© 2026 the EV-SIENT contributors



EV-SIENT: Right-Sized Public Charging

Yiguedan Philippe Andy Yaro

View GitHub Repository: github.com/andyyaro/EV-SIENT

Let's connect! Email: yyaro@gmu.edu

LinkedIn: linkedin.com/in/andy-yaro



Five of us built EV-SIENT, the Site Evaluation and Navigation Tool: planning support for public electric vehicle (EV) charging in Fairfax City, Fairfax County and Prince William County, Virginia. It screens 12 real candidate sites and sizes how many charging ports each should get now, and later.

It also ships a planner that runs the same model, unmodified, in a web browser on anyone's own sites. Their file never leaves the device.

We are a George Mason University Future City Fellowship team: Anna, Laylas, Andy, Adama and Paula. I am Yiguedan Philippe Andy Yaro. The code is open source under Apache-2.0.

A. Why EV-SIENT Exists



A commuter garage in Prince William County installed 70 charging ports. Seventeen were later taken offline: each station cost about $1,000 a year to keep, and too few people used them.

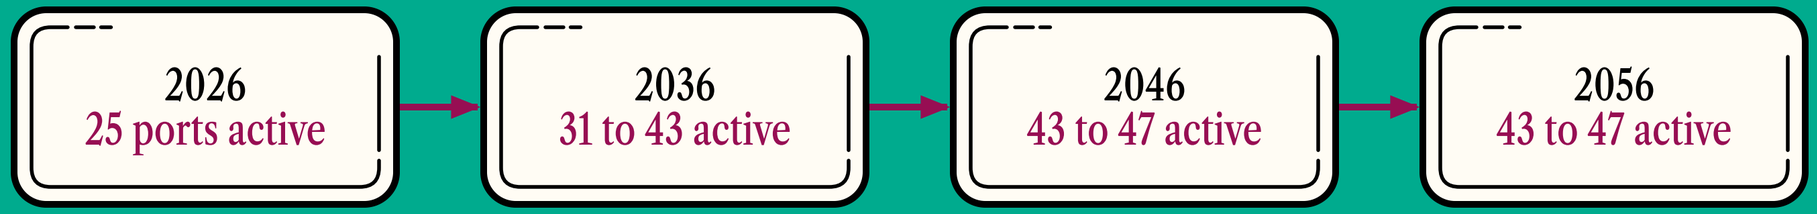
At another county site, a ten-port job ran over budget because the trenching was longer than planned.

Build too much and you pay for idle chargers. Build too little and adding more means digging up ground you already paid to close.

1. Staged, Evidence-Gated Deployment

Conduit and electrical pathway for 47 ports, laid at the 2026 build while the ground is already open

Pathway



Each step up passes a **trigger review**: six months of observed data, 97% uptime, then use or waits above the bar, then a person approves the budget.

Active ports on the base adoption path, then the high path. On the high path the 47-port ceiling binds from 2046.

Net Present Cost: Staged Against a Broad Buildout Now

By	Staged plan	Broad buildout now	Staged avoids
2036	$1.49M	$1.85M	$0.36M
2046	$2.21M	$2.60M	$0.38M
2056	$2.92M	$3.36M	$0.44M

Portfolio net present cost from contract v3, at a 3% real discount rate with base escalation. Screening-class estimates for comparing cases: never a budget, quote or bid.

i

M = million US dollars

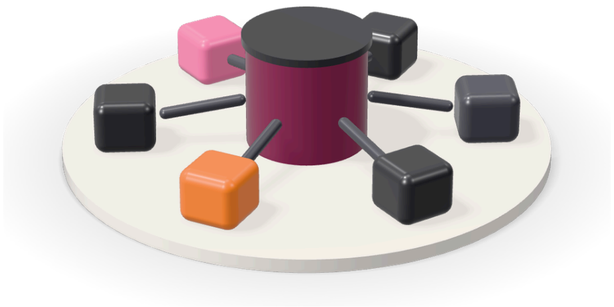
1. Start with 25 ports at the strongest of the 12 sites.
2. Lay conduit for 47 while the trench is open, and energise later.
3. Grow only on observed evidence, and only with a person's approval.

25 and 47 are the official totals, set by decision record DEC-027. The earlier 48-port pathway is history, still reproducible in the legacy mode.

