Towards Near-Real-Time Carbon Accounting for Internet Data Transmission: A Resolution-Based Attribution Framework

Kanhu Charan Pattnayak[1,2,*] and Akshae Golekar[1]

[1]*Carbon Tech International Limited, London, United Kingdom*
[2]*Cambridge Judge Business School, University of Cambridge, Cambridge, United Kingdom*

[*] **Corresponding Author:** kcpattnayak@gmail.com

Abstract

Digital infrastructure and internet-based services contribute an estimated 3-4% of global CO_2e emissions, a share comparable to aviation and projected to grow. Current carbon accounting practices for internet data transmission rely on annualised, batch-based estimation methods, which lack temporal granularity, omit critical components of the transmission circuit, and introduce structural bias into emissions reporting. This paper proposes and evaluates three complementary attribution methodologies designed to overcome these limitations: (1) an operational hybrid resolution method for compliance-grade reporting; (2) a hybrid deterministic-batch method for improved accuracy under real-world telemetry constraints; and (3) a fully deterministic event-level framework representing the theoretical upper bound of attribution fidelity. Using synthetic hourly datasets for 200 workloads of a representative mid-size website, built with the emission factors of a production carbon accounting platform, we show that the hybrid deterministic-batch method reduces hourly attribution error from 35% under conventional batch accounting to 5%, statistically indistinguishable from the fully deterministic benchmark, and still by more than half when only a quarter of its deterministic components are covered by telemetry. The operational hybrid method ensures complete, auditable coverage of emission sources and removes most of the annual bias of batch accounting. Our findings indicate that near-real-time carbon accounting for internet data transmission is technically feasible, significantly more accurate than current practices, and essential for supporting evolving regulatory and operational decision-making.

Keywords: Carbon Accounting; Internet Emissions; Data Transmission; Temporal Resolution; Attribution Accuracy; Hybrid Methodology; Regulatory Compliance

1. Introduction

The internet and associated digital infrastructure now account for an estimated 3-4% of global CO_2e emissions, a share comparable to that of the aviation sector (Belkhir and Elmeligi, 2018; Malmodin and Lundén, 2018; Freitag et al., 2021) and expected to rise with continued growth in cloud computing, artificial intelligence, and data-intensive digital services (Andrae and Edler, 2015; Obringer et al., 2021). Projections suggest that, in the absence of targeted methodological and operational interventions, emissions from digital infrastructure will continue to increase in both absolute and relative terms, amplifying the importance of robust carbon accounting for internet activity. Data centre electricity demand is projected to rise substantially with the expansion of AI workloads (IEA, 2025), although estimates of current and future energy use remain highly sensitive to methodological assumptions and data sources (Masanet et al., 2020; Mytton and Ashtine, 2022). Sector-specific assessments, for example of online advertising, likewise indicate a substantial but highly uncertain footprint (Pärssinen et al., 2018).

Despite this growing significance, emissions accounting for internet data transmission remains largely retrospective and coarse-grained, typically relying on annual or quarterly batch-based estimates derived from aggregated activity metrics and static emission factors. Published estimates of the electricity intensity of internet data transmission consequently span several orders of magnitude, largely because of differences in system boundaries, reference years and allocation methods (Coroama and Hilty, 2014; Aslan et al., 2018). Such methods fail to capture temporal variation in grid carbon intensity, dynamic changes in traffic volume and routing, and load-dependent compute efficiency, which leads to systematic misallocation and under- or over-estimation of emissions across time and services. In parallel, most existing frameworks focus on isolated subsystems such as cloud servers (e.g. Lannelongue et al., 2021) while excluding network routing, backbone transmission, CDNs, idle infrastructure capacity, and protocol or data-transfer overheads, thereby introducing structural bias into reported emissions. The effect of boundary choice is material in practice: in a production-scale measurement of AI serving, a narrow boundary covering only active accelerator energy underestimated per-request energy by a factor of 2.4 relative to a comprehensive boundary that also included host CPU and memory, idle provisioned capacity and data centre overhead (Elsworth et al., 2025).

These methodological limitations create a misalignment with regulatory developments such as the EU Corporate Sustainability Reporting Directive (CSRD), climate disclosure regimes such

as California's Climate Corporate Data Accountability Act (SB 253) and the ISSB standards (IFRS S2), and expanding carbon pricing and assurance frameworks, all of which increasingly demand emissions data that are accurate, auditable, and granular (European Union, 2022; IFRS Foundation, 2023; State of California, 2023). For internet-native organisations whose emissions are distributed across heterogeneous infrastructure layers, third-party service providers, and geographically diverse electricity grids, the combination of low temporal resolution and incomplete system boundaries constrains both regulatory compliance and operational decision-making.

This paper therefore addresses the fundamental methodological question: How can carbon emissions from internet data transmission be measured with sufficient accuracy and temporal resolution to support both regulatory compliance and operational decision-making? To answer this question, the study pursues two objectives: (1) to develop and formalise deployable and analytical attribution methodologies operating at multiple resolution levels for near-real-time carbon accounting across the full internet data transmission circuit; and (2) to empirically evaluate the accuracy improvements of these methodologies relative to conventional batch-based approaches. The core hypothesis is that attribution accuracy is fundamentally constrained by temporal resolution and system boundary completeness rather than by any single emission factor or modelling choice, and that improving these dimensions yields substantial gains even under realistic telemetry constraints.

To operationalise this hypothesis, the paper introduces a resolution ladder that distinguishes three levels of attribution resolution an operational hybrid method, a hybrid deterministic batch analytical framework, and a fully deterministic event-level model evaluated alongside conventional batch accounting. Together, these approaches clarify what is achievable today and what may become possible as infrastructure telemetry and transparency improve, providing both a conceptual map and a practical roadmap for the evolution of internet emissions accounting.

2. Data and Methodology

2.1 Data

The evaluation uses a synthetic but operationally plausible dataset representing a mid-size software-as-a-service (SaaS) and marketing website, a common type of platform for which

internet-related emissions are reported. Emission and energy factors are taken directly from the production implementation of the CarbonCut carbon accounting platform, so that the analysis reflects factors used in practice: network data transfer of 0.006 kWh/GB, origin and CDN processing of 0.004 Wh per visit, a default page weight of 2 MB per visit, CDN delivery of 0.30 g CO_2e/GB, end-user device energy of 0.0012 Wh (mobile) and 0.0028 Wh (desktop) per minute of viewing, a component power model for servers (5 W per vCPU, 0.375 W per GB of memory and 1.2 W per TB of SSD storage) and data centre PUE presets between 1.1 and 2.0 (ISO/IEC, 2016).

Two hundred workloads were generated, each simulated hourly for one year (8,760 hours). Workload parameters were sampled to span the range typical of this platform class: monthly visits log-uniformly between 20,000 and 120,000 (median 48,555), page weight of 1.1-3.3 MB per visit, a CDN cache-hit ratio of 75-95%, one to four cloud instances of 2-8 vCPU, 0.2-2 TB of triple-replicated storage, and a traffic mix across four markets (United Kingdom, United States, India and Germany) drawn from a Dirichlet distribution centred on 40/30/15/15%. Hourly traffic combines a working-day diurnal profile in each market's local time (peak-to-mean ratio 2.0), a weekend reduction, seasonal troughs in August and late December, six two-day marketing campaigns per year and Poisson noise. The origin infrastructure is located in a single London cloud region.

Hourly grid carbon intensity was synthesised for each market by combining the annual averages used by the platform (United Kingdom 162, United States 342, India 589 and Germany 298 gCO_2/kWh) with market-specific diurnal profiles (evening peaks and, where solar generation is significant, midday troughs) of ±25% amplitude, seasonal variation and autocorrelated weather-driven variability, consistent with the documented within-day and between-region variation in grid intensity (Hawkes, 2010; Siler-Evans et al., 2012; Miller et al., 2022). The resulting hourly intensities (5th-95th percentile) range from 102 to 221, 264 to 430, 455 to 738 and 190 to 401 gCO_2/kWh respectively. The simulated workloads have median annual operational emissions of 72 kg CO_2e (5th-95th percentile 27-200 kg), of which cloud compute accounts for 82%, storage 12%, network transmission 4%, end-user devices 1% and CDN delivery 0.5% (Figure 4).

2.2 Methods

This study develops a resolution-based attribution framework for quantifying carbon emissions from internet data transmission across the full transmission circuit, from end-user devices and access networks through backbone and content delivery infrastructure to cloud compute and storage. Emissions are estimated by converting component-level energy use into CO_2e using the carbon intensity of the electricity supplying each component, consistent with the GHG Protocol (WRI and WBCSD, 2004; WRI, 2015). Three attribution methods are formulated that differ in temporal resolution, system boundary coverage and the extent to which components are treated deterministically rather than through batch allocation, and these are compared against conventional annualised batch accounting. The methods are evaluated using synthetic datasets representative of a live digital platform, with performance assessed in terms of attribution error, component-level error behaviour and the ability to capture temporal variation in emissions.

2.2.1 Conceptualisation of the Internet Transmission Circuit

Internet data transmission is conceptualised as a multi-stage circuit that includes end-user devices, access networks and local ISPs, backbone and transit networks, CDN edge infrastructure where applicable, origin servers and cloud compute, storage and database services, and the return transmission path (Baliga et al., 2011; Schien et al., 2015). Each component contributes to energy consumption and associated emissions through compute, data transfer, cooling, and idle capacity overheads, but exposes distinct levels of observability and controllability to the reporting organisation. Recent first-party measurement frameworks illustrate this: Elsworth et al. (2025) define their boundary by operational control and therefore exclude external networking and end-user devices, which are central components of the transmission circuit considered here.

Annualised batch methods tend to smooth emissions over time and focus on a subset of this circuit, typically omitting structural elements such as backbone energy and protocol overhead, which hides both load-driven and grid-intensity-driven variability. Because the degree to which emissions can be attributed deterministically varies by component, the methodology adopts mixed-resolution attribution: high variance, well observed components are treated at event level where possible, while low variance or opaque components are incorporated using conservative, auditable batch proxies, consistent with the GHG Protocol principles of completeness, consistency and transparency (WRI and WBCSD, 2004; ISO, 2018). This

mixed-resolution treatment underpins all three proposed methods and ensures internal consistency across the resolution ladder.

