

EVERYONE EATS

A Systems Architecture for the End of Hunger

V2.1 THESIS The missing piece is not another food program. It is a common protocol that lets independent institutions coordinate supply, need, freshness, fulfillment, and waste - without surrendering ownership of their operations or data.

Jebb Filz - The Architect

Version 2.1 | September 2026 | Evidence + Protocol Release

PROTOCOL FIRST. EVIDENCE SECOND. INFRASTRUCTURE THIRD.

EXECUTIVE SUMMARY

Everyone Eats is an open interoperability protocol and implementation framework for food-system coordination. It is not a proposal for one company, one government, or one algorithm to control food. It is a way for farms, grocers, schools, food banks, kitchens, clinics, processors, carriers, and public programs to exchange the minimum operational signals needed to coordinate while remaining independent.

Version 2.1 moves the concept from architecture toward testable specification. It adds a verified evidence ledger, a concrete v0.1 data model and API surface, a minimum viable network implementation plan, a financial model that separates cash costs from social benefits, and a formal evidence ladder so modeled scenarios cannot masquerade as demonstrated outcomes.

v2.1 upgrade	What changes
Evidence ledger	High-salience quantitative claims are tied to current or clearly dated primary/authoritative sources.
Protocol schema	Canonical objects, required fields, lifecycle states, conformance rules, and example endpoints are specified.
Financial model	Editable workbook maps payer, cost category, timing, scenario, and measurable operating benefits.
Pilot design	90-day baseline + 9-month intervention; capital spending follows measured constraints.
Evidence ladder	Every important claim is labeled as observed fact, external evidence, pilot hypothesis, modeled scenario, or normative design choice.

THE CLAIM Everyone Eats does not need to prove that software can end hunger. Phase 0 needs to prove something narrower: shared, auditable coordination can improve measurable performance across existing institutions.

I. EVIDENCE BASELINE - WHAT WE CAN SAY NOW

v2.1 deliberately narrows the factual foundation to claims that can be stated precisely. Global and U.S. measures use different definitions and populations; they should not be blended into one headline number.

Claim	Status / interpretation	Source
Global hunger	About 673 million people, 8.2% of the global population, experienced hunger in 2024. FAO reports a global decline from 2023, with regional deterioration in much of Africa and western Asia.	FAO, SOFI 2025
U.S. household food insecurity	13.7% of U.S. households - 18.3 million households - were food insecure at some time in 2024. 47.9 million people lived in food-insecure households.	USDA ERS, 2024 report
Global food waste	1.05 billion tonnes were wasted at	UNEP Food Waste Index 2024

Claim	Status / interpretation	Source
	retail, food-service, and household levels in 2022, equal to about 19% of food available to consumers.	
Global food loss	FAO reports 13.2% of food is lost after harvest and before retail. This is distinct from UNEP consumer-level food waste.	FAO / UNEP
U.S. retail + consumer food loss	USDA ERS estimated 31% (133 billion lb) of the available U.S. food supply at retail and consumer levels went uneaten in 2010. This is useful historical evidence, not a current 2026 estimate.	USDA ERS EIB-121
Nutrition incentives	USDA evaluations show incentives can increase fruit/vegetable spending; evidence on consumption varies by program and evaluation. GusNIP continues to fund and evaluate these models.	USDA FNS / NIFA
Free school meals in Finland	Finland identifies 1948 as the start of free school meals for pupils nationwide.	Finnish National Agency for Education

EVIDENCE DISCIPLINE A sourced fact is not automatically evidence for Everyone Eats. External evidence establishes that the problem or component exists. The pilot must still test whether this particular coordination architecture improves outcomes.

Evidence ladder

Level	Meaning	Allowed language
A - Observed pilot fact	Measured directly in the Everyone Eats pilot.	"The pilot measured..."
B - External demonstrated evidence	Measured by an independent study/program relevant to a component.	"External evidence shows..."
C - Pilot hypothesis	Target set before intervention.	"The pilot will test whether..."
D - Modeled scenario	Calculated from explicit assumptions.	"Under these assumptions..."
E - Design choice / value	Governance or normative decision, not empirical fact.	"The architecture requires..."

II. THE FIVE-LAYER ECOSYSTEM

Layer 01 - Production

Treat local and regional production as pluggable supply nodes. Controlled-environment agriculture, greenhouses, urban plots, conventional farms, aquaculture, mushrooms, eggs, and broadacre agriculture participate where their unit economics and environmental constraints make sense.

Layer 02 - Distribution

Expose inventory, capacity, time windows, cold-chain requirements, and need. Existing carriers, kitchens, pickup sites, food banks, and new micro-hubs can compete or cooperate to fulfill transfers.

Layer 03 - Nutrition

Start with a universal operational baseline: culturally appropriate food, allergies, household composition, and dietary constraints can improve matching without biometrics. Clinical or wearable data is optional and separately governed.