: Built at the start: the pathway

: A step up, only through a trigger review

B. Architecture

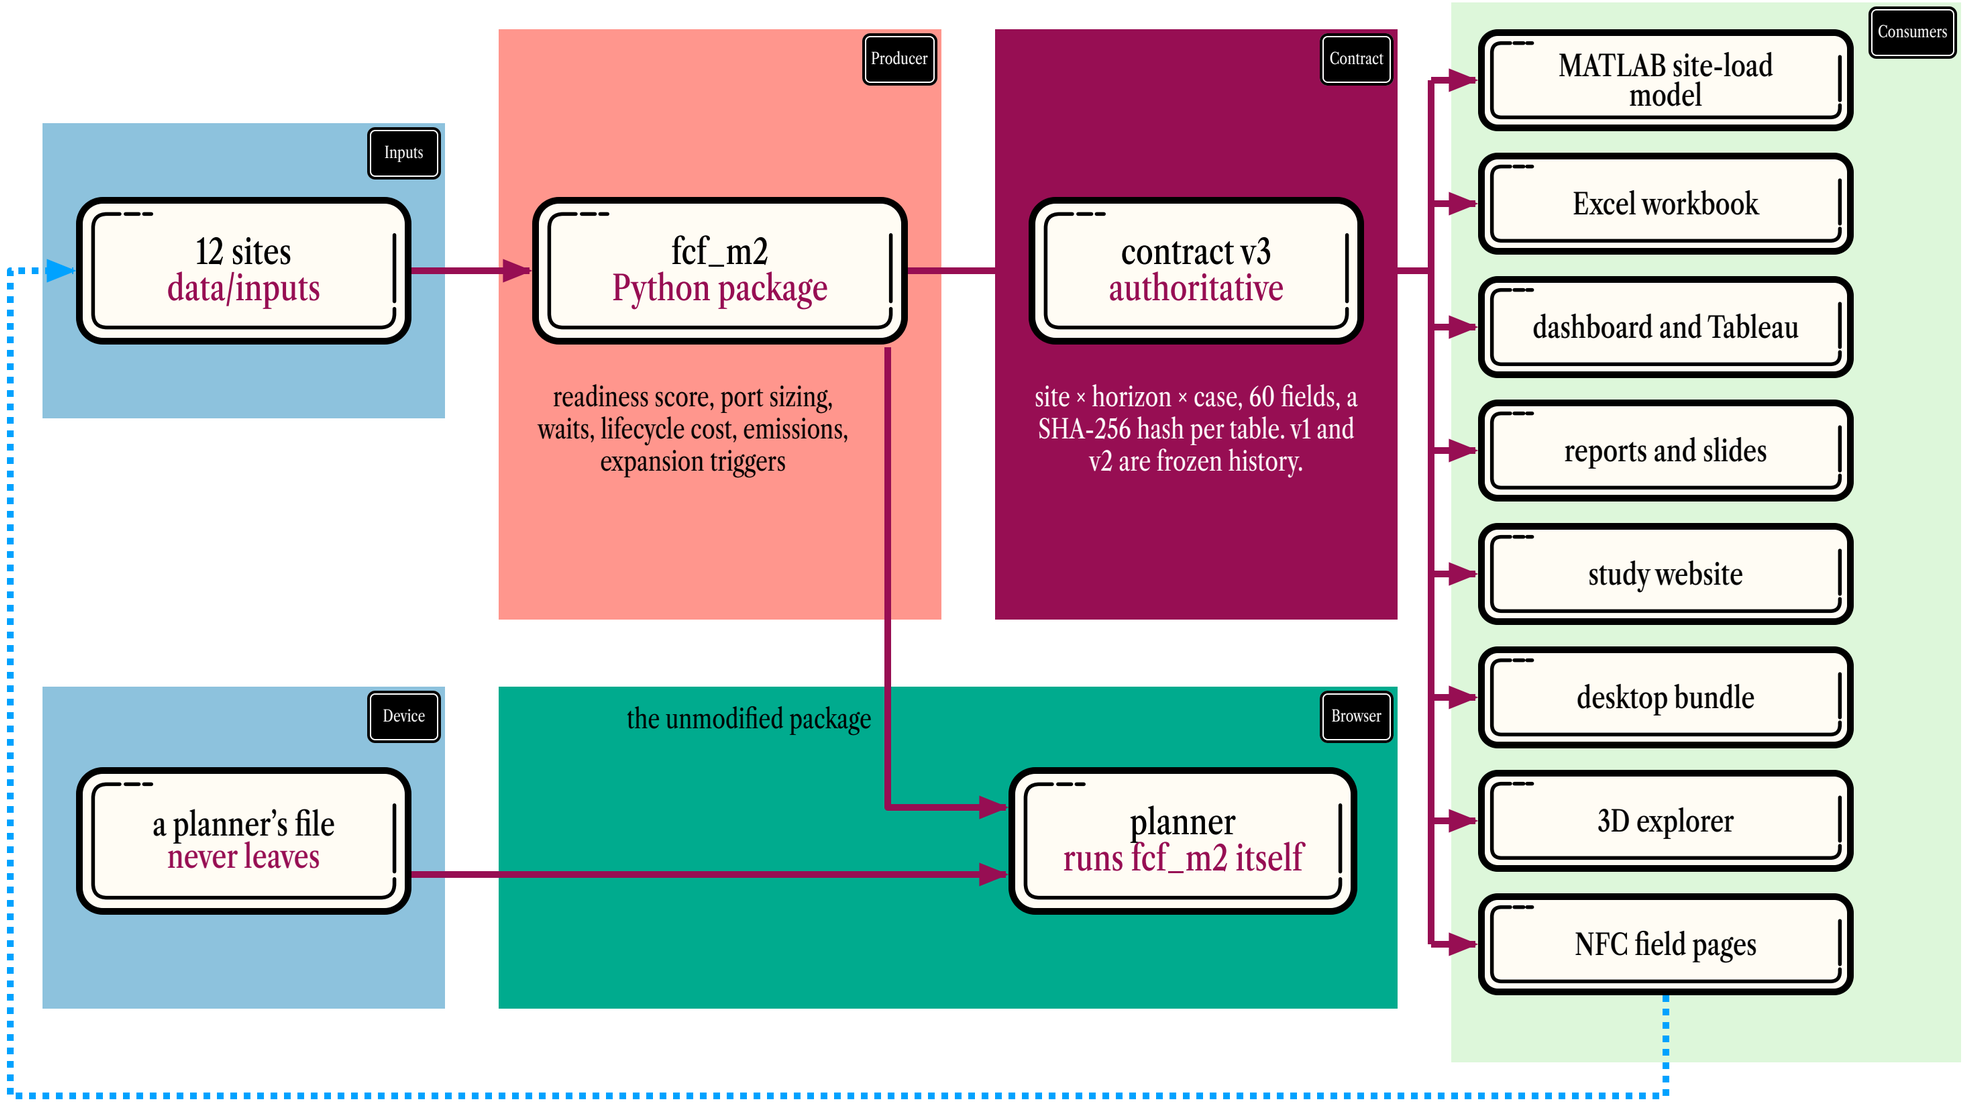


Python is the single producer. It writes one validated data contract, and MATLAB, Excel, Tableau, both websites, the three-dimensional explorer and every chart read it instead of recomputing.

So no two surfaces can disagree about a port count, a year, or whether capacity is a placeholder.

The planner is the one surface that runs the model rather than reading its output.

1. One Producer, Many Consumers



3D = three-dimensional
NFC = near-field communication
SHA = Secure Hash Algorithm

Every consumer quotes the contract and re-derives nothing. capacity_basis travels on every row.

Three surfaces still read the frozen v1 until they are rebuilt: the Excel workbook, the desktop bundle and the published Tableau workbook.

The planner never carries the study's real inputs: it runs on a planner's own sites.

: Inputs : Producer : Contract : Consumers : Browser : Field corrections

2. The Canonical Data Contract

data/contract/v3 · model version 3.1.0

Contract v3

- scenario_table.csv
  120 rows × 60 fields
- site_table.csv
  one canonical spelling per site
- source_provenance.csv
  where each field came from
- standards_register.csv
  which standards apply where
- assumptions.json
  every labelled assumption
- schema.json
  JSON Schema for every row
- contract_manifest.json
  versions and a SHA-256 per table

What the Validator Checks

RT-01	every site_id resolves; one spelling per site
RT-02	12 sites × every horizon and case, no duplicates
RT-03	active ports never fall, never exceed the pathway
RT-04	reference totals: 25 starter, 47 pathway
RT-05	units in range, such as 0 to 24 charging hours
RT-06	*not implemented, and the code says so*
RT-07	placeholder capacity cannot claim high confidence
RT-08	energy delivered never exceeds energy required
RT-09	every emissions row names its factor vintage
RT-10	every table's hash matches the manifest
plus	every row carries a source and a date

Rules for Every Consumer

Join on site_id only; site_name is display text.

An empty observed_* field means not observed. It never means zero.

Quote contract values; never re-derive them.