2.2.2 Baseline: Conventional Batch Accounting

Conventional internet emissions accounting is expressed as a batch calculation in which annual or quarterly activity metrics (for example, total data transferred or compute hours) are multiplied by static, often region averaged emission factors (WRI, 2015). This formulation ignores time specific grid carbon intensity, which varies substantially within a day and between regions (Hawkes, 2010; Siler-Evans et al., 2012; Miller et al., 2022), routing variability, and component level energy attribution, and it typically restricts the system boundary to a subset of the full circuit, for example, cloud services alone. Even comprehensive first-party disclosures apply annual average, market-based emission factors to measured energy (Elsworth et al., 2025), so temporal variation in grid carbon intensity is not reflected in reported emissions. As a result, conventional batch accounting serves as a useful reference baseline but not as an adequate representation of temporally resolved, circuit wide emissions.

2.2.3 Method 1: Operational Hybrid Resolution

Method 1 is designed as a deployable, production-grade attribution resolution that supports regulatory compliance, organisational reporting, and operational visibility under real-world data constraints. It combines near-real-time event capture for high-frequency sources (such as website traffic) with monthly or otherwise periodic batch aggregation for structurally opaque or low-frequency sources (such as cloud provider reports, CDN bills, and workforce-related emissions), intentionally allowing heterogeneous temporal resolution by source.

User emissions are captured through a staged operational pipeline beginning with onboarding and configuration of emission sources, followed by continuous event ingestion where available and monthly ingestion of provider or organisational datasets, with emissions persisted at the source level alongside metadata on scope classification, data provenance, and attribution accuracy. Each emission record is normalised to source_type, scope, $kgCO_2$, period, accuracy, and sources are labelled according to assurance level, distinguishing high-assurance provider data, medium-assurance cost-based proxies, and conservative estimates, with scope classification following the GHG Protocol corporate and value chain standards (WRI and WBCSD, 2004, 2011).

Source-specific models include: high-assurance ingestion of cloud provider emissions where available, cost-to-carbon proxies when provider data are absent (using provider and region specific emission factors), event-level website emissions derived from observed data transfer and grid intensity, and CDN emissions calculated from monthly data volumes using provider-specific kWh/GB, PUE, and regional carbon intensity (ISO/IEC, 2016; Aslan et al., 2018). Method 1 prioritises completeness, reproducibility, and auditability over maximal attribution accuracy and explicitly avoids forcing event-level attribution on structurally opaque components where this would introduce false precision.

2.2.4 Method 2: Hybrid Deterministic–Batch Attribution

Method 2 is an analytical methodology that combines deterministic event-level attribution for high-variance, well-observed components with batch allocation for low-variance or poorly observed components, aiming to minimise attribution error under realistic telemetry constraints. High-confidence, real-time signals such as compute utilisation, data transfer volume, and geographic region are treated deterministically, using time-resolved grid carbon-intensity signals of the kind produced by consumption-based, flow-traced electricity accounting (Tranberg et al., 2019), while components such as backbone network energy and idle infrastructure are allocated using short-interval averages or conservative proxies.

At system level, Method 2 extends the effective boundary beyond that of legacy batch accounting and increases temporal resolution, as illustrated in Figure 8, which contrasts the two approaches in terms of coverage and allocation logic. Formally, emissions are decomposed into deterministically attributed components and batch-allocated residuals, with the allocation philosophy giving priority to deterministic treatment where variability and observability are greatest, thereby capturing most of the achievable accuracy gains without requiring fully pervasive telemetry.

2.2.5 Method 3: Full Event Level Deterministic Attribution

Method 3 defines a fully deterministic event-level attribution framework in which emissions are attributed to individual transmission events across the entire circuit using timestamped, component-level energy estimates. For a given event e, emissions are calculated as

$$CO_2e_e = \sum_{k \in circuit} \left(E_{k,e} \times CI_{k,t} \right),$$

where $E_{k,e}$ denotes the energy attributable to component k for event e, and $CI_{k,t}$ is the carbon intensity of the electricity supplying component k at time t.

This method assumes access to detailed telemetry including routing paths, per-component energy use, and allocation of idle capacity, which is not currently available across most infrastructure, so it is presented as a theoretical upper bound on achievable attribution fidelity rather than a universally deployable solution. The method nonetheless serves as an important benchmark: analysis of its behaviour indicates that the largest incremental accuracy gains arise from deterministic treatment of compute and CDN components and from fine-grained temporal alignment with grid carbon intensity. This is consistent with evidence that hourly rather than annual accounting materially changes the emissions attributed to electricity consumption (Miller et al., 2022).

2.3 Experimental Setup

To evaluate the methodologies, the dataset described in Section 2.1 is used to construct, for every workload and hour, a reference emissions series and the estimates produced by each approach. All results are based on synthetic but operationally plausible data, selected to preserve the essential statistical and structural characteristics of production systems while avoiding disclosure of proprietary information.

Experimental evaluation focuses on comparing Method 2 against Method 3, with conventional batch accounting serving as the reference baseline; Method 1 is treated as an operational standard and is reported for completeness but not optimised for error. The experimental architecture (Figure 11) includes user devices in four markets, access and backbone network segments, CDN edges in each market and a single-region cloud origin, providing a realistic environment in which to explore attribution accuracy and temporal responsiveness.

The reference ("true") emissions are generated by a component-level hourly model that is richer than any of the estimators. Compute energy follows a linear server power model with an idle floor of 35% of peak power, driven by origin request load after CDN caching and multiplied by PUE (1.3 at baseline), plus a constant idle provisioned-capacity term sized to 10% of total energy. Storage energy combines replicated SSD power with a load-dependent input/output term. Network energy applies the network intensity to transferred volume plus a load-dependent protocol and retransmission overhead of 5-11%, and is attributed at the carbon

intensity of the routing path (70% user market, 30% origin region). CDN energy is attributed at the intensity of the edge market and device energy at the intensity of the user market.

The four approaches observe this system through different data. Conventional batch accounting multiplies annual instance-hours at an assumed 50% utilisation, annual stored volume and annual transferred data by static factors (the annual average of the cloud region, the platform storage factor and the platform default of 387 gCO_2/kWh), excludes CDN, device, protocol-overhead and idle-capacity energy, and spreads the annual total evenly over the year. Method 1 uses event-level transfer volumes from the page-view tracker with the hourly grid signal, monthly provider reports for compute and storage, and monthly CDN volumes multiplied by the platform CDN factor. Method 2 attributes compute (from hourly utilisation) and CDN (from hourly logs) deterministically with the hourly grid signal, uses hourly transfer volumes with a mean overhead and daily average intensity for the network, and daily proxies for storage and devices. Method 3 attributes every component hourly from component telemetry and path-level intensity. All observed data carry realistic measurement error: the hourly grid-intensity signal ±5%, utilisation-based power and CDN logs ±5%, path-level network energy ±8%, device-class power ±15% and monthly provider reports ±3% (Figure 14).

Attribution error is measured for each workload as the weighted absolute percentage error of the hourly emissions series, that is, the sum of absolute hourly deviations divided by total reference emissions, which avoids instability in low-traffic hours. MAE denotes the mean of this quantity across the 200 workloads and the error range its 5th-95th percentile. Bias is the signed error in annual totals. Uncertainty in the MAE is quantified with 10,000 bootstrap resamples of workloads, and paired differences between methods are tested with the Wilcoxon signed-rank test. Sensitivity is assessed one parameter at a time on 100 workloads (Section 3.9). The simulation code is written in Python (NumPy, SciPy) with a fixed random seed.

3. Results

3.1 Overall Attribution Accuracy

Attribution accuracy is assessed using mean absolute error (MAE) and error ranges for four approaches: conventional batch accounting, Method 1, Method 2 and Method 3, each operating at a different point on the resolution ladder (Figure 6). Conventional batch accounting, with annual temporal resolution and partial system boundary coverage, exhibits a MAE of 35.0% with a range of 28.9-47.3%, indicating substantial deviation from the hourly reference. Most of this error is structural rather than temporal: batch accounting understates annual emissions

by 32.6% on average, because it omits CDN, device, protocol-overhead and idle-capacity energy and relies on a storage factor and a utilisation assumption that understate actual energy use (Figure 3).

Method 1, which combines event-level and monthly signals within a conservative, compliance-oriented framework, reduces MAE to 20.0% (18.3-21.5%). Its annual totals are close to the reference (mean bias -0.6%), so its remaining hourly error reflects the uniform distribution of monthly compute, storage and CDN totals across each month rather than incomplete coverage. Method 2 reduces MAE to 5.0% (4.8-5.6%), a 7.0-fold reduction relative to conventional batch accounting, by attributing compute and CDN emissions hourly using utilisation data, CDN logs and the hourly grid signal. Method 3, the full event-level deterministic model, achieves the lowest error, with a MAE of 4.9% (4.3-5.2%), which reflects only the measurement error assumed for component telemetry and grid-intensity signals.

The pattern across the four approaches confirms that attribution performance improves monotonically with increased temporal resolution and system boundary completeness, with sharply diminishing returns once compute and CDN components are treated at event level: for the platform profile studied here, the step from Method 2 to Method 3 reduces MAE by only 0.2 percentage points.

3.2 Component-Level Error Behaviour

Beyond aggregate MAE, the analysis examines how attribution error varies across components of the transmission circuit (Figure 9). Under conventional batch accounting, CDN and device emissions are excluded entirely (100% component error), storage is understated by the platform's per-gigabyte storage factor (98% error), and network emissions carry a 53% error from the use of a static default intensity and the omission of protocol overhead. Because compute represents 82% of emissions, however, the 30% error in compute dominates the aggregate. Under Method 1, residual error is concentrated in compute (21%) and CDN (46%), whose monthly totals are correct but temporally flattened, while event-level treatment of data transfer keeps network error at 8%.