Layer 04 - Waste

Represent surplus and spoilage risk as routable states. Human consumption is preferred where safe and appropriate; processing, animal feed, composting, and digestion become downstream options rather than invisible disposal.

Layer 05 - Economics and policy

Public benefits, institutional procurement, philanthropy, grants, commercial sales, and community programs can coexist. The protocol coordinates transactions; local governance decides eligibility and funding rules.

III. EVERYONE EATS PROTOCOL v0.1

The protocol is a common language, not a central brain. Conforming implementations may use different databases, vendors, optimization methods, and user interfaces as long as they preserve the required semantics, safety boundaries, audit trail, and portability.

Object	Purpose	Minimum fields
SupplyOffer	Declares food or capacity available for transfer.	offer_id; actor_id; item_class; quantity; unit; available_from/to; location_zone; storage_class; use_by; disposition; status
DemandRequest	Declares a quantified need.	request_id; actor_id; item_class; quantity; unit; need_by; location_zone; constraints; priority_basis; status
FreshnessScore	Records an interpretable freshness/spoilage estimate.	score_id; object_id; method; observed_at; score; confidence; evidence_ref
TransferRequest	Proposes movement from supply to demand.	transfer_id; offer_id; request_id; quantity; pickup/dropoff window; carrier; safety_rules; status
FulfillmentReceipt	Closes the transaction with what actually happened.	receipt_id; transfer_id; accepted_quantity; delivered_at; condition; exceptions; signer

Object	Purpose	Minimum fields
WasteEvent	Records food that exits human-consumption flow.	event_id; object_id; quantity; reason; destination; occurred_at
OverrideEvent	Records a human rejection/change of a system recommendation.	override_id; recommendation_id; actor_id; reason_code; free_text_optional; occurred_at

Lifecycle states

Objects use explicit state transitions so audits can reconstruct decisions. A minimal TransferRequest lifecycle is: PROPOSED -> ACCEPTED -> PICKED_UP -> DELIVERED -> CLOSED, with CANCELLED and REJECTED terminal alternatives. State transitions are timestamped and attributed to an actor.

Reference API surface

Method + path	Function	Notes
POST /v0.1/supply-offers	Publish supply	Idempotency key required.
POST /v0.1/demand-requests	Publish need	No person-level identity required for institutional demand.
POST /v0.1/matches	Request/propose match	Returns explanation fields and constraint checks.
POST /v0.1/transfers/{id}/accept	Accept proposed transfer	Human or authorized institutional system.
POST /v0.1/transfers/{id}/events	Record lifecycle event	Append-only event semantics.
POST /v0.1/overrides	Reject/modify recommendation	Reason code required; narrative optional.
GET /v0.1/audit/{id}	Reconstruct transaction	Role-based view; privacy filtering applied.
GET /v0.1/metrics	Aggregate operational metrics	No sensitive person-level public output.

CONFORMANCE RULE A vendor does not become "Everyone Eats" by implementing the endpoints. Conformance requires data portability, documented state transitions, reconstructable decisions, minimum safety checks, role-based access, and successful interoperability testing with at least one independent implementation.

Decision stack

Layer	Use	Default rule
Deterministic rules	Food safety, authorization, eligibility, required fields, hard constraints.	Must be explicit and testable.

Layer	Use	Default rule
Optimization	Routing, allocation, scheduling, capacity, spoilage minimization.	Objective function and constraints documented.
Machine learning	Forecast uncertain demand, supply, travel time, spoilage risk.	Only where performance beats a simpler baseline.
Generative AI	Operator assistance, summarization, documentation.	Never final authority for food safety, eligibility, or clinical decisions.

IV. GOVERNANCE AND TRUST

Participation. Open standards; no single required vendor.

Data boundaries. Minimum necessary operational exchange; sensitive personal data remains local or explicitly consented.

Dispute resolution. Local institutions can reject recommendations without punishment.

Versioning. Public change logs, review periods, semantic versioning, migration guidance.

Auditability. Fulfillment, waste, latency, cost, parity, and overrides are reconstructable.

Exit and fork rights. Participants can export data and implement the standard elsewhere.

Stewardship rotation. Term-limited, multi-stakeholder stewardship including recipient-community representation.

Safety boundary. Validated rules and accountable humans control food safety, clinical advice, and eligibility.

V. PHASE 0 - MINIMUM VIABLE NETWORK

Phase 0 tests coordination before capital intensity. The preferred pilot is one bounded service area over twelve months: 90 days of baseline observation followed by nine months of intervention.

