CarbonKit

Working Standard v0.1

Architecture, rollout and research reviews

Proposal for discussion | 1 November 2026

Copyright © 2026 Stephen Makonin

An open platform for personal carbon understanding. A shared language for tools, transparent methods and voluntary participation.

Research foundation: Laura Guzman, Stephen Makonin and Alex Clapp (2019), CarbonKit: Designing A Personal Carbon Tracking Platform. Fourth International Workshop on Social Sensing, Montreal. DOI: 10.1145/3313294.3313385.

This document develops new proposals from the paper. It does not report a production CarbonKit platform, a ratified standard or demonstrated behavioural, health or emissions benefits.

Project: https://github.com/smakonin/CarbonKit

Contact: Stephen Makonin, smakonin@sfu.ca

Contents: Working Standard (record model, calculation, corrections, consent and conformance); reference architecture; StreamingCalc candidate; staged rollout; volunteer roles; paper reviews, including the 2014 Consumer Bill of Rights; domain checklist.

Copyright © 2026 Stephen Makonin

Purpose and status

Working Standard v0.1 is a proposal for discussion, dated 1 November 2026. It is not a ratified standard, certification programme, implemented API or production service. MUST, SHOULD and MAY below express proposed conformance requirements only. No implementation is currently certified against this draft.

CarbonKit began with the 2019 design proposal by Laura Guzman, Stephen Makonin and Alex Clapp. This draft develops an interoperability profile around its common API, personal dashboard, voluntary participation and privacy principles. The paper does not specify the record structures, calculation rules, governance or test suite proposed here.

A Consumer Bill of Rights for Energy Conservation (Makonin, Guzman Flores, Gill, Clapp, Bartram and Gill, 2014) is an additional informative design reference for access, privacy, security, equality, comfort and staged participation. Its proposed rights inform the discussion; this draft does not claim that they are enacted law or already implemented by CarbonKit.

The core scope is voluntary personal activity tracking and transparent emissions estimates. Health and fitness, money-saving tools, social functions and incentives are optional extensions. Each extension needs separate consent and a separate review before implementation.

- Inventory describes estimated emissions associated with recorded activity within a declared boundary.
- A baseline describes a counterfactual or reference scenario; a comparison describes a difference. Neither is an additional inventory item.
- Points, badges and rewards are separate from emissions, money, verified offsets and tradable allowances. No credit issuance, trading or offset verification is defined by v0.1.

Copyright © 2026 Stephen Makonin

Record model and ownership

Providers MUST retain their underlying raw records. CarbonKit coordination MUST be limited to connections, permissions, preferences, goals, progress and the minimum derived summaries needed for the dashboard. Detailed records in this profile are exchanged between an authorised adapter and a private user context; their existence does not authorise central collection.

A separately consented user-controlled vault MAY hold a user export or reproducibility bundle. That vault is an explicit extension to the paper's provider-retention model, not the default CarbonKit database. Its operator, access, retention, recovery and deletion rules MUST be disclosed.

Record	Required meaning
Activity	A time-bounded observation or estimate: stable ID; pairwise subject reference; provider and source reference; start/end; category; quantity and unit; observed/estimated/missing quality; coverage; boundary; creation time; consent reference; optional supersession.
Estimate	A derived amount in kgCO2e linked to activity IDs; inventory/baseline/comparison kind; complete factor snapshots and sources; method/version/expression; input digest; rounding; boundary; uncertainty; limitations; creation time; optional supersession.
Consent receipt	Provider, requesting client, recipient, purpose, scopes, grant/expiry times, retention policy/version, revocation time and deletion/export instructions. A receipt is evidence of a decision, not a bearer credential.
Summary	Period, boundary, included record IDs or private manifest reference, calculation version, incomplete/missing coverage, total and generation time. Central summaries MUST NOT expose private evidence URLs or tokens.

IDs MUST be opaque, stable within their declared scope and free of direct identifiers. A source reference MUST support deduplication within a provider. Pairwise subject references SHOULD differ by provider/client. Pseudonymisation does not make a record anonymous. Private evidence references MUST require authorisation even if the reference is discovered.