Method 2's hybrid design targets the components that drive this residual: deterministic attribution of compute and CDN emissions reduces their errors to 5% and 4% respectively. Its largest remaining component errors are for devices (27%), network (16%) and storage (12%), which are allocated through daily proxies, but these components together account for only 17% of emissions. Under Method 3, which assumes ideal telemetry, component errors fall to 4-8% and are determined by assumed measurement error rather than by structural limitations in system boundary or resolution.

3.3 Temporal Responsiveness and Peak Capture

Temporal responsiveness is evaluated by comparing how each method represents emissions during periods of high load and periods of elevated grid carbon intensity. Batch methods, which aggregate activity and apply static emission factors, smooth emissions over time and fail to reproduce peaks that occur when high traffic coincides with carbon-intensive grid conditions (Figure 2), thereby under-representing the operational importance of shifting or reducing demand during these windows. In the simulated workloads, the 10% of hours with the highest effective grid intensity account for 14.4% of annual emissions (13.9-14.8%), whereas a time-invariant batch estimate assigns them exactly 10%.

All three proposed methods capture temporal variation to some extent, with the event-level approach providing the highest sensitivity to short-term fluctuations (Figure 10). Method 1 preserves high-frequency signals for sources where event data are available, enabling near-real-time representation of emissions associated with website traffic, while Method 2 goes further by aligning deterministic attribution of compute and CDN with time-resolved grid intensity, making peak emissions visible and actionable. Emissions time series show a tight coupling between grid carbon intensity and attributed emissions under Methods 2 and 3 (Figure 12); at daily resolution their errors fall to 1.2% and 1.0%, compared with 13.7% for Method 1 and 33.3% for conventional batch accounting.

3.4 Resolution Ladder and Marginal Gains

The results support the conceptual resolution ladder by showing that improvements in temporal resolution and system boundary completeness drive most of the observed accuracy gains, particularly when moving from conventional batch accounting to the hybrid deterministic–batch framework. Method 1, while not optimised for fidelity, illustrates that simply expanding the system boundary and introducing mixed temporal resolution can materially reduce structural bias even when parts of the circuit remain covered by batch proxies.

The comparison between Method 2 and Method 3 shows that marginal gains from full determinism are concentrated in components where routing paths, path-level grid intensity and device behaviour influence emissions, namely network transmission and end-user devices. For a website whose emissions are dominated by compute, these components are a small share of the total, so Method 2 captures 99% of the achievable reduction in error without the complexity and data requirements of Method 3. Organisations can therefore realise most of the improvement in attribution accuracy without waiting for complete infrastructure transparency; for network-heavy services such as video streaming or large file transfer, the gap between the two methods would be expected to be larger.

3.5 Marginal Gains along the Resolution Ladder

To move beyond aggregate error values, Table 1 expresses the performance of each method relative to conventional batch accounting and to the full event-level benchmark. Method 2 reduces MAE by 30.0 percentage points (85.7%) relative to batch accounting, an error ratio of 6.98, and closes 99.5% of the gap between batch accounting and Method 3. Method 1 closes 50% of this gap despite not being optimised for fidelity, indicating that boundary completeness and the use of provider data alone account for about half of the achievable improvement. The final step from Method 2 to Method 3 yields a further 0.2 percentage points while requiring full telemetry for every component of the circuit. Figure 13 shows how this outcome depends on telemetry availability: the error of Method 2 falls steadily as coverage of its deterministic components increases, from 18.6% with no coverage to 5.0% with full coverage.

Table 1. Attribution performance of each method relative to conventional batch accounting and to the full event-level benchmark, derived from the MAE values reported in Section 3.1. The share of achievable gain is the reduction in MAE relative to batch accounting divided by the reduction achieved by Method 3.

Approach	MAE (%)	Error range (%)	Reduction vs batch (pp)	Relative reduction vs batch (%)	Error ratio (batch / method)	Share of achievable gain (%)
Conventional batch	35.0	28.9-47.3	-	-	1.00	0
Method 1: Operational hybrid	20.0	18.3-21.5	15.1	43.0	1.76	50
Method 2: Hybrid deterministic-batch	5.0	4.8-5.6	30.0	85.7	6.98	99
Method 3: Full event-level	4.9	4.3-5.2	30.2	86.1	7.21	100

Expressed as proportional reductions between successive rungs, the gains are uneven: the step from batch accounting to Method 1 reduces the remaining error by 43.0%, Method 1 to Method 2 by 74.9%, and Method 2 to Method 3 by only 3.1%. The diminishing returns at the top of the ladder reflect the fact that, once compute and CDN are treated deterministically, the remaining error is dominated by measurement noise in telemetry and grid signals, which full determinism does not remove. This distinction has practical consequences. Where the absolute size of the residual error is the decision-relevant quantity, as in compliance reporting against materiality thresholds, Method 2 is sufficient for this platform class; further investment in telemetry is justified mainly for services in which network and device emissions form a larger share of the total, or where better grid-intensity signals reduce the measurement error shared by all high-resolution methods.

3.6 Separability and Uncertainty of Method Performance

The bootstrap 95% confidence intervals for the MAE are 34.2-35.9% for conventional batch accounting, 19.8-20.1% for Method 1, 5.0-5.1% for Method 2 and 4.8-4.9% for Method 3. The intervals for the first three rungs are widely separated, and Wilcoxon signed-rank tests confirm that Method 2 outperforms both batch accounting and Method 1 in every workload ($p < 0.001$). The difference between Method 2 and Method 3, by contrast, is not statistically significant (Wilcoxon $p = 0.30$; median paired difference -0.04 percentage points; Method 2 has the higher error in 44% of workloads), and their error ranges (4.8-5.6% and 4.3-5.2%) overlap. Within the platform class studied, the advantage of full determinism over the hybrid approach is therefore not distinguishable from measurement noise. The error distribution is also much narrower for the higher-resolution methods (Figure 6), because their remaining error is driven by telemetry noise that is similar across workloads, whereas batch error depends on workload-specific boundary omissions such as the shares of storage and CDN.

3.7 Component Contributions to Accuracy Gains

Figure 9 decomposes the error reduction achieved by the hybrid approach by component of the transmission circuit. In relative terms, the largest reductions occur for CDN (96%), storage (88%) and compute (82%), followed by end-user devices (73%) and network transmission (69%). Because compute accounts for 82% of emissions, its reduction from 30% to 5% contributes most of the aggregate improvement. This ranking is consistent with the design logic of Method 2: compute and CDN energy use is both variable, responding to load, caching behaviour and time of day, and well observed through utilisation metrics and CDN logs, so deterministic treatment yields the largest benefit (Figure 7). Network transmission shows the smallest relative improvement because backbone and access-network energy remain allocated through proxies and daily average intensity in the absence of routing-level telemetry. Together with end-user devices, it therefore represents the main source of residual error under current data availability, and the component where improved provider transparency would most increase accuracy. The importance of components that are often omitted is supported by production measurements of AI serving, in which idle provisioned capacity and facility overhead together accounted for about 18% of per-request energy and were excluded by narrower measurement boundaries (Elsworth et al., 2025).

3.8 Error Direction and Implications for Reporting

The practical consequence of attribution error for reporting depends on whether it affects annual totals or only their distribution in time. For an organisation reporting annual internet-related emissions of 1,000 tCO_2e, the mean absolute error in annual totals implies an expected deviation of approximately ±326 tCO_2e under conventional batch accounting, compared with

±9, ±6 and less than ±1 tCO_2e under Methods 1, 2 and 3 respectively. Against the materiality thresholds of around 5% commonly applied in greenhouse gas verification, conventional batch accounting fails by a wide margin, whereas all three proposed methods meet the threshold for annual totals. For daily or hourly reporting, which carbon-aware operation requires, only Methods 2 and 3 approach this level of precision (daily errors of 1.2% and 1.0%, against 13.7% for Method 1). This reinforces the role of the provenance and assurance tagging built into Method 1, which makes the remaining temporal uncertainty explicit to auditors rather than concealing it within a single annual figure.

MAE does not indicate whether an approach systematically over- or under-estimates emissions. In the simulated workloads, the mean signed error in annual totals is -32.6% for conventional batch accounting (5th-95th percentile -47.0% to -24.2%), -0.6% for Method 1, +0.6% for Method 2 and +0.01% for Method 3. The batch bias is driven mainly by boundary omissions and by factors that understate storage and idle energy, rather than by temporal averaging. Temporal averaging introduces a further, profile-dependent bias: because batch methods apply annual average carbon intensity, they understate emissions when demand coincides with carbon-intensive grid periods and overstate them when demand is concentrated in low-carbon periods (Miller et al., 2022). The 10% of hours with the highest grid carbon intensity hold 14.4% of emissions, so a time-invariant factor understates the emissions of these hours by about 31%. The direction of bias therefore depends on how an organisation's traffic profile aligns with grid conditions (Figure 12), which has direct implications for the credibility of reported reductions and for the design of demand-shifting strategies (Wiesner et al., 2021; Radovanović et al., 2023).

3.9 Sensitivity to Modelling Assumptions

Because the results are derived from synthetic datasets, their robustness to the assumptions used to generate those datasets must be established. The key test is whether the advantage of Method 2 over conventional batch accounting is preserved across plausible ranges of the parameters that most influence attributed emissions (Table 2). Published estimates of network electricity intensity, for example, vary by orders of magnitude depending on boundary and reference year (Coroama and Hilty, 2014; Aslan et al., 2018), and grid carbon intensity varies substantially by hour and region (Hawkes, 2010; Siler-Evans et al., 2012). Each parameter was set to the lower and upper ends of its range with all others held at baseline, on 100 workloads. Across the five physical parameters, the ratio of batch to Method 2 MAE remains between 6.1 and 8.4, so the advantage of the hybrid approach is not an artefact of any single physical assumption. Batch error is most sensitive to the idle capacity share (31.4-39.8%) and network

intensity (32.8-45.3%), because both increase the share of energy that lies outside the batch boundary or is estimated with static factors. The assumption that materially changes the result is telemetry coverage: when only 25% of the deterministic components are covered, Method 2 MAE rises to 14.8% and the error ratio falls to 2.3, close to the approximately two-fold advantage expected under sparse telemetry.

Table 2. Sensitivity of attribution error to key modelling assumptions (one-at-a-time, 100 workloads). Each cell gives the value at the lower / upper end of the range tested; error ratio = MAE batch / MAE Method 2.