Role	Illustrative minimum
Supply	3-5 farms, grocers, distributors, kitchens, or processors
Demand	2 schools + 1 food-bank/pantry network + 1 senior/community institution
Distribution	Existing carriers or partner fleets; one or more cold-storage nodes
Evaluation	Independent university, research group, or evaluator
Governance	Recipient/community representative + operators + public-health/food-safety + technical steward
Technology	Reference API, adapters, dashboard, rules/optimizer, audit/event store

90-day baseline

- Measure supply offered, demand requested, fulfilled quantity, rejected quantity, waste, latency, delivery cost, staff time, reliability, and neighborhood/service-area parity.

- Do not optimize during baseline except where safety or ordinary operations require intervention.
- Publish metric definitions before intervention so the denominator cannot change after results are known.

Nine-month intervention

- Activate standardized supply/demand exchange and matching.
- Start with deterministic rules + transparent optimization.
- Introduce ML forecasting only after a simple benchmark exists and only if it improves out-of-sample performance.
- Track partner burden; a system that saves food by creating clerical work can still fail.
- Publish negative and null results.

VI. MEASUREMENT AND FALSIFIABILITY

Metric	Definition	Primary question
Fulfillment rate	Accepted nutritious quantity / eligible requested quantity	Did more legitimate demand get fulfilled?
Preventable waste rate	Eligible human-edible quantity wasted / eligible supply handled	Did less edible food exit the human food pathway?
Median match-to-delivery time	Median elapsed time from accepted match to confirmed delivery	Did coordination become faster?
Cost per fulfilled unit	Relevant operating cost / standardized fulfilled quantity	Did the network become cheaper to operate?
Partner burden	Staff minutes per completed transfer + exception workload	Did software reduce or increase work?
Reliability	Completed transfers / accepted transfers	Can partners trust the network?
Coverage parity	Performance dispersion across predefined service areas/groups	Did improvement reach the whole pilot area?
Override rate	Human overrides / recommendations	Where is the system systematically wrong?

TARGETS v2.1 does not hard-code 25%, 30%, or 40% improvement claims. Targets should be set after baseline measurement and before intervention analysis, with the rationale and statistical plan published in advance.

Stop / redesign triggers

- Cost per fulfilled unit remains flat or rises beyond the agreed learning period.
- Reliability is worse than baseline.
- Partner administrative burden materially increases.
- Coverage parity deteriorates.
- Privacy or surveillance complaints exceed a published threshold.
- Safety/access decisions cannot be reconstructed.
- A vendor becomes operationally indispensable despite portability requirements.

VII. ECONOMICS - CASH FLOW BEFORE GRAND TOTALS

The accompanying workbook is the v2.1 financial backbone. It intentionally does not claim that avoided healthcare spending automatically funds food-system operations. Costs and benefits are assigned to actors and time periods. Social value can be real while cash savings accrue somewhere else.

Category	Examples	Treatment
Direct pilot cash cost	Staff, integration, hosting, evaluation, legal/privacy, training, contingency	Budgeted cash outflow
Partner in-kind contribution	Existing staff time, vehicles, storage, donated food	Tracked separately; not treated as free
Operating efficiency	Reduced waste disposal, fewer emergency transfers, lower routing miles, less staff coordination time	Count only when measured and monetizable
Transfers	Food assistance or household food-spend reduction	Reported separately from system efficiency
Avoided downstream costs	Potential health, absenteeism, or emergency-service effects	Do not book as pilot cash savings without causal evidence and payer capture
Capital expansion	Cold storage, processing, production, hubs	Triggered by measured bottlenecks; outside Phase 0 base case

BUDGET HYPOTHESIS The prior \$500,000-\$2 million Phase 0 range remains a planning envelope, not a validated cost estimate. The workbook forces the user to replace assumptions with quotes and commitments before presenting a funded proposal.

VIII. SCALE ROADMAP - EVIDENCE BEFORE INFRASTRUCTURE

Phase 0 | 0-12 months

Baseline, protocol v0.1, governance charter, adapters, intervention, independent evaluation.

Year 1-2 | Expand network

Add partners and targeted cold-chain/processing capacity only where measured constraints justify it. Publish protocol v1.0 and demonstrate interoperability with an independent implementation.

Year 3 | Integrate infrastructure

Add production, processing, waste recovery, and regional redundancy where unit economics and demand evidence support them. Optional nutrition personalization remains separately consented and governed.

Year 5 | City-scale proof

Set city-scale targets from accumulated evidence. Publish a transparent payer/savings map and distinguish measured outcomes from modeled scenarios.

Year 10 | Replication

Success is not one giant platform. Success is multiple communities running interoperable implementations that can exchange standards without surrendering local control.

IX. A DAY IN THE NETWORK - YEAR 10

You wake up. If you have opted into personalized nutrition, your chosen tools may suggest an iron-rich meal based on information you chose to share. If you have not opted in, nothing about your food access changes.