Copyright © 2026 Stephen Makonin

Calculation and provenance

- An activity MUST use an explicit unit. Version 0.1 core examples use kWh; other units require a documented method profile and dimensional validation. Carbon output MUST use kgCO2e, not an unlabelled carbon score.
- An estimate MUST retain the exact input values or a recoverable authorised snapshot, their units and a canonical input digest. The method MUST identify the exact input objects and a versioned serialisation rule, including key order, string escaping, number representation and encoding. A digest alone is not sufficient to reproduce a result.
- Each factor MUST record its identifier, version, value and unit, source URI, source-artifact digest, geographical scope, valid period, gas coverage and global-warming-potential basis. Invented factors MUST be visibly labelled synthetic and MUST NOT enter real user inventories.
- The method MUST record the implementation version, expression or resolvable algorithm, allocation choices, boundary, conversions and rounding. Round only for display; preserve unrounded calculation values.
- The method MUST check time, geography and unit compatibility before use. A mismatch MUST fail or produce a separately labelled scenario with an explicit override rationale; it MUST NOT silently replace the applicable factor.
- The estimate MUST declare whether it covers operational emissions, embodied emissions or another boundary. Gaps and exclusions MUST be visible. Different boundaries or GWP bases MUST NOT be silently merged.
- Uncertainty MUST state quantified bounds and their interpretation, or explicitly say unknown with a rationale. Missing uncertainty is not zero uncertainty. Methods MUST NOT manufacture precision from coarse assumptions.

A reproducibility bundle SHOULD contain the method version, factor snapshot, authorised input snapshot, conversion/allocation rules and expected output. If the provider cannot make evidence available, the estimate MUST carry a reproducibility limitation and MUST NOT claim full reproducibility.

The synthetic example distributed with this draft is 10 kWh multiplied by an invented 0.1 kgCO2e/kWh factor, yielding 1 kgCO2e. Its geography is EXAMPLE-ONLY. For this ASCII-only fixture, input_sha256 hashes the complete activity object serialised with Python json.dumps(activity, sort_keys=True, separators=(comma, colon), ensure_ascii=False, allow_nan=False), encoded as UTF-8, with no trailing newline. Comma and colon mean the literal single-character separators. This fixture rule is not a general canonicalisation profile for other implementations. The factor artifact digest hashes the exact UTF-8 bytes of examples/synthetic-factor.txt, also without a trailing newline. It is a schema fixture, not an electricity factor recommendation or an estimate of any person's activity.

Copyright © 2026 Stephen Makonin

Missing data, correction and totals

- Missing activity MUST use a null value and quality=missing with a reason. Zero MUST be a reported or explicitly estimated zero with provenance. A null MUST NOT be coerced to zero in totals or charts.
- Time intervals MUST be half-open [start,end), with RFC 3339 timestamps and offsets. End MUST follow start. Summaries MUST declare the reporting timezone and interval. Methods MUST document allocation across interval boundaries.
- Records MUST be immutable revisions. A correction creates a new ID and a supersedes link to the prior record in the same subject/provider lineage. Cycles, forks and cross-subject supersession MUST be rejected or quarantined for resolution.
- Recalculation under a changed factor/method creates a new estimate revision. The system MUST retain the old method/factor provenance while it remains lawfully retained. A current summary MUST include only the active revision. Historical exports MUST identify the requested as-of time.
- Adapters MUST deduplicate retries with a provider-scoped source reference and idempotency key. Unresolved collisions MUST be shown, not automatically summed.
- Overlapping coverage MUST be resolved by a declared aggregation policy. Whole-home electricity and appliance estimates, or measured electricity and modelled streaming electricity, MUST NOT both be counted for the same consumption. Disaggregation MAY annotate the whole, not enlarge it.
- Inventory, baseline and comparison records MUST be stored and presented separately. Baseline subtraction requires matching periods, boundaries, allocation and assumptions and a named comparison method.
- Summaries MUST show coverage and unresolved overlap alongside a total. Partial coverage MUST be labelled partial; the product MUST NOT imply a complete personal footprint.

Deletion may remove prior revisions and their evidence. In that case an authorised history export MUST report the gap without retaining deleted personal data merely to maintain an audit trail. Non-personal aggregate service metrics require a separate, disclosed retention policy.