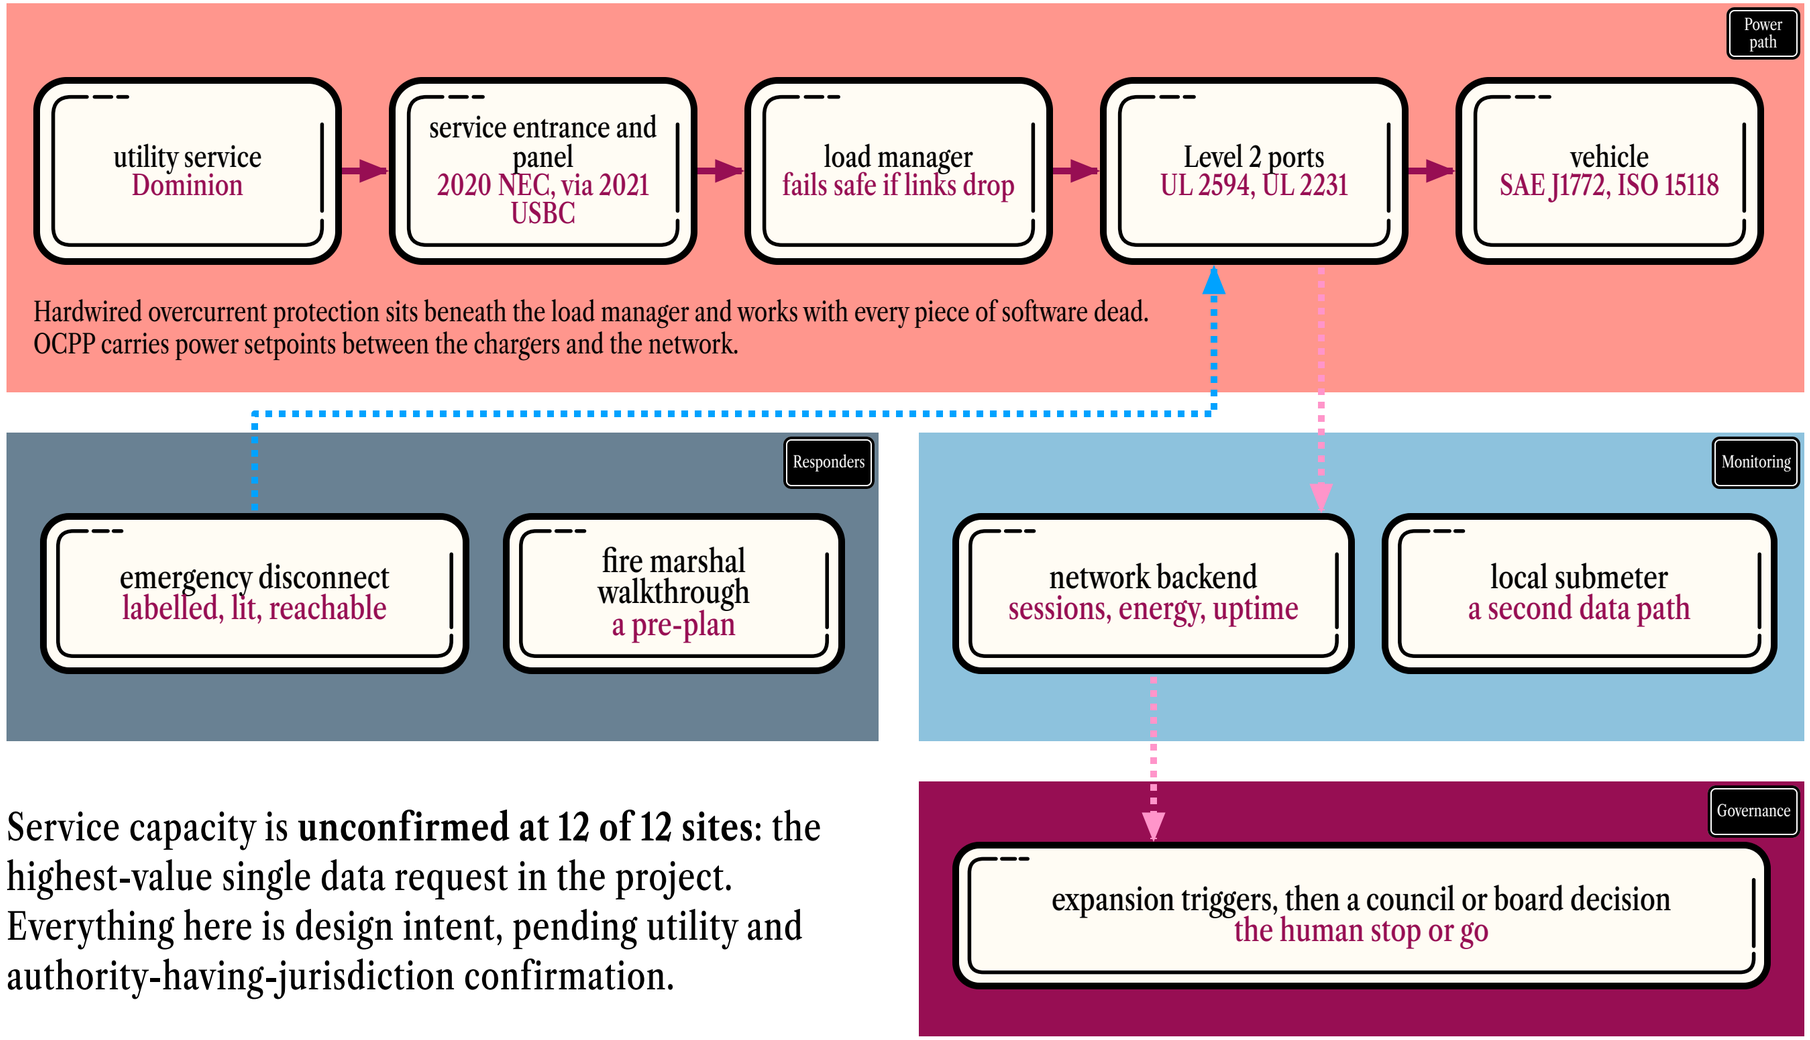
A schema change bumps the version, and consumers assert it.

Today all 120 rows read capacity_basis placeholder, data confidence Low, and trigger status awaiting observed data. Validation passes for v3, v2 and v1.