The neighborhood pickup point was restocked before dawn. Some produce came from a controlled-environment farm, some from regional growers, some from commercial surplus. The protocol does not care who owns the supply node. It cares whether a valid SupplyOffer can safely satisfy a real DemandRequest.

Your kids eat breakfast at school. The kitchen can publish aggregate demand, receive supply, and record fulfillment without creating a technology profile for each child.

Across town, a senior center runs ahead of predicted demand. A restaurant has excess root vegetables. The system proposes a transfer. A human accepts it. A FulfillmentReceipt closes the loop. If the recommendation is wrong, the human rejects it and an OverrideEvent records why.

Weather still breaks things. Trucks still fail. Food still spoils. People still disagree. The difference is that failures become visible earlier, alternatives become easier to discover, and institutions can cooperate without being absorbed into one platform.

That is the difference between charity and architecture: not replacing human institutions, but giving them rails on which cooperation can scale.

Everyone eats. Not because one machine controls the food supply. Because thousands of independent actors can finally coordinate without surrendering themselves to the network.

X. BUILD ORDER

1. Define the geography - small enough to measure; large enough for meaningful supply/demand variation.
2. Recruit anchor partners - supply, demand, distribution, evaluation, governance.
3. Measure baseline - 90 days of waste, unmet demand, cost, latency, burden, reliability, and parity.
4. Sign governance - participation, data, audit, dispute, safety, versioning, and exit rules.
5. Freeze protocol v0.1 - object schemas, state transitions, endpoints, and conformance tests.
6. Build adapters - connect existing systems before replacing them.
7. Activate matching - rules and optimization first; ML only when benchmarked.
8. Evaluate independently - publish positive, negative, and null results.
9. Fix bottlenecks - spend capital where measured constraints justify it.
10. Scale by replication - publish the implementation playbook and reference code.

APPENDIX A - PROTOCOL FIELD NOTES

v0.1 uses opaque IDs and coarse location zones by default. Person-level identity is not part of the core supply/demand protocol. Implementations may maintain local identity mappings where legally necessary, but those mappings are outside the interoperable core unless explicitly required for a transaction.

Field principle	Requirement
Time	ISO 8601 timestamps with timezone.
Quantity	Value + explicit unit; conversions logged.
Location	Service zone or facility ID by default; exact address only when fulfillment requires it.

Field principle	Requirement
Item classification	Versioned taxonomy; local labels may map to canonical classes.
Constraints	Machine-readable hard constraints separated from preferences.
Provenance	Every score/recommendation names method/version and observation time.
Privacy	No public endpoint exposes person-level sensitive data.
Audit	Append-only event history for transaction-critical changes.

APPENDIX B - FINANCIAL MODEL LOGIC

The accompanying XLSX contains four sheets: Assumptions, Phase0 Budget, Benefits Map, and Scenarios. Blue values are user-editable assumptions; formulas are black; linked cells are green. The base case is intentionally illustrative. Replace all yellow-highlighted assumptions with local quotes or signed partner commitments before external fundraising.

APPENDIX C - VERIFIED SOURCES

- 1. FAO, SOFI 2025.** Global hunger: 673 million people / 8.2% in 2024.
<https://www.fao.org/newsroom/detail/global-hunger-declines--but-rises-in-africa-and-western-asia--un-report/>
- 2. USDA ERS, Household Food Security in the United States in 2024.** 13.7% / 18.3 million U.S. households food insecure in 2024; 47.9 million people lived in food-insecure households.
<https://ers.usda.gov/publications/113622>
- 3. UNEP, Food Waste Index Report 2024.** 1.05 billion tonnes wasted in 2022; 19% of food available to consumers.
<https://www.unep.org/resources/publication/food-waste-index-report-2024>
- 4. FAO, Food Loss and Food Waste.** 13.2% lost after harvest and before retail; 19% wasted at retail/food service/household levels.
<https://www.fao.org/policy-support/policy-themes/food-loss-and-food-waste/>
- 5. USDA ERS, EIB-121.** 2010 U.S. estimate: 31% / 133 billion lb uneaten at retail and consumer levels.
<https://ers.usda.gov/publications/43836>
- 6. USDA FNS, FINI evaluation.** In three of four study groups, FINI increased monthly household fruit/vegetable expenditures 12-16%; no measurable consumption impact in that evaluation.
<https://www.fna.usda.gov/research/snap/evaluation-fini-interim>
- 7. USDA NIFA, GusNIP.** Program purpose and current federal nutrition-incentive framework.
<https://www.nifa.usda.gov/grants/programs/gus-schumacher-nutrition-incentive-program-gusnip>
- 8. Finnish National Agency for Education.** Free school meals since 1948.
<https://www.oph.fi/en/education-and-qualifications/basic-information-about-primary-and-lower-secondary-education>

EVERYONE EATS v2.1

Open protocol. Measurable pilot. Replaceable implementations. Evidence before scale.