Copyright © 2026 Stephen Makonin

Consent, security and lifecycle

- Sharing MUST be opt-in and specific to provider, purpose, recipient and data scope. Health, social, money and incentive access MUST be independently selectable. Refusing an extension MUST NOT block the core tracker.
- Grant and renewal screens MUST explain the data requested, retention, expiry, recipients and effects of revocation in plain language. Clients MUST request only the minimum scope needed. Sensitive disclosure and public sharing MUST be off by default.
- Every data operation MUST enforce the current consent and authorisation at the responsible boundary. Remote services MUST enforce authorisation server-side; local components MUST check the locally valid grant before processing. Revoke MUST stop new synchronisation and block subsequent reads under that grant immediately; queued jobs and caches MUST check revocation before use. Previously exported copies cannot be remotely recalled; this limitation MUST be disclosed.
- Transport MUST be encrypted. Private records MUST be encrypted at rest with per-user isolation and a documented key/recovery design. Secrets MUST stay out of browsers, public repositories, URLs and logs. A profile ID MUST NOT function as an access token.
- Audit records MUST capture grant, access, correction, export and deletion events with minimal personal detail. Operations staff MUST use least privilege. Retention periods, backup purge windows and lawful exceptions MUST be disclosed before collection.
- Users MUST be able to export their authorised activity, estimates, factors and provenance in a documented machine-readable format, including coverage, revisions and schema version. Exports MUST omit reusable access tokens.
- Users MUST be able to request deletion and see completion status. The operator MUST publish a deletion target, maximum backup expiry, provider handoff process and exception policy before a pilot; unresolved deletion failures MUST remain visible. Deletion from CarbonKit MUST NOT be presented as deletion from an independent provider unless confirmed.

This public website is static project documentation. It has no account system, activity ingestion, carbon calculator, analytics script or application database. GitHub serves the pages under its own service policies; contacting volunteers uses the visitor's email application.

Copyright © 2026 Stephen Makonin

Interchange, conformance and governance

The proposed transport is HTTPS with UTF-8 JSON and an explicit spec_version. A future API profile should define resource paths, scoped authentication, pagination, idempotency, error codes and retry rules. Those interfaces are design work, not existing endpoints. Version 0.1 here provides record schemas and fixtures only.

Required fields MUST NOT change meaning within a minor version. Breaking changes need a new major profile and a migration guide. The 0.1 shape profile rejects unknown fields and unsupported versions. Extension namespaces are deferred to a later reviewed profile. Clients MUST never silently reinterpret unknown units or categories. JSON Schemas check record shape; semantic checks below are also required.

Proposed acceptance test	Required result
Reproduction	Pinned inputs/factors/method reproduce the expected unrounded result within a declared tolerance; a missing artifact yields a visible limitation.
Units and validity	Incompatible units, out-of-period factors and wrong geography are rejected in ordinary inventory mode.
Missing versus zero	Null remains missing through export, summary and display; an evidenced zero remains zero.
Revision and replay	A corrected record replaces the prior revision in current totals; replaying the same provider source does not add emissions.
Overlap	A whole-home total plus a streaming/appliance estimate does not double count the covered electricity.
Private lifecycle	Cross-user access fails; revocation blocks reads and queued sync; export contains provenance; deletion and backup expiry are demonstrated.
Extension isolation	An account can use core tracking with health/social/reward scopes declined; no hidden sharing occurs.

Conformance claims require an identified implementation version, profile version, test evidence, exceptions and independent review. Passing the distributed shape/fixture checks alone does not establish full conformance. No conformance badge is offered by this draft.

Proposed governance: discuss changes through repository issues and pull requests; publish rationale, compatibility and privacy/method impacts; invite at least one software reviewer and one carbon-method or privacy reviewer for affected changes. Maintainer roles and decision rules remain to be agreed. Volunteer interest does not constitute an appointment or partnership.

Open decisions before v0.2: pilot geography and categories; approved factor sources and licensing; GWP basis; user-vault custody and keys; consent protocol; deletion deadlines; governance; accessibility review; and evidence thresholds for modelled estimates.

Copyright © 2026 Stephen Makonin

Reference architecture: proposed boundaries

This architecture is a staged implementation proposal. The live public site publishes documents only. No user vault, adapter service, consent service or private dashboard described below has been deployed.