CSV = comma-separated values
JSON = JavaScript Object Notation
RT = reconciliation test
SHA = Secure Hash Algorithm

: The files under data/contract/v3, one row per site, horizon and case

3. The Physical Chain It Plans For



NEC = National Electrical Code
USBC = Uniform Statewide Building Code
UL = Underwriters Laboratories
SAE = Society of Automotive Engineers
ISO = International Organization for Standardization
OCPP = Open Charge Point Protocol

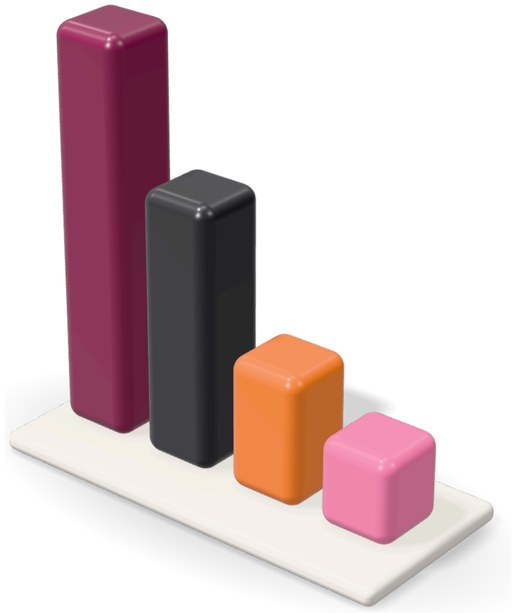
Both jurisdictions are Dominion territory, so interconnection review runs through its Level 2 make-ready program.

Three failures the client has lived through drive the design: overbuilt and underused, cost blown by trenching, and an unclear way to shut power off.

Service capacity is **unconfirmed at 12 of 12 sites**: the highest-value single data request in the project. Everything here is design intent, pending utility and authority-having-jurisdiction confirmation.

: Power path : Responder access : Monitoring : Governance : Emergency isolation : Usage data

C. The Model



Each site gets a readiness score out of 100 from four weighted parts: demand 45, feasibility 30, safety 15, equity 10. Scores use fixed anchors, so adding or removing a site moves no one else's score.

The model then sizes ports with a queueing screen, checks peak waits in a seeded simulation, prices a staged plan against a broad build, and counts emissions avoided.

1. Four Weighted Parts, Fixed Anchors

Demand: 45 Points

Households within 1 mile	8.1
Jobs within 1 mile	8.1
Daily traffic nearby (AADT)	5.9
Amenities within half a mile	5.0
Parking spaces	4.5
Hours open	3.6
Local EV share	3.6
Visit length	3.2
Existing chargers nearby lower is better	3.2

Feasibility: 30 Points

Spare electrical capacity	10.5
Distance to electrical service lower is better	7.5
Estimated install cost lower is better	7.5
Trenching needed lower is better	4.5

Safety: 15 Points

ADA layout possible	5.3
Fire access adequate	5.3
Enclosed garage lower is better	2.3
Emergency disconnect visible	1.5
Stormwater risk lower is better	0.8

Equity: 10 Points

Equity-gap index	10.0

AADT = annual average daily traffic
ADA = Americans with Disabilities Act
EV = electric vehicle

Points are each driver's maximum, rounded to one decimal. A driver marked lower is better scores (1 − x).

Fixed anchors. Each driver is scaled against a versioned benchmark, so a site's score does not depend on which other sites are in the run.

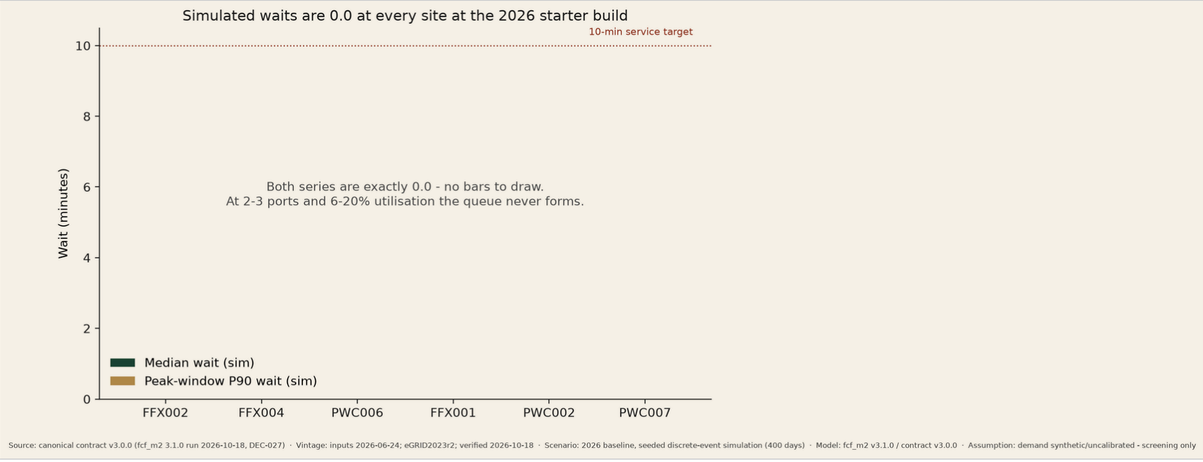
Absent is absent. A blank, n/a, unknown or not recorded cell earns no credit, in either direction.