Parameter	Baseline	Range tested	MAE batch (%)	MAE Method 2 (%)	Error ratio
Network electricity intensity (kWh/GB)	0.006	0.002-0.06	32.8 / 45.3	5.0 / 5.6	6.5 / 8.1
Diurnal amplitude of grid carbon intensity	±25%	±10% to ±50%	33.0 / 36.3	4.8 / 5.5	6.9 / 6.6
Traffic peak-to-mean ratio	2.0	1.5-3.0	34.8 / 33.7	4.9 / 5.1	7.0 / 6.6
Idle capacity share of total energy	10%	5-20%	31.4 / 39.8	5.1 / 4.7	6.1 / 8.4
Power usage effectiveness (PUE)	1.3	1.1-1.6	33.7 / 34.1	5.0 / 5.0	6.8 / 6.8
Telemetry coverage of deterministic components	100%	25-100%	33.9 / 33.9	14.8 / 5.0	2.3 / 6.8

3.10 Critical Appraisal of the Results

Several features of the experimental design qualify the interpretation of these results. First, all errors are measured against a reference derived from the same synthetic data-generating process that underpins the deterministic methods; methods that share more structure with that process will tend to perform better, so the absolute error values should be read as indicators of relative performance rather than as estimates of real-world accuracy. Second, Method 3 itself shows a non-zero MAE, which implies that the reference emissions are not identical to the Method 3 formulation. The reference emissions are generated by a richer component-level model (Section 2.3) that includes an idle power floor, load-dependent protocol overhead and path-level grid intensity, and Method 3 departs from it only through the measurement error assumed for telemetry and grid-intensity signals. The absolute errors of Methods 2 and 3 are therefore set largely by these assumed error levels (Figure 14) and should not be read as achievable real-world accuracy. In addition, part of the batch error arises from platform default factors, notably a per-gigabyte storage factor that understates storage energy relative to a power-based estimate; batch accounting with better calibrated static factors would show a smaller bias, although its temporal error would remain. Third, the analysis covers a single platform architecture and traffic profile; services dominated by video streaming, large file transfer or AI inference differ in the relative weight of network and compute components, and the gains from each rung of the ladder may differ accordingly (Aslan et al., 2018; Elsworth et

al., 2025). Finally, the evaluation is limited to operational electricity-related emissions and does not consider the embodied emissions of network and end-user equipment, which can be material over the life cycle of digital services (Gupta et al., 2021). Validation against metered data from a production environment, as has recently been demonstrated for AI serving (Elsworth et al., 2025), is therefore the most important next step.

4. Discussion

The expanded results confirm that attribution accuracy for internet emissions is primarily a function of temporal resolution and system boundary completeness rather than any single emission factor choice, and that hybrid approaches can capture most of the attainable gains. By combining deterministic event-level attribution for high-variance components with conservative batch treatment of structurally opaque elements, the hybrid deterministic batch method substantially outperforms conventional batch accounting while remaining deployable at scale under current data availability. This mirrors evidence from production measurement of AI serving, where differences in measurement boundary alone account for order-of-magnitude variation in published per-request energy estimates, and a narrow boundary understated energy by a factor of 2.4 relative to a comprehensive boundary for the same workload (Elsworth et al., 2025).

From a policy and operational perspective, these findings support a staged transition: organisations can first adopt an operational hybrid baseline to ensure comprehensive, auditable coverage, then progressively introduce hybrid deterministic–batch methods as telemetry and data integration improve, while viewing full event-level determinism as a longer-term target. Time-resolved attribution also enables operational levers such as carbon-aware scheduling and temporal workload shifting, which have been shown to reduce emissions in cloud and data centre settings (Wiesner et al., 2021; Radovanović et al., 2023), and aligns with emerging software carbon-intensity standards (ISO/IEC, 2024). The resolution ladder provides a structured framework for this evolution, aligning methodological ambition with practical constraints and regulatory needs.

The study is limited by incomplete access to ISP-level and backbone energy data, partial reliance on inferred routing paths, and the use of synthetic pilot datasets rather than cross-provider empirical datasets. The wider literature on computing emissions shows that reported footprints are highly sensitive to measurement boundary, hardware utilisation and the treatment of embodied emissions (Strubell et al., 2019; Gupta et al., 2021; Patterson et al., 2021; Luccioni et al., 2023), underlining the need for empirical validation of the present framework. These limitations are structural rather than methodological and primarily affect the

feasibility and calibration of Methods 2 and 3, while leaving the compliance-oriented design of Method 1 largely unaffected. In addition, the framework considers operational electricity-related emissions only; embodied hardware emissions, which can be included through amortised Scope 1 and Scope 3 contributions (Elsworth et al., 2025), and the choice between location-based and market-based emission factors should be addressed in future work.

5. Conclusion

Near-real-time carbon accounting for internet data transmission can provide a more accurate and operationally useful view of digital emissions, but its practical value depends on how temporal resolution, system boundary completeness, and data availability are balanced in concrete methodologies. The resolution ladder developed in this work shows that organisations do not need to jump directly from annual batch estimates to fully deterministic event-level models; instead, they can progress through staged approaches that yield significant accuracy gains while remaining compatible with current telemetry and reporting infrastructures.

At the first stage, an operational hybrid resolution offers a conservative, compliance-grade baseline that consolidates all material emission sources across the internet transmission circuit using a mixture of near-real-time event signals and audited batch inputs. This design corrects major omissions of traditional methods by explicitly incorporating network and CDN components and by preserving temporal variability wherever event-level data are available, while clearly tagging each emission record with scope, provenance, and assurance level. For many organisations facing regulatory requirements under frameworks such as CSRD, California SB 253 and IFRS S2, this level represents an actionable near-term target that strengthens auditability without requiring speculative data or intrusive telemetry.

The second stage, the hybrid deterministic-batch approach, brings reported emissions even closer to operational reality by treating high-variance, well-observed components deterministically and reserving batch allocation for low-variance or structurally opaque elements. Experimental results for a representative website indicate that this design reduces hourly attribution error by about 86% relative to conventional batch accounting, and still by more than half when only a quarter of its deterministic components are covered by telemetry, capturing most of the achievable accuracy gains attributable to finer temporal resolution and expanded system boundaries. This enables higher-fidelity applications such as time-of-use emissions optimisation, carbon-aware workload shifting, and more robust product-level footprints for digital services, all without demanding universal, high-frequency telemetry across the entire circuit.

The final stage full event-level deterministic attribution serves primarily as a directional benchmark, clarifying the theoretical upper bound on attribution fidelity and indicating where expanded transparency (for example in compute, CDN, and routing data) would most improve accuracy. As infrastructure providers gradually expose richer energy and emissions signals, elements of this upper-bound framework can be integrated into operational systems, narrowing the gap between theoretical and practical performance. Overall, the findings suggest that the key challenge is methodological design rather than perfect data: by structuring accounting systems around a resolution ladder, policymakers and practitioners can drive continuous improvement in internet emissions reporting while ensuring that methods remain realistic, auditable, and aligned with decarbonisation goals.

Acknowledgement: The authors acknowledge their respective Heads of the institutions for providing support and encouraging collaborative research. Their guidance and facilitation significantly contributed to the successful completion of this study.

CRediT Taxonomy/Authors' contribution: KCP: Conceptualization; Formal analysis; Visualization; Software, Validation, Writing - original draft; and Writing - review &editing, VS: Visualization; Validation, **AG**: Conceptualization; Formal analysis; Visualization; Software, Validation, Writing - original draft; Writing - review & editing, Supervision.

Funding: No specific funding is received for this study.

Data Availability: The synthetic datasets analysed in this study, and the code used to generate them and compute the results, are available at [repository link and DOI to be added].

Conflict of interest: All authors are affiliated with Carbon Tech International Limited, [which develops commercial carbon accounting products related to the methods described in this paper]. The authors declare no other financial or personal relationships that could have appeared to influence the work reported in this manuscript.

References:

Andrae, A.S.G., Edler, T., 2015. On global electricity usage of communication technology: trends to 2030. Challenges 6, 117–157.

Aslan, J., Mayers, K., Koomey, J.G., France, C., 2018. Electricity intensity of Internet data transmission: untangling the estimates. J. Ind. Ecol. 22, 785–798.

Baliga, J., Ayre, R.W.A., Hinton, K., Tucker, R.S., 2011. Green cloud computing: balancing energy in processing, storage, and transport. Proc. IEEE 99, 149–167.

Belkhir, L., Elmeligi, A., 2018. Assessing ICT global emissions footprint: trends to 2040 & recommendations. J. Clean. Prod. 177, 448–463.

Coroama, V.C., Hilty, L.M., 2014. Assessing Internet energy intensity: a review of methods and results. Environ. Impact Assess. Rev. 45, 63–68.

Elsworth, C., Huang, K., Patterson, D., Schneider, I., Sedivy, R., Goodman, S., Townsend, B., Ranganathan, P., Dean, J., Vahdat, A., Gomes, B., Manyika, J., 2025. Measuring the environmental impact of delivering AI at Google scale. arXiv preprint arXiv:2508.15734.

European Union, 2022. Directive (EU) 2022/2464 of the European Parliament and of the Council of 14 December 2022 as regards corporate sustainability reporting (Corporate Sustainability Reporting Directive). Off. J. Eur. Union L 322, 15–80.

Freitag, C., Berners-Lee, M., Widdicks, K., Knowles, B., Blair, G.S., Friday, A., 2021. The real climate and transformative impact of ICT: a critique of estimates, trends, and regulations. Patterns 2, 100340.

Gupta, U., Kim, Y.G., Lee, S., Tse, J., Lee, H.-H.S., Wei, G.-Y., Brooks, D., Wu, C.-J., 2021. Chasing carbon: the elusive environmental footprint of computing. In: 2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA), pp. 854 –867.

Hawkes, A.D., 2010. Estimating marginal CO2 emissions rates for national electricity systems. Energy Policy 38, 5977–5987.

IEA, 2025. Energy and AI. International Energy Agency, Paris.

IFRS Foundation, 2023. IFRS S2 Climate-related Disclosures. International Sustainability Standards Board, London.