Component	Responsibility and boundary
Public project site	Static GitHub Pages content: paper attribution, draft, schemas, examples and volunteer contact. No private records or credentials.
Provider and adapter	Raw data remains with the provider. An authorised adapter maps minimum required observations into a private user context, validates units and coverage, and supplies provenance. It may run provider-side or locally.
Private user context	User device or separately consented vault holds detailed authorised records and evidence. The dashboard reads these within the user boundary. Device loss and key recovery need an explicit design.
Calculation module	Runs locally or in the consented private context. Uses pinned methods and versioned factors; produces immutable estimates with uncertainty and exclusions. No provider-to-provider raw data sharing.
Minimal coordination	Stores connections, consent state, preferences, goals, progress and necessary derived summaries. It does not act as a central raw activity warehouse. Opaque pairwise identifiers link authorised contexts.
Factor/method registry	Versioned non-personal artifacts, source licenses, checksums, units, valid times, geography and method profiles. A factor revision never silently rewrites historical estimates.
Optional extensions	Health, money, social and incentives are separate capabilities, off by default. Each has purpose-limited scopes, separate storage and an agreed review before enablement.

Proposed flow: provider raw record -> consent check -> adapter -> private activity record -> versioned calculation -> private estimate -> minimum derived dashboard summary. Factors flow from the registry into the calculation. Revocation signals reach adapters, calculations, caches and coordination before further processing.

The dashboard must explain what was counted, what is missing, why an estimate changed and which connections have access. It should support the paper's motivate, track, inform, adjust, reward loop without promising a measured behavioural or health effect.

Copyright © 2026 Stephen Makonin

Privacy, operations and failure behaviour

- Threats to test include cross-user data access, token theft, public evidence links, insider access, provider correlation, reidentification from detailed timestamps, duplicate/replayed data and unsafe factor substitution.
- Use short-lived scoped provider tokens where supported, isolated user keys, auditable access and explicit recovery choices. The pilot must document who can decrypt records; encryption alone is not a complete privacy design.
- A provider outage yields stale/partial coverage and last-success time, not fabricated zeros. Failed jobs retry within bounds and use idempotency keys. Provider rate limits and revocation take precedence over retries.
- Validate incoming records before calculation. Quarantine invalid units, impossible intervals, missing provenance and unresolved overlap. User-visible summaries distinguish stale, missing, synthetic and scenario-only data.
- Exports use authenticated, time-limited delivery with an expiry. Logs redact credentials and raw activity. Backups follow the published purge schedule. Restore tests must prove deleted grants do not resume sync.
- A documented incident process must identify an owner, stop affected connectors, revoke credentials, assess exposed data and communicate with affected participants through an approved channel. No incident response team is currently claimed.

Design review gate: a data-flow inventory, threat model, retention/deletion schedule, consent screens, key custody/recovery plan and cross-user access tests must be complete before any real personal data pilot. Applicable legal obligations need jurisdiction-specific review before collection; this draft is a technical proposal, not legal advice.

The 2014 Consumer Bill of Rights for Energy Conservation adds a consumer-centred design reference alongside StreamingCalc. Its proposals support affordable, understandable feedback, local processing, secure communications and conservation choices that respect essential comfort. Before a pilot, review access barriers and options for staged participation and re-entry. These are proposed CarbonKit applications of the paper, not an implemented service or a statement of current law. See the consumer-rights research note for all seven articles and the full citation.

Copyright © 2026 Stephen Makonin

StreamingCalc: a concrete integration candidate

The inspected StreamingCalc main revision is f75d50924d1cc621f40d3a80052e54983787b424 (27 May 2022). It provides a macOS Numbers model for estimating streaming-video electricity consumption, accompanying research material and references. It is not an executable web API or a CarbonKit connector. Its README seeks volunteers for a Python version.

A suitable first adapter would accept duration and streaming assumptions, compute a documented electricity estimate, emit an activity record with quality=estimated and kWh units, and attach a versioned model manifest. A separately reviewed carbon method would apply appropriate electricity factors to create kgCO2e estimates. kWh is not kgCO2e.