Capacity earns nothing yet. Spare capacity is unconfirmed at all 12 sites, so its 10.5 points go unused everywhere.

2. Sizing Ports and Checking Waits

sessions per day × a growth multiplier → Erlang-C chooser the fewest ports → floors and lot cap 2 to 12 ports → peak window 55% of a day in 4 hours → seeded simulation M/G/c, 400 days → pathway ceiling never exceeded

1. Modelled demand times a provisional multiplier: 1.9 or 2.6 by 2036, 3.0 or 4.5 by 2046, 4.0 or 6.0 by 2056, base or high. The stress case is high × 1.15.
2. The smallest port count with utilisation at or below 70% and a mean wait of 10 minutes or less, searching up to 20 ports.
3. At least 2 ports, at most 5% of parking spaces rounded up to an even number. A lot nobody counted gets 2, not 6.
4. Erlang-C again, at the arrival rate of the busiest four hours, when 55% of the day's sessions arrive.
5. A discrete-event simulation with a fixed seed, 20260719, reports the 10th, 50th and 90th percentile waits.
6. Active ports stay within each site's pathway at every horizon. Contract rule RT-03 fails any row that breaks this.



Repository chart dashboard/static/07_peak_wait.png, generated from contract v3.

An honest zero, and its limit. At the 2026 starter build, 2 or 3 ports a site at 6 to 20% utilisation, no queue forms: the simulated median and 90th-percentile waits are 0.0 at all 12 sites.

The zero does not carry forward. The contract's queueing estimates climb to about 69 minutes on average (PWC008, 2056) and 215 at peak (PWC002, 2046), and five PWC002 rows have no finite peak wait: that queue never clears.

i

M/G/c = random arrivals, general dwell times, c ports
PWC = Prince William County site
RT = reconciliation test

3. Cost, Emissions and Expansion Triggers

Lifecycle Cost and Emissions

Net present cost at a 3% real discount rate with 2.5% base escalation: the staged plan against a broad buildout now.

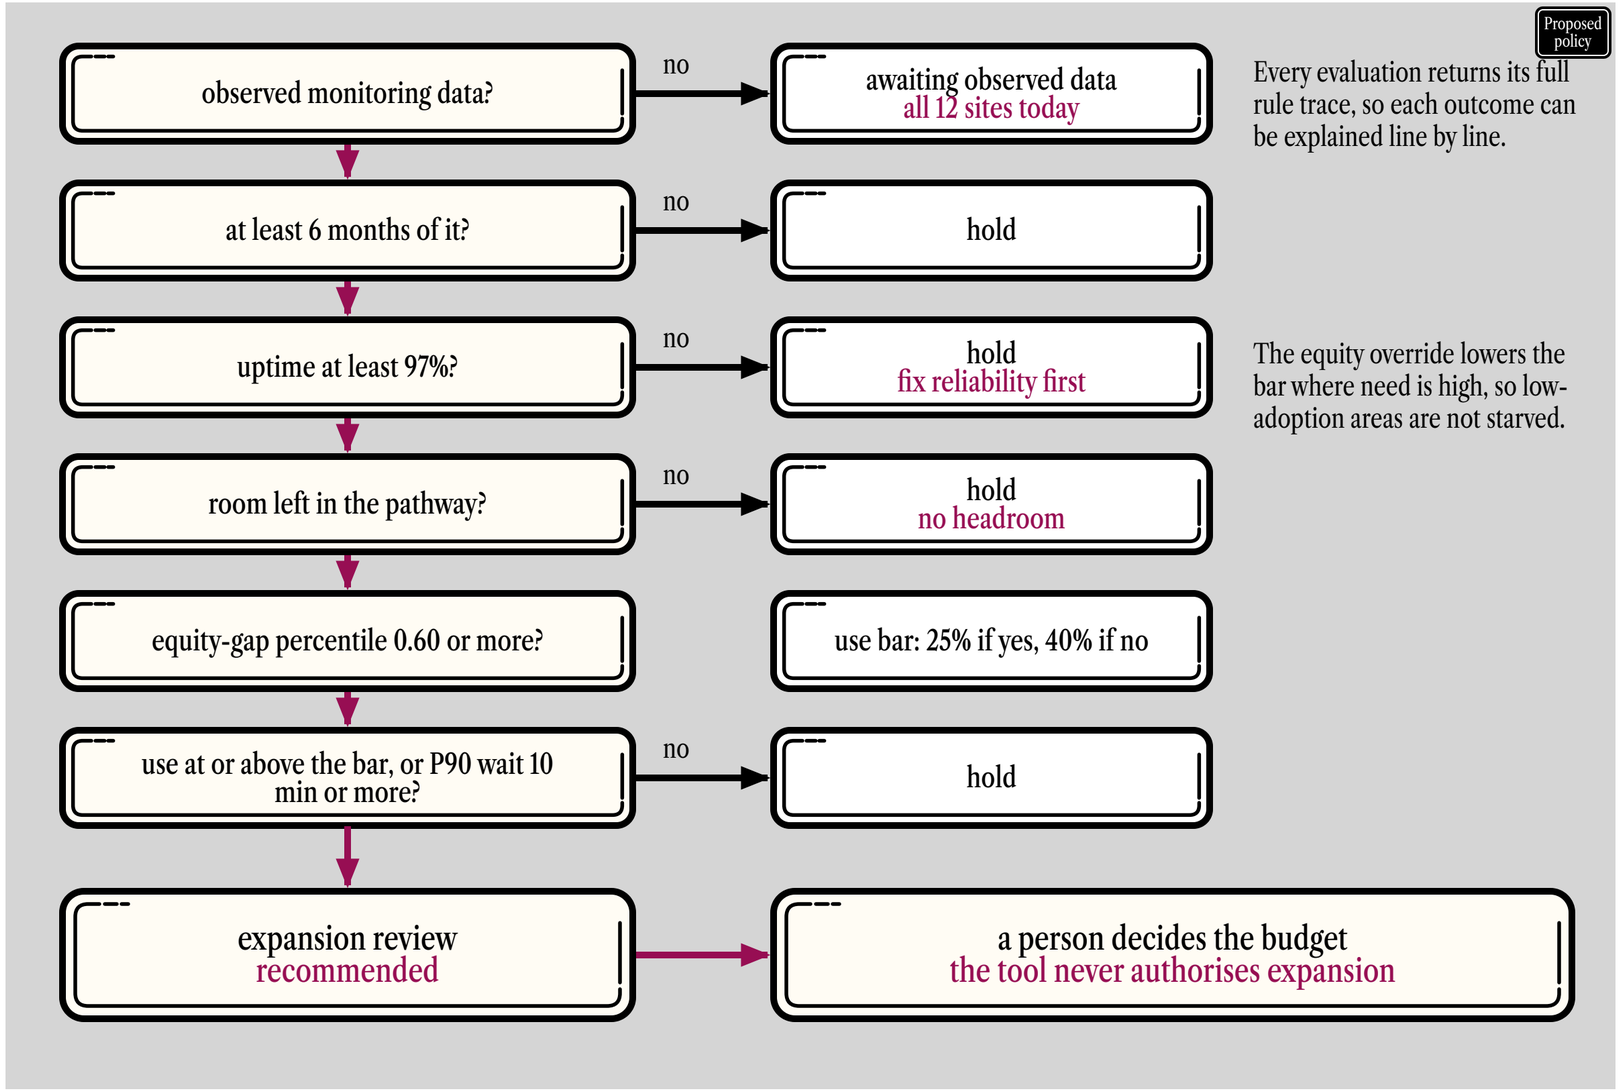
Screening-class benchmarks for comparing cases. Never budgets, quotes or bids.