ISO, 2018. ISO 14064-1:2018 Greenhouse gases — Part 1: Specification with guidance at the organization level for quantification and reporting of greenhouse gas emissions and removals. International Organization for Standardization, Geneva.

ISO/IEC, 2016. ISO/IEC 30134-2:2016 Information technology — Data centres — Key performance indicators — Part 2: Power usage effectiveness (PUE). International Organization for Standardization, Geneva.

ISO/IEC, 2024. ISO/IEC 21031:2024 Information technology — Software Carbon Intensity (SCI) specification. International Organization for Standardization, Geneva.

Lannelongue, L., Grealey, J., Inouye, M., 2021. Green Algorithms: quantifying the carbon footprint of computation. Adv. Sci. 8, 2100707.

Luccioni, A.S., Viguier, S., Ligozat, A.-L., 2023. Estimating the carbon footprint of BLOOM, a 176B parameter language model. J. Mach. Learn. Res. 24 (253), 1–15.

Malmodin, J., Lundén, D., 2018. The energy and carbon footprint of the global ICT and E&M sectors 2010–2015. Sustainability 10, 3027.

Masanet, E., Shehabi, A., Lei, N., Smith, S., Koomey, J., 2020. Recalibrating global data center energy-use estimates. Science 367, 984–986.

Miller, G.J., Novan, K., Jenn, A., 2022. Hourly accounting of carbon emissions from electricity consumption. Environ. Res. Lett. 17, 044073.

Mytton, D., Ashtine, M., 2022. Sources of data center energy estimates: a comprehensive review. Joule 6, 2032–2056.

Obringer, R., Rachunok, B., Maia-Silva, D., Arbabzadeh, M., Nateghi, R., Madani, K., 2021. The overlooked environmental footprint of increasing Internet use. Resour. Conserv. Recycl. 167, 105389.

Pärssinen, M., Kotila, M., Cuevas, R., Phansalkar, A., Manner, J., 2018. Environmental impact assessment of online advertising. Environ. Impact Assess. Rev. 73, 177–200.

Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., Dean, J., 2021. Carbon emissions and large neural network training. arXiv preprint arXiv:2104.10350.

Radovanović, A., Koningstein, R., Schneider, I., et al., 2023. Carbon-aware computing for datacenters. IEEE Trans. Power Syst. 38, 1270–1280.

Schien, D., Coroama, V.C., Hilty, L.M., Preist, C., 2015. The energy intensity of the Internet: edge and core networks. In: Hilty, L.M., Aebischer, B. (Eds.), ICT Innovations for Sustainability. Advances in Intelligent Systems and Computing 310. Springer, Cham, pp. 157–170.

Siler-Evans, K., Azevedo, I.L., Morgan, M.G., 2012. Marginal emissions factors for the U.S. electricity system. Environ. Sci. Technol. 46, 4742–4748.

State of California, 2023. Senate Bill No. 253: Climate Corporate Data Accountability Act. Approved by the Governor 7 October 2023.

Strubell, E., Ganesh, A., McCallum, A., 2019. Energy and policy considerations for deep learning in NLP. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 3645–3650.

Tranberg, B., Corradi, O., Lajoie, B., Gibon, T., Staffell, I., Andresen, G.B., 2019. Real-time carbon accounting method for the European electricity markets. Energy Strateg. Rev. 26, 100367.

Wiesner, P., Behnke, I., Scheinert, D., Gontarska, K., Thamsen, L., 2021. Let's wait awhile: how temporal workload shifting can reduce carbon emissions in the cloud. In: Proceedings of the 22nd International Middleware Conference, pp. 260–272.

WRI, 2015. GHG Protocol Scope 2 Guidance. World Resources Institute, Washington, DC.

WRI and WBCSD, 2004. The Greenhouse Gas Protocol: A Corporate Accounting and Reporting Standard, revised edition. World Resources Institute and World Business Council for Sustainable Development, Washington, DC and Geneva.

WRI and WBCSD, 2011. Corporate Value Chain (Scope 3) Accounting and Reporting Standard. World Resources Institute and World Business Council for Sustainable Development, Washington, DC and Geneva.

Figures

Figure 1 Estimated contribution of digital infrastructure and internet-based services to global CO_2e emissions.

Figure 2 annualised batch methods smooth emissions over time, masking both load driven and grid-intensity-driven variability.

Figure 3 illustrates the structural omissions introduced by conventional subsystem-focused accounting.

Figure 4 Relative emissions contribution across the full transmission circuit.

Figure 5 Compares methodological capabilities across existing and proposed approaches.

Figure 6 The resulting attribution error variability across workloads.

Figure 7 illustrates the allocation philosophy underpinning the hybrid approach.

Figure 8A and 8B. Comparison of legacy batch-based internet carbon accounting and the proposed hybrid deterministic–batch framework, highlighting differences in system boundary coverage, temporal resolution, and attribution logic.

Figure 9 the largest accuracy gains arise from deterministic treatment of compute and CDN components.

Figure 10 illustrates the temporal granularity achievable under full event-level attribution.

Figure 11 The experimental architecture used to generate representative pilot data.

Figure 12 shows the coupling between grid carbon intensity and attributed emissions under near-real-time methods.

Figure 13 highlights why the hybrid approach captures most achievable accuracy gains under current telemetry constraints.

Figure 14 Heatmap of data confidence across components and signal types.

Figure 15 A practical roadmap from batch estimation to full deterministic accounting.

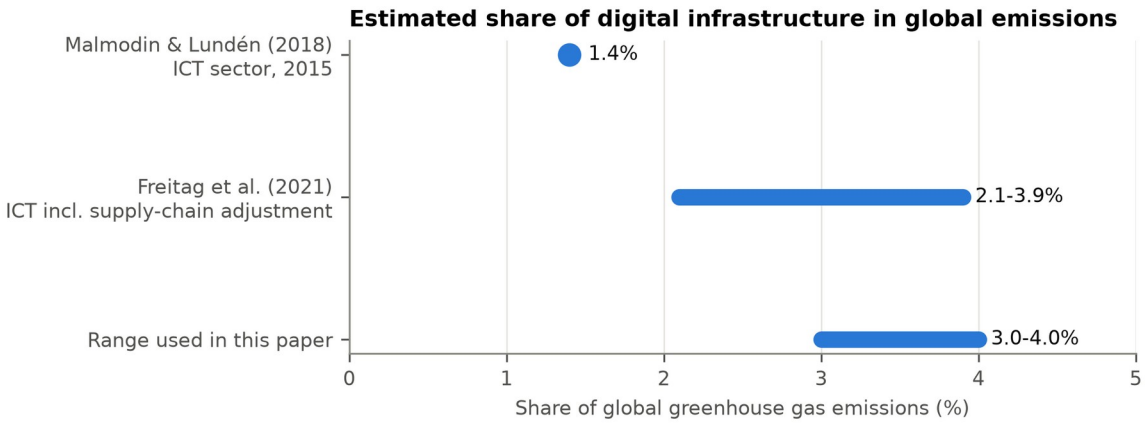


Figure 1 Estimated contribution of digital infrastructure and internet-based services to global CO_2e emissions.

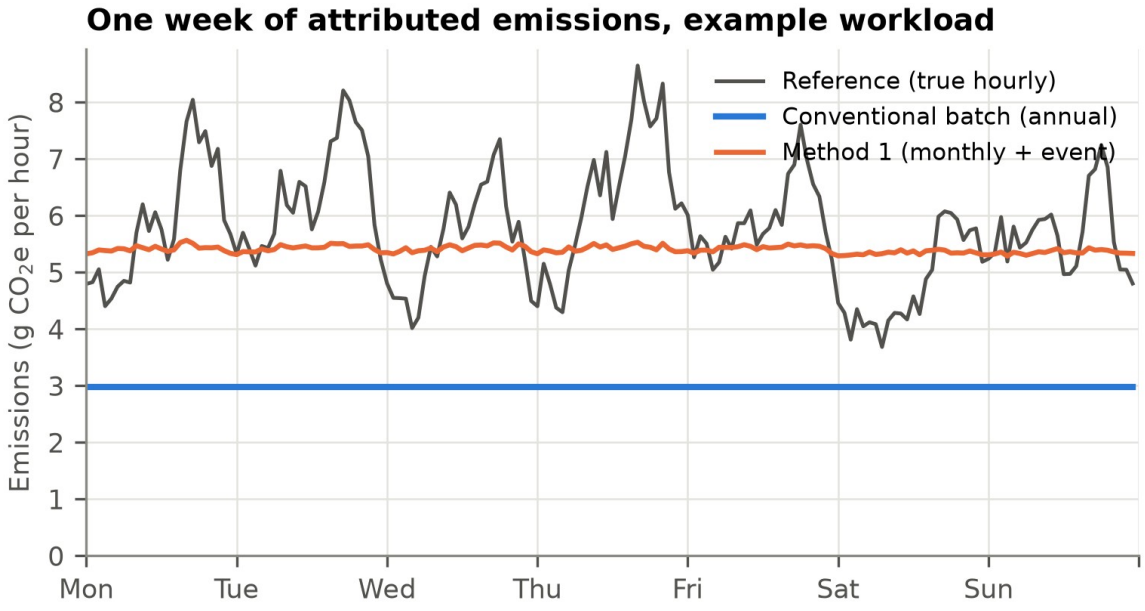


Figure 2 annualised batch methods smooth emissions over time, masking both load driven and grid-intensity-driven variability.