Mapping	Proposed treatment
Inputs	Duration, download rate, simultaneous streams, data-centre demand, core/local-network intensity or demand, device power and manufacturing assumptions. Preserve units and the exact model version.
Boundary	Separate data centre, core network, local network, device and embodied proxy assumptions. Geography and electricity-factor coverage may differ by component; do not assign the user's grid factor to all components without justification.
Output	Estimated electricity per stream plus model assumptions/limitations. For carbon, use component-aware factors or declare a limited scenario. No claim of measured household consumption or avoided emissions.
Overlap	If whole-home electricity already includes local network/device use, allocate or annotate those components rather than adding them again. Remote components need a separately declared inventory boundary.
Validation	Audit workbook formulas and sources, then port with permission-compatible attribution; compare outputs against frozen workbook fixtures within a declared tolerance. Reproduce component totals and allocation under multiple stream counts.

StreamingCalc carries an MIT license (copyright 2022 SFU Computational Sustainability Lab); CarbonKit's existing AGPL-3.0 license is preserved. Retain StreamingCalc notices and the requested paper citation in any port, and review copied material's individual licenses. The paper and third-party reference PDFs are not relicensed by this website.

StreamingCalc research citation: S. Makonin, L. U. Marks, R. Przedpelski, A. Rodriguez-Silva and R. ElMallah (2022), Calculating the Carbon Footprint of Streaming Media: Beyond the Myth of Efficiency. Eighth Workshop on Computing within Limits. DOI: 10.21428/bf6fb269.7625cc76. Its historical UK scenario and assumptions are not unqualified current defaults.

Copyright © 2026 Stephen Makonin

Rollout with evidence gates

The following sequence is proposed and depends on volunteer capacity. No delivery dates, participant numbers, funding or current partners are asserted. Advance only when the prior gate has evidence; record what was learnt before expanding scope.

Stage	Deliverable and acceptance gate
0 - Public foundation	Publish this site, paper review, draft, schemas and source. Check links, mobile/keyboard access, Pages deployment and domain HTTPS. Invite issues and volunteers; agree maintainers and review responsibilities.
1 - Interoperability sandbox	Use synthetic records only. Implement record validation, reproducible calculations, correction, duplicate/overlap rules and export. Gate: all core conformance scenarios pass in two independently written adapters; disagreement is documented and resolved.
2 - StreamingCalc adapter	Audit/port the workbook with a method manifest and frozen fixtures. Gate: independent review of formulas, component allocation, units, uncertainty and carbon mapping; repeatable results and clear overlap treatment. Electricity-only output is acceptable until carbon methods are reviewed.
3 - Consented pilot	Choose geography, category and approved factors; complete privacy/security and accessibility reviews. Gate before collection: consent/revoke/export/delete demonstrations, cross-user isolation, backup expiry, incident owner, participant information and approvals required for the chosen setting.
4 - Evaluate and revise	Measure coverage, reproducibility, error/correction rates, lifecycle completion and participant understanding. Publish limitations and changes. Behavioural/health benefits require an appropriate study design and evidence; usage alone does not establish them.
5 - Optional extensions	Consider health, money, social and non-tradable rewards one at a time, with separate consent, method and equity review. Allowances, trading and verified offsets remain outside v0.1 and would require a separate policy and assurance framework.

Pilot operating measures should report provenance completeness, calculation replay success, duplicate/overlap resolution, consent-revocation latency, deletion completion/backup expiry, accessibility task success and participant comprehension. Set targets before recruitment; publish denominator, missingness and method with each result.

Copyright © 2026 Stephen Makonin

Call for volunteers

CarbonKit invites collaborators to help turn a research proposal into a careful, useful open platform. Interest is welcome from people with lived experience and from people learning a field, as well as established specialists. Roles are proposed work areas, not claims of an existing team or funded positions.

Work area	A useful first contribution
Standards and interoperability	Review one record shape, field definition or correction/overlap case; propose a fixture.
Carbon methods	Review units, factor provenance, boundaries, geography and uncertainty for one activity category.
Software	Audit a StreamingCalc formula, prototype an adapter or improve validation and reproducibility.
Accessibility and usability	Try navigation, keyboard use, screen-reader structure and plain-language consent explanations.
Privacy and security	Review a data flow, key/recovery choice, revocation case or deletion/backup procedure.
Community and governance	Help define inclusive participation, documentation, review processes and what success should mean.

Contact Stephen Makonin at smakonin@sfu.ca with your area of interest, a possible first contribution and your availability. The website's volunteer link opens a draft in your email application; it does not send a message automatically. Please do not include sensitive personal data or provider credentials.