Grid factor: EPA eGRID2023 Rev 2, SRVC subregion, 596.3 lb CO2e per MWh, re-pulled every year.

Displacement: 0.30 kWh per mile in place of 400 g CO2 per gasoline mile, both labelled assumptions.

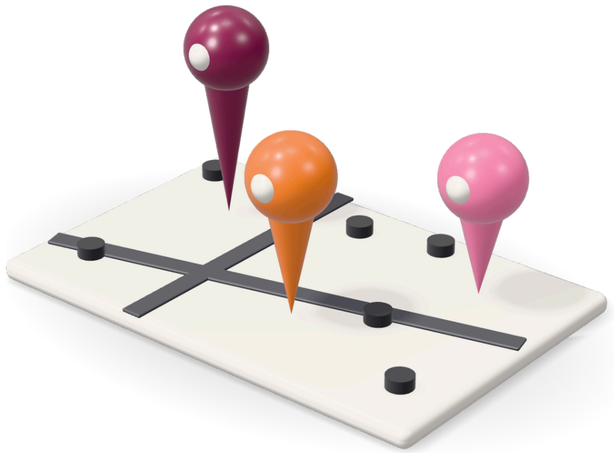
Negative results are reported as computed, never clipped.

EPA = Environmental Protection Agency
eGRID = Emissions & Generation Resource Integrated Database
SRVC = the Virginia–Carolina eGRID subregion
CO2e = carbon dioxide equivalent
MWh = megawatt-hour
kWh = kilowatt-hour
P90 = 90th-percentile wait



: Proposed trigger policy, not adopted : Yes: continue : No: stop there

D. Results



Under the canonical model, Fairfax County Government Center (FFX004) ranks first at 53.4, and the City of Fairfax Regional Library (FFX002) second at 51.8. Two of the 12 sites reach Tier 1.

The plan: 25 ports now inside a 47-port pathway. Data confidence is Low at all 12 sites.

These are screening values from contract v3: comparisons, not approvals.

1. The Twelve-Site Ranking, on the Study Site



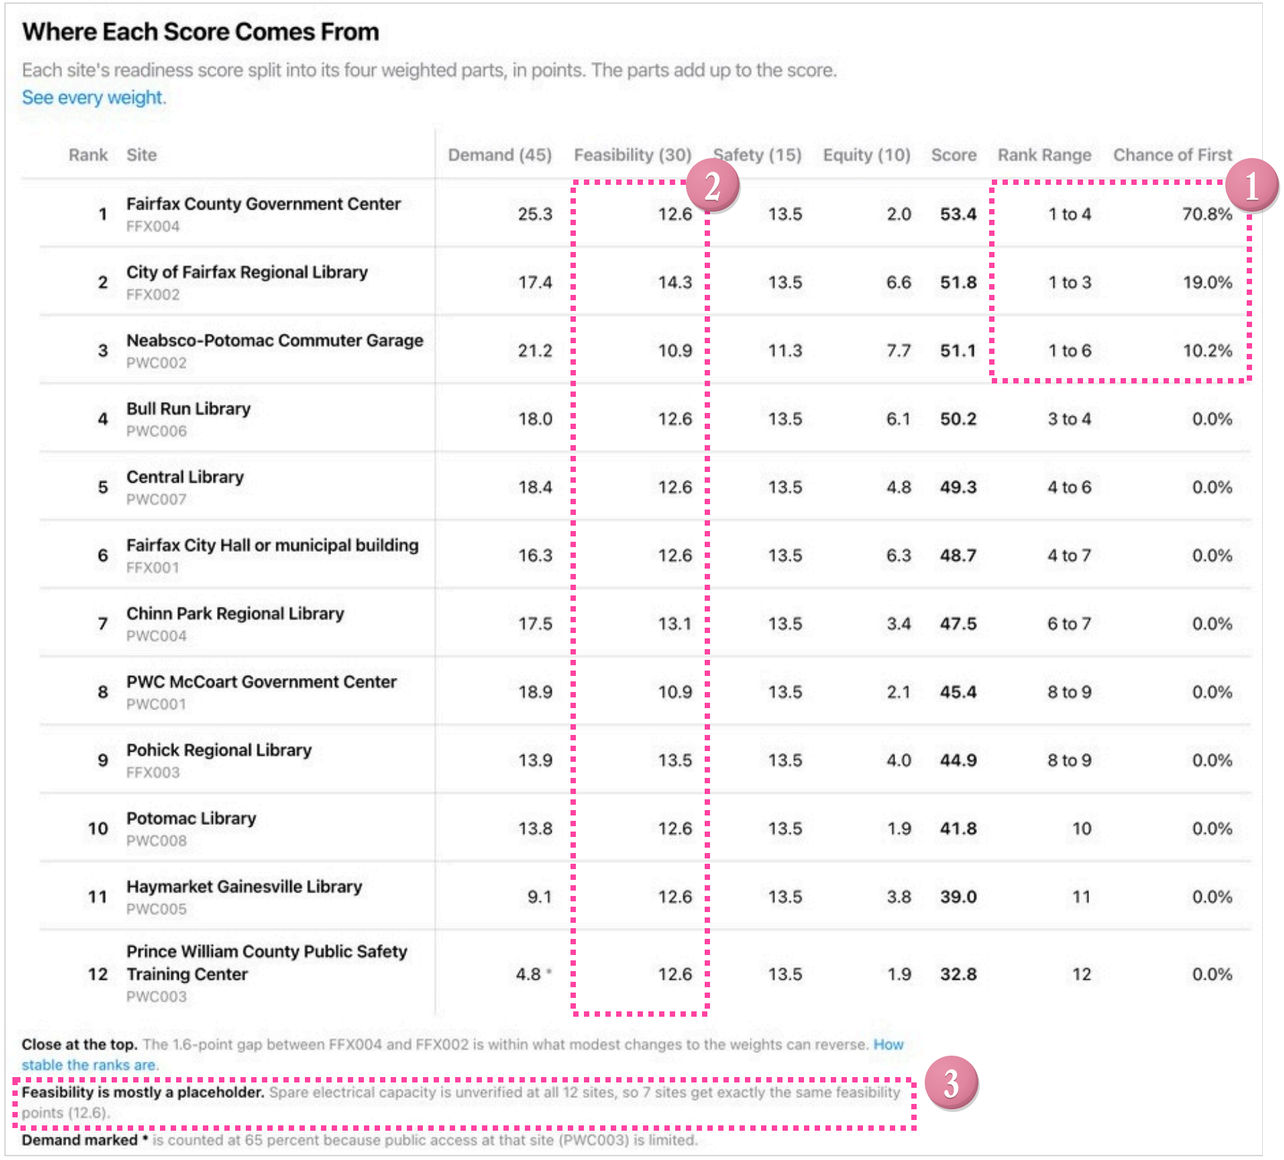
1. FFX004 leads FFX002 by **1.6 points**, 53.4 to 51.8. Both are Tier 1: advance to site screening.

2. The page itself warns that modest changes to the weights can reverse that gap.

Tiers, by the chart's own legend: Tier 1 advance, high-review candidate, Tier 2 