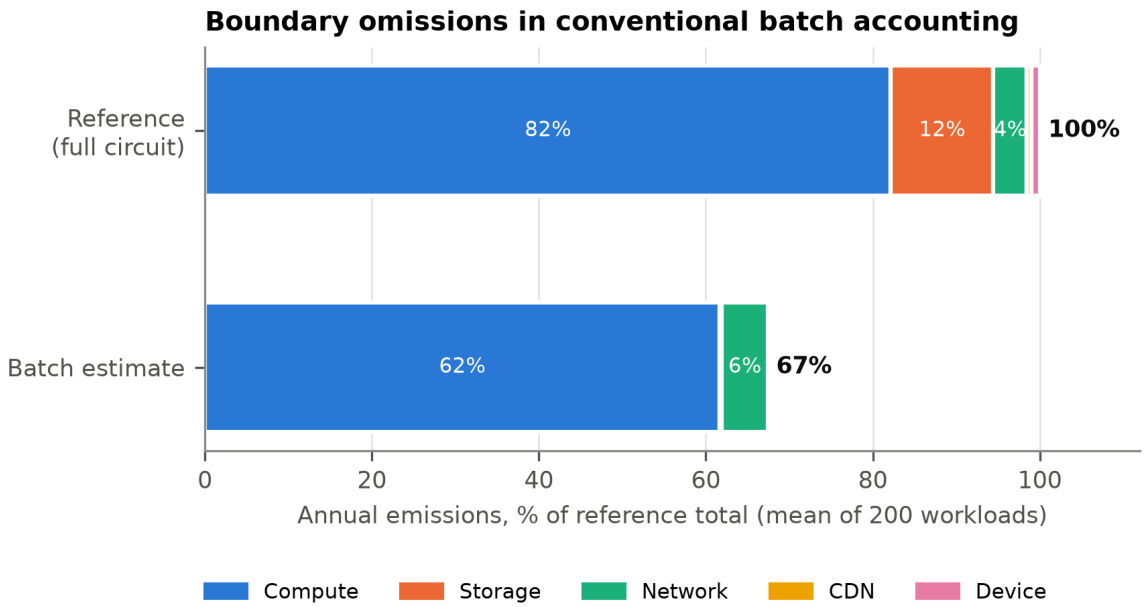


Figure 3 illustrates the structural omissions introduced by conventional subsystem-focused accounting.

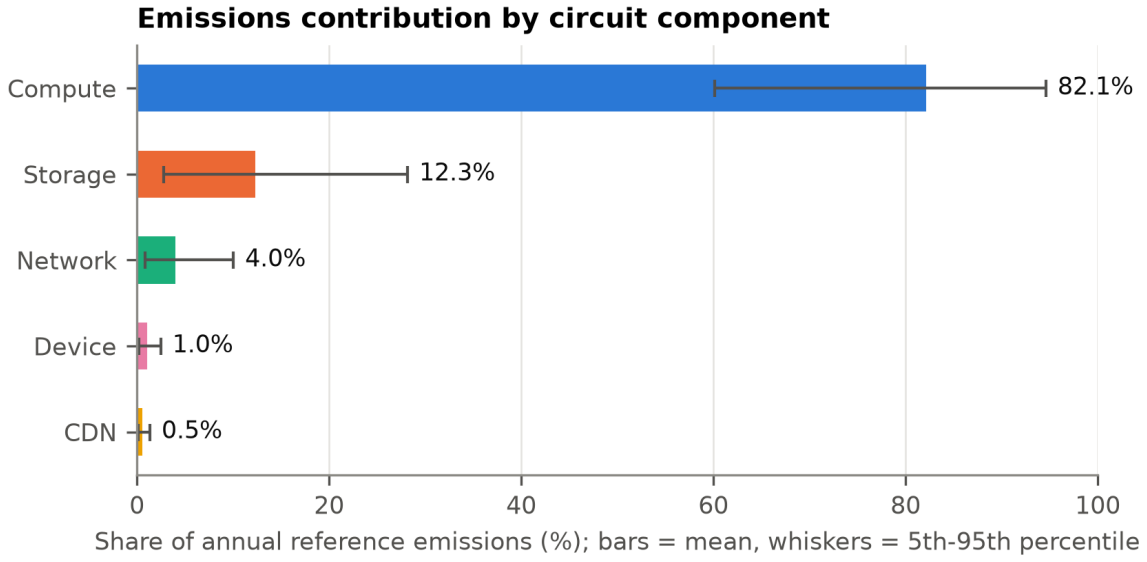


Figure 4 Relative emissions contribution across the full transmission circuit.

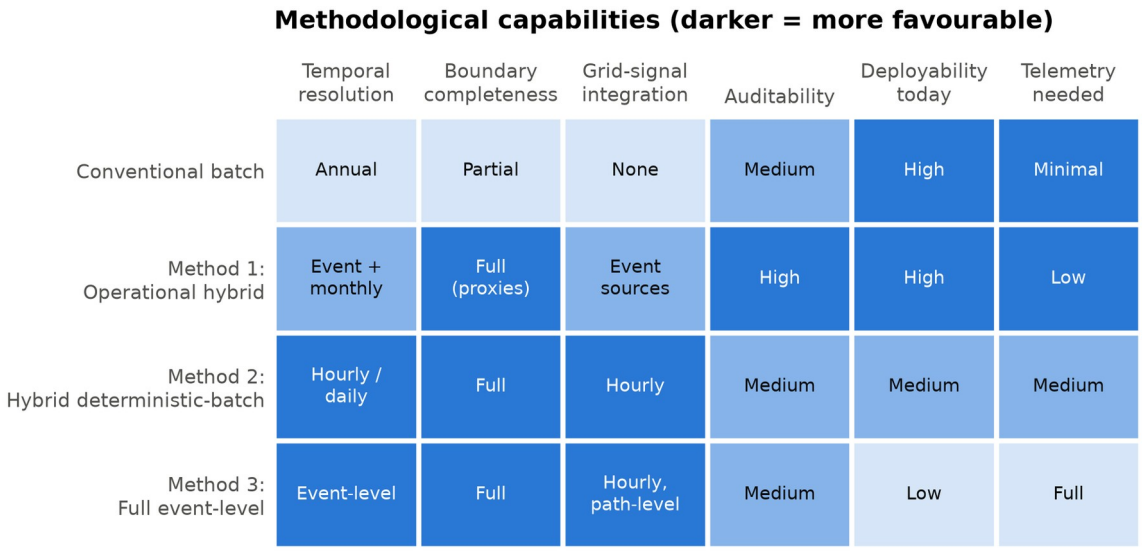


Figure 5 Compares methodological capabilities across existing and proposed approaches.

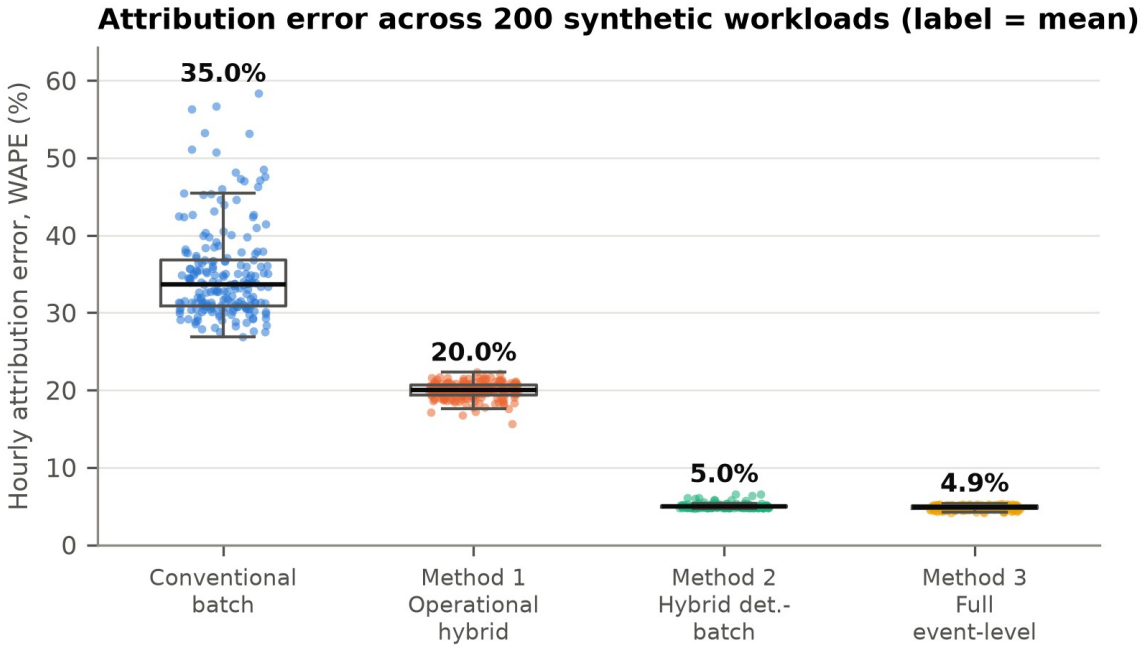


Figure 6 The resulting attribution error variability across workloads.

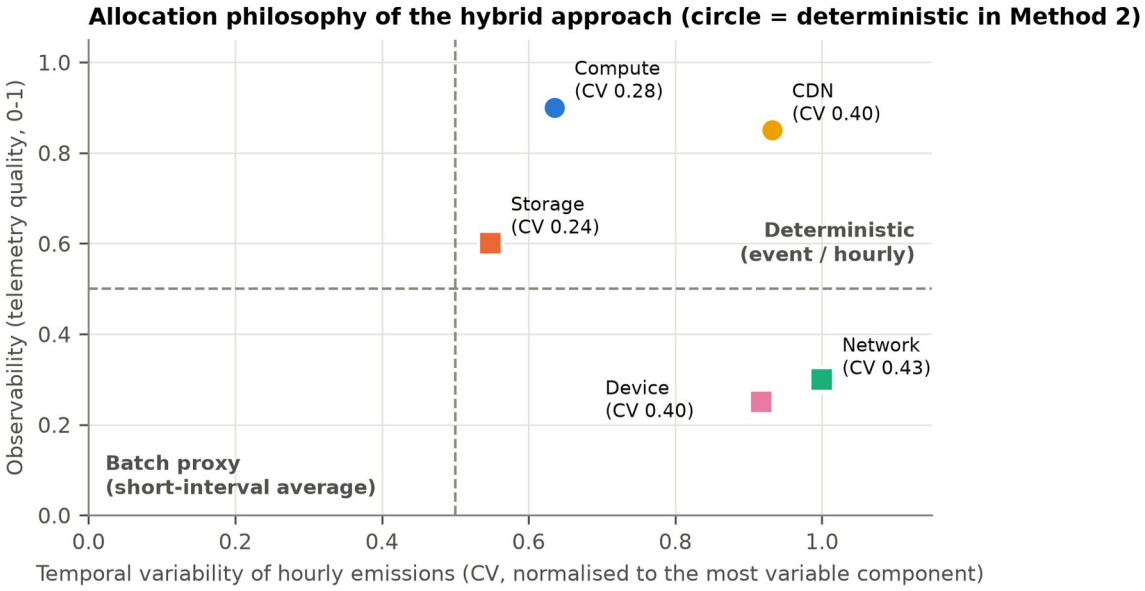


Figure 7 illustrates the allocation philosophy underpinning the hybrid approach.