Copyright © 2026 Stephen Makonin

What the workshop paper establishes

Source: Laura Guzman, Stephen Makonin and Alex Clapp (2019), "CarbonKit: Designing A Personal Carbon Tracking Platform," Fourth International Workshop on Social Sensing, Montreal, 6 pages. DOI: 10.1145/3313294.3313385. All six supplied PDF pages were read and visually reviewed.

Paper location	Evidence and implication
pp. 1-3; Sections 1-3	Proposes a platform connecting computing applications, information sources and incentives. A shared structure should let different participants integrate services; the site should not reduce CarbonKit to one calculator.
p. 3; Section 3.1 and Figure 3	Motivate, track, inform, adjust and reward form a proposed feedback loop. Goals can vary and change. The loop is a design rationale, not a demonstrated outcome.
pp. 3-4; Section 4.1	Five modules: personal carbon tracking, health and fitness, money-saving tools, social media and incentives. A common API and configurable personal dashboard connect them.
p. 4; Section 4.2	Encrypt communications and individually protect stored information. Keep provider-collected data with the provider; store minimal central information; preserve user control and need-to-know access.
pp. 4-5; Section 5	Begin with voluntary tracking for personal information. Allowances, trading and eco-currency are later hypothetical directions, not existing programme authority.
pp. 5-6; Section 6	British Columbia is a design case identifying possible stakeholder roles. Historical institutions, programmes and brands do not establish current partnerships.

Copyright © 2026 Stephen Makonin

Limits and the resulting draft decisions

The paper is conceptual: it reports no deployed CarbonKit service, interoperable implementation, accuracy evaluation, user trial or demonstrated emissions, behaviour, health or financial benefits. Figure 4 is a dashboard mockup with illustrative values. The smartphone statistics and stakeholder examples are historical.

The paper does not provide data schemas, a calculation methodology, an emissions-factor registry, uncertainty rules, lifecycle procedures or conformance tests. Version 0.1 proposes those missing pieces, and labels them as new design work. Current laws, institutional responsibilities and programme availability were not inferred from the paper.

- Preserved: open integration, common API ambition, user goals, voluntary participation, provider-held raw data, minimal coordination, encrypted communication and user control.
- Sequenced: core carbon tracking first, with the other four modules as optional later capabilities. This is a proposed rollout choice rather than a claim that the paper required optional modules.
- Added for review: record and factor versioning, unit/boundary checks, explicit unknowns, correction/supersession, overlap handling, consent lifecycle, export/deletion and test evidence.
- Excluded from v0.1: offset verification, tradable credits, mandatory participation and health/behaviour benefit claims. Reward points must remain distinguishable from emissions quantities.
- Requires an explicit decision: any detailed user vault beyond provider storage. It must be separately consented and must not become a default central raw-data warehouse.

The supplied ACM paper is retained privately in the project papers folder. This site links the DOI and attributes all coauthors rather than republishing the PDF or its third-party photograph. The project repository's software license does not override the paper's publication notice.

Copyright © 2026 Stephen Makonin

A Consumer Bill of Rights for Energy Conservation

Stephen Makonin, Laura Guzman Flores, Robyn Gill, Roger Alex Clapp, Lyn Bartram and Bob Gill (2014). A Consumer Bill of Rights for Energy Conservation. 2014 IEEE Canada International Humanitarian Technology Conference (IHTC), 6 pages. DOI: 10.1109/IHTC.2014.7147535.

The paper argues that useful energy information and opportunities to conserve should be available regardless of household income. It proposes consumer rights and discusses appliance-level feedback, local processing, informative displays and better appliance information. Its seven articles appear in Section V, pages 5-6. They are the authors' research proposals.

Proposed right	Meaning in the paper
I - Access (p. 5)	Consumers should know appliance power demands and signatures, including average, peak and multiple-state consumption, through an informative display that does not depend on a computer or Internet access.
II - Privacy (pp. 5-6)	Detailed appliance-use or disaggregated information belongs solely to the consumer; aggregate usage is jointly owned by consumer and producer. The proposal conditions third-party distribution on anonymisation, aggregation and express written consumer consent.
III - Security (p. 6)	Smart-meter communications to the utility and devices in the home should use a secure network.
IV - Equality (p. 6)	Every consumer should have an equal opportunity to participate in energy conservation, whatever their socio-economic situation.
V - Comfort (p. 6)	Consumers should be able to use energy for basic comforts without shame or penalisation.
VI - Re-entry (p. 6)	Consumers should be able to re-enter conservation retrofitting programmes and complete improvements in stages as residences and budgets change.
VII - Maximization (p. 6)	Listing these rights should not deny or diminish other consumer rights. Maximization here does not mean maximising energy reduction.

