)
Mistral	2.0%	43 (26 / 60)	24 (17 / 31)	10 (7 / 14)
Poe	1.0%	21 (13 / 30)	12 (8 / 15)	5 (3 / 7)
You.com	0.5%	11 (6 / 15)	6 (4 / 8)	3 (2 / 4)
HuggingChat	0.3%	6 (4 / 9)	4 (2 / 5)	2 (1 / 2)
Figma AI	0.1%	2 (1 / 3)	1 (1 / 2)	1 (0 / 1)
Bolt.new	0.1%	2 (1 / 3)	1 (1 / 2)	1 (0 / 1)
NotebookLM	0.0%	0 (0 / 1)	0 (0 / 0)	0 (0 / 0)
Google AI Studio	0.0%	0 (0 / 1)	0 (0 / 0)	0 (0 / 0)
Total	100.0%	2,140	1,181	522

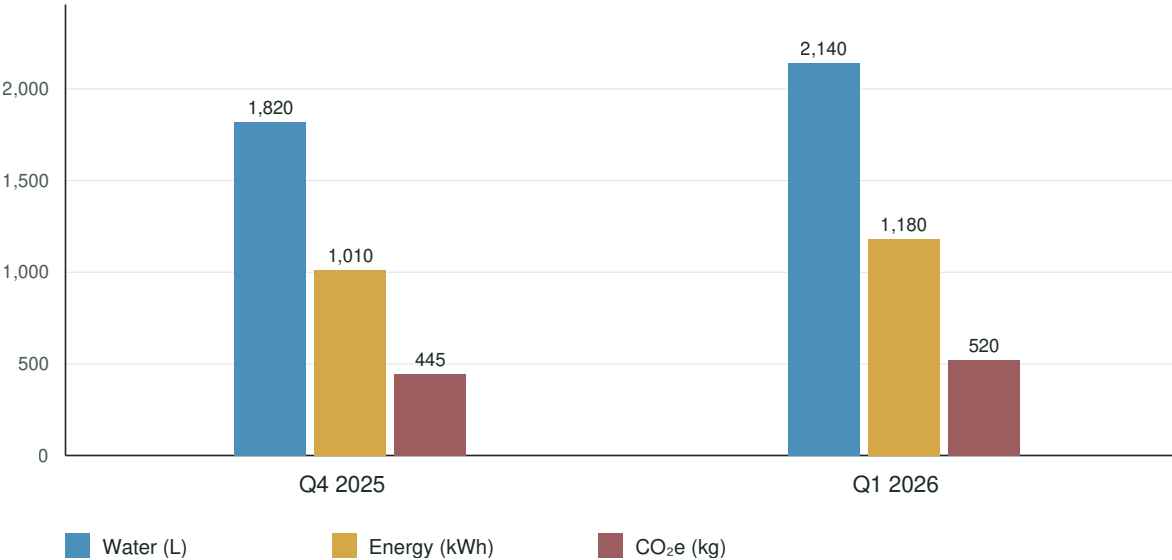
Note: percentages may not sum to exactly 100.0% due to rounding. Total impact rounds to within 1% of the team total in the executive summary; small differences are caused by per-platform rounding.

All numbers in this section pinned to ByteThirst methodology v2.0 (released 3 March 2026). Estimates carry explicit ± ranges per FTC Green Guides framing and IPCC AR6 uncertainty practice. 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; ranges per the methodology section. The chart is an SVG (vector) drawing with embedded title and description for screen readers; the same data is reproduced in tabular form immediately below the chart.

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



Quarter-over-quarter trend chart. Blue bars: water in litres. Yellow bars: energy in kilowatt-hours. Brown bars: CO₂-equivalent in kilograms.

Metric	Q4 2025 (prior)	Q1 2026 (this report)	Δ
Water (L)	1,820	2,140	+320 (+18%)
Energy (kWh)	1,010	1,180	+170 (+17%)
CO ₂ e (kg)	445	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 pinned to ByteThirst methodology v2.0 (released 3 March 2026). Estimates carry explicit ± ranges per FTC Green Guides framing and IPCC AR6 uncertainty practice. 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 CO₂e 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 CO₂e 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: ISPs, 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 v2.0 (released 3 March 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, and in the page footer of every numeric section. This pinning supports GHG Protocol §3.5 (consistency) and ESRS E1 §AR-7 (traceability) requirements.

All numbers in this section pinned to ByteThirst methodology v2.0 (released 3 March 2026). Estimates carry explicit \pm ranges per FTC Green Guides framing and IPCC AR6 uncertainty practice. 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 v2.0 from the current vs. swap model-tier energy/water multipliers.

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

Observed pattern: 412 sessions averaging 8,200 chars in / 5,400 chars out; ~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 **154 L water, 85 kWh energy, 38 kg CO₂e** per quarter at current pace. Per ADR 0005c, the recommendation-card surface displays the mid-point only; underlying ± ranges are preserved internally and remain accessible via the methodology link.

#2. Long-context Claude sessions for documentation

Observed pattern: 287 sessions averaging 18,400 chars in / 12,100 chars out on Claude Opus; ~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 **98 L water, 54 kWh energy, 24 kg CO₂e** per quarter at current pace. Per ADR 0005c, the recommendation-card surface displays the mid-point only; underlying ± ranges are preserved internally and remain accessible via the methodology link.

#3. Repetitive Copilot prompts on similar boilerplate

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

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

Estimated savings (mid-point): approximately 58% lower impact on this workflow → roughly **72 L water, 40 kWh energy, 18 kg CO₂e** per quarter at current pace. Per ADR 0005c, the recommendation-card surface displays the mid-point only; underlying ± ranges are preserved internally and remain accessible via the methodology link.

Aggregate savings if all three are adopted

324 L water + 179 kWh energy + 80 kg CO₂e per quarter — equivalent to a 15% reduction on water, 15% on energy, and 15% on CO₂e 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.

All numbers in this section pinned to ByteThirst methodology v2.0 (released 3 March 2026). Estimates carry explicit ± ranges per FTC Green Guides framing and IPCC AR6 uncertainty practice. 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 uncertainty ranges 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 CO₂e 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 v2.0; 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, uncertainty ranges, 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 section pinned to ByteThirst methodology v2.0 (released 3 March 2026). Estimates carry explicit ± ranges per FTC Green Guides framing and IPCC AR6 uncertainty practice. 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
AI	Artificial Intelligence	Refers throughout to the AI platforms ByteThirst supports
AR-7	Application Requirement 7	ESRS-internal clause numbering
AR6	Sixth Assessment Report	IPCC; cited for uncertainty terminology (Annex A)
CDP	CDP (formerly Carbon Disclosure Project)	Voluntary climate disclosure platform
CO₂e	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

Acronym	Spelled out	Notes
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	
ISO	International Organization for Standardization	Standards 14064-1 (GHG quantification) and 14289-1 (PDF/UA-1) referenced here
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
TPU	Tensor Processing Unit	Google's custom AI inference hardware
UI	User Interface	
UK	United Kingdom	
US	United States	