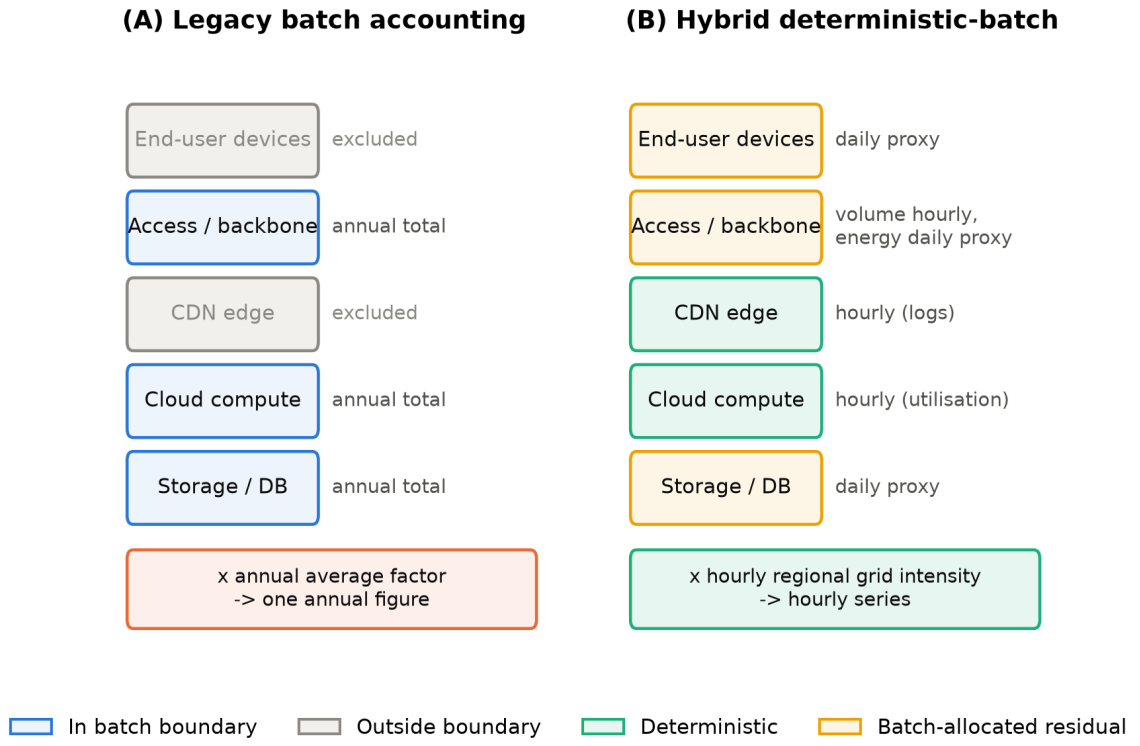


Figure 8A and 8B. Comparison of legacy batch-based internet carbon accounting and the proposed hybrid deterministic–batch framework, highlighting differences in system boundary coverage, temporal resolution, and attribution logic.

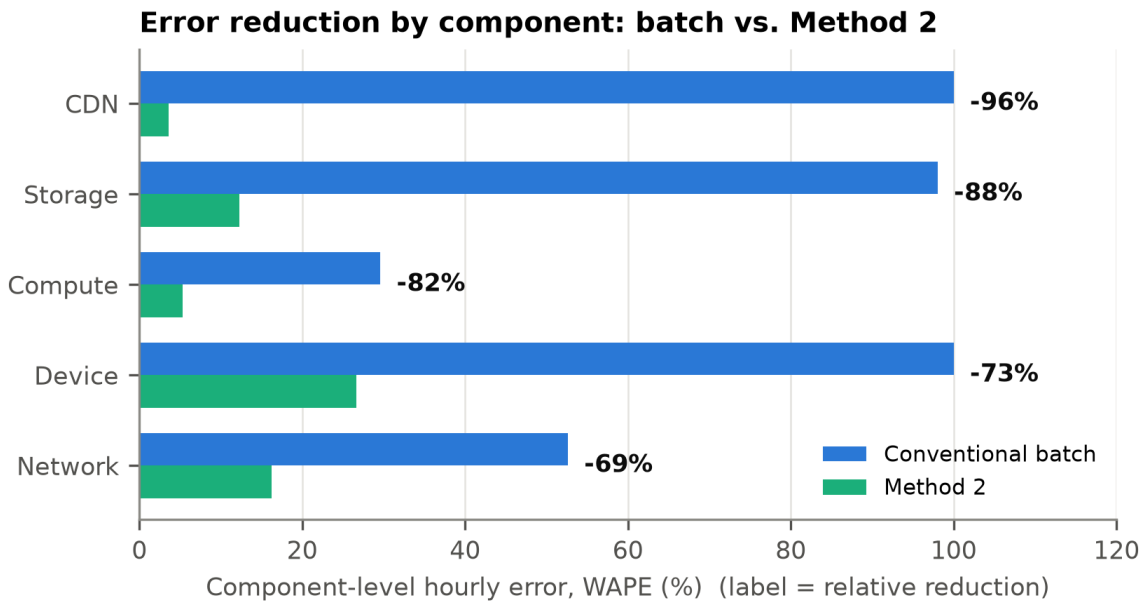


Figure 9 the largest accuracy gains arise from deterministic treatment of compute and CDN components.

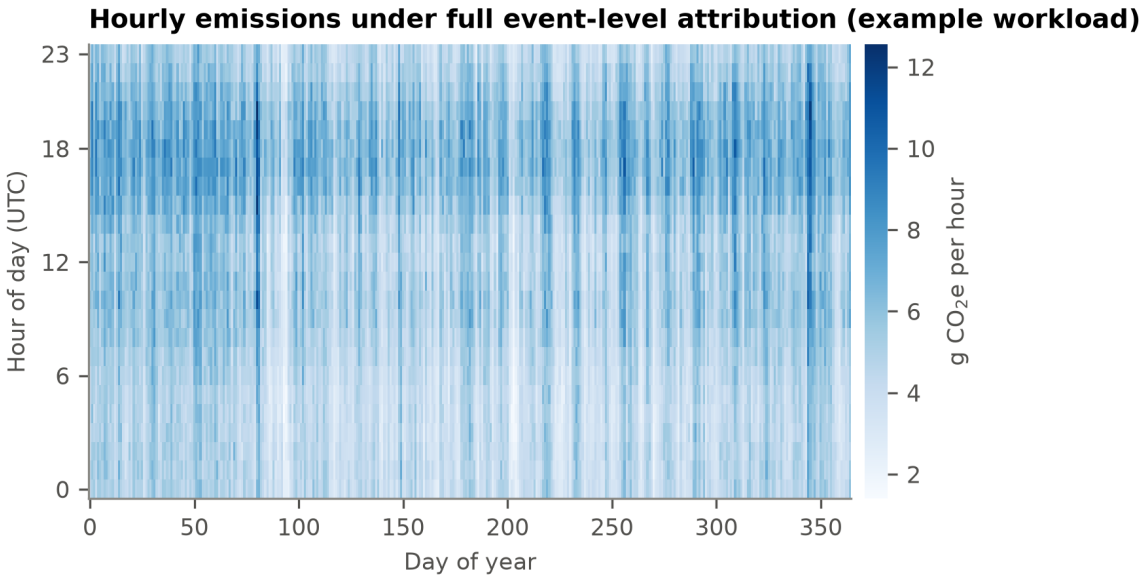


Figure 10 illustrates the temporal granularity achievable under full event-level attribution.

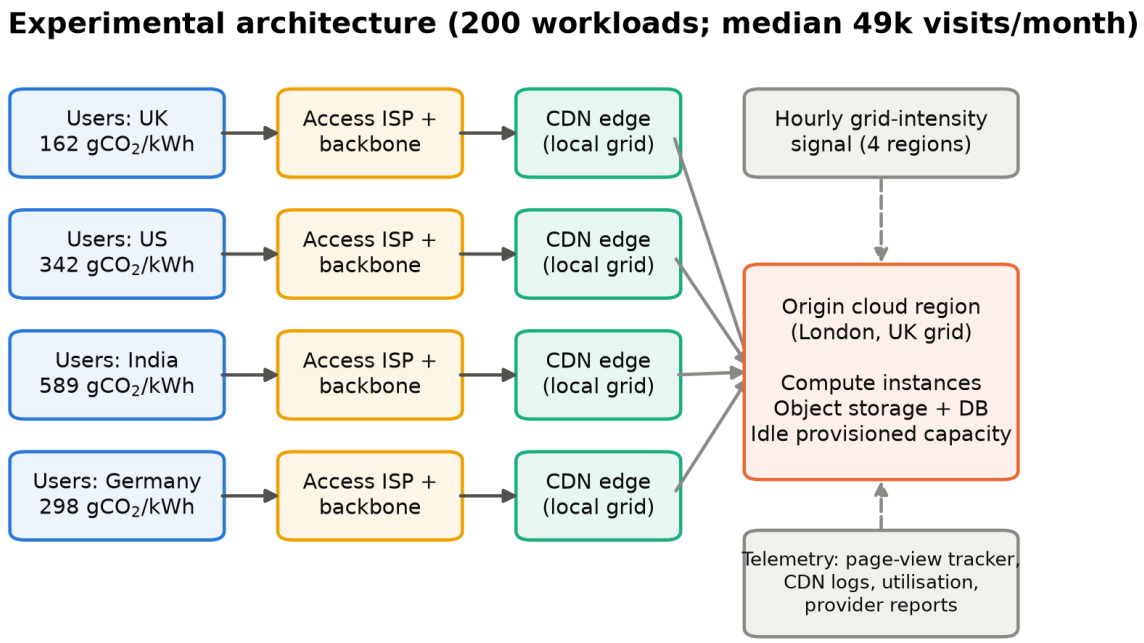


Figure 11 The experimental architecture used to generate representative pilot data.