Copyright © 2026 Stephen Makonin

Applying the ideas to CarbonKit

The following are proposed CarbonKit applications of the paper, not API requirements defined by its authors. The 2014 ownership language is a proposed policy position, not a claim about current law. Provider-held storage in CarbonKit describes custody; it does not by itself determine ownership or grant permission to share.

- Access and equality: review affordability, language, literacy, disability, connectivity and rental or mobility barriers before a pilot. Consider accessible, low-bandwidth, offline or assisted feedback; do not assume a household can buy new equipment or maintain a smartphone and Internet subscription.
- Privacy and security: favour local processing and minimum central information. The paper places disaggregation within the home network (p. 3). Document any sharing policy that differs from Article II rather than describing generic opt-in consent as equivalent to that article.
- Useful feedback: show data freshness and limitations, and test whether people understand which activity an estimate describes. The paper emphasises timely, conveniently placed information matched to consumption (p. 4).
- Comfort and re-entry: use voluntary goals and neutral language about essential energy needs. Support staged participation and rejoining as circumstances change, while respecting export and deletion choices.
- Open appliance information: the paper proposes public descriptions of appliance states, power demands and durations (pp. 4-5). Any future CarbonKit registry or adapter should preserve source and method provenance; no existing registry or certification programme is asserted here.

Alongside the 2019 CarbonKit platform proposal and the StreamingCalc integration candidate, this paper supplies a consumer-centred design reference. Its computational illustration and cited feedback studies do not validate a CarbonKit service or guarantee savings for future users. Its BC rates, programmes and policy examples are historical.

The supplied IEEE author manuscript is retained as a private project source. The public website links the publication DOI and credits all six authors rather than republishing the PDF.

Copyright © 2026 Stephen Makonin

Domain connection checklist

Host: GitHub Pages for smakonin/CarbonKit, serving main at the repository root. Target custom domain: carbonkit.org. This checklist records intended DNS changes, not confirmation that they have been applied.

Public DNS observed on 1 November 2026: carbonkit.org A records 3.33.251.168 and 15.197.225.128; www CNAME carbonkit.org; nameservers ns41.domaincontrol.com and ns42.domaincontrol.com. No authenticated GoDaddy editing capability was available when this checklist was prepared.

In GoDaddy, open the domain portfolio, select carbonkit.org, and open DNS. Save the current zone records before changing them. Configure the GitHub custom domain before switching DNS. GitHub recommends domain verification first; use the TXT value generated by the account's Pages settings rather than inventing one.

Record	Intended value
A @	185.199.108.153
A @	185.199.109.153
A @	185.199.110.153
A @	185.199.111.153
CNAME www	smakonin.github.io (no /CarbonKit path)

Replace the two old apex A values with the four GitHub Pages values above; replace the current www target with smakonin.github.io. Keep unrelated MX, TXT and other service records. Check for conflicting AAAA/forwarding records; change only conflicts with this website target. Leave nameservers in place. Use GoDaddy's default TTL or 1 hour where offered.

After public DNS resolves to GitHub, check repository Settings > Pages for successful DNS and certificate provisioning, then enable Enforce HTTPS when available. Check https://carbonkit.org and https://www.carbonkit.org, the canonical redirect, the document download and a nested page. GitHub notes DNS changes may take up to 24 hours.

Rollback, if needed: restore the saved prior A and www values and review GitHub custom-domain settings. A CNAME file or a GitHub setting alone does not change GoDaddy DNS. Domain completion requires successful DNS and HTTPS checks.

Official sources: GitHub Pages custom-domain management; GoDaddy Help, edit an A record and edit a CNAME record. The web appendix links these instructions. DNS values were verified against GitHub documentation on 1 November 2026.

Copyright © 2026 Stephen Makonin