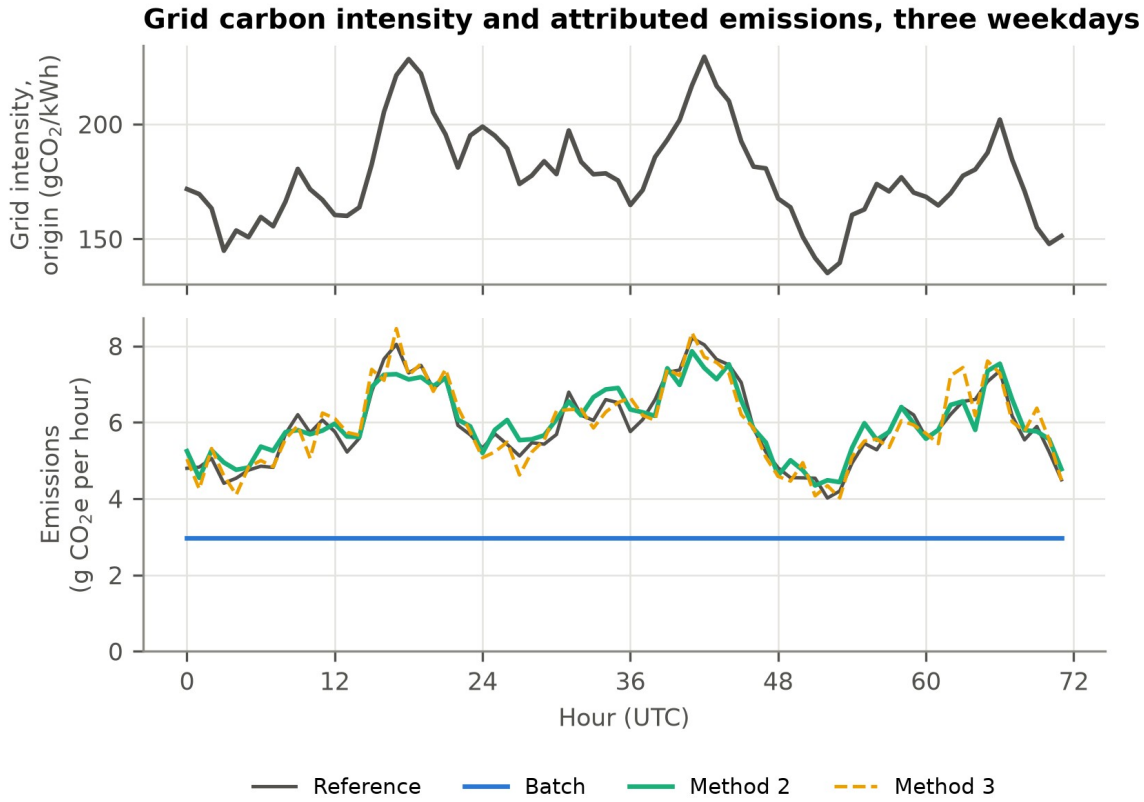


Figure 12 shows the coupling between grid carbon intensity and attributed emissions under near-real-time methods.

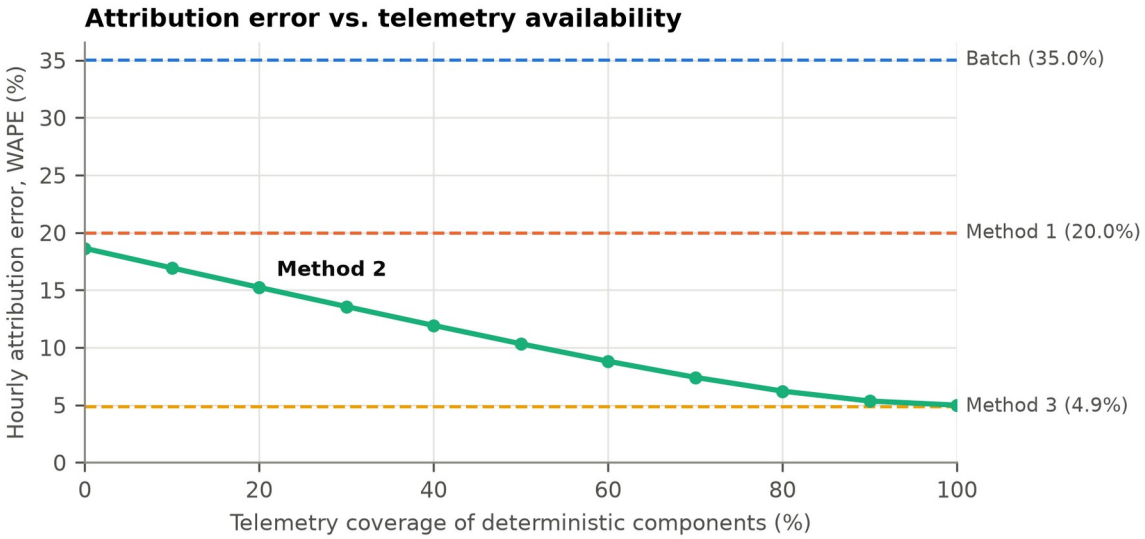


Figure 13 highlights why the hybrid approach captures most achievable accuracy gains under current telemetry constraints.

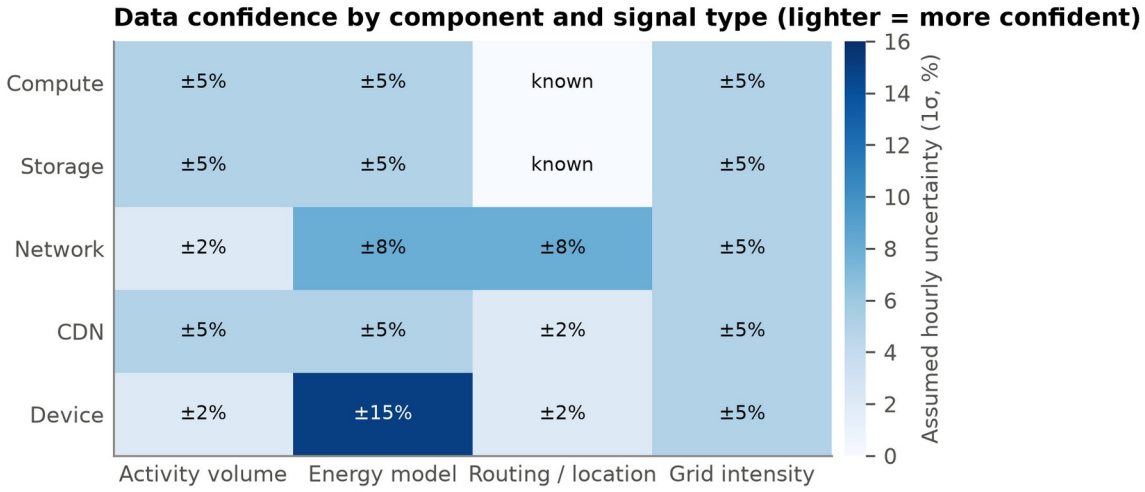


Figure 14 Heatmap of data confidence across components and signal types.

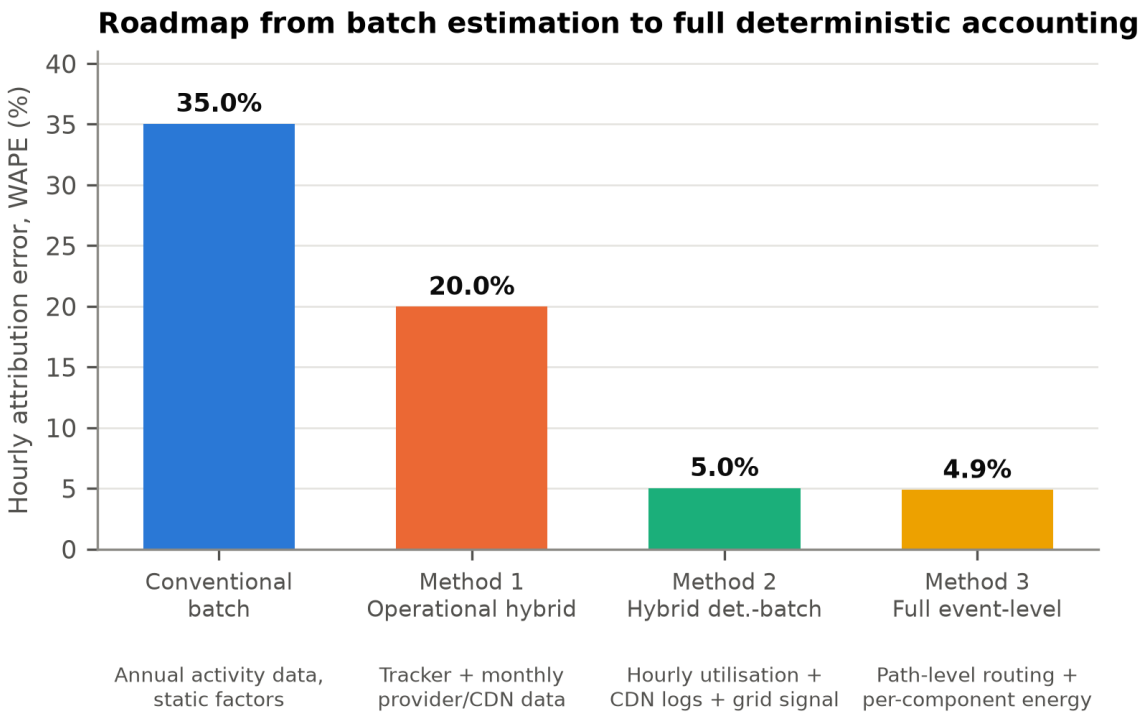


Figure 15 A practical roadmap from batch estimation to full deterministic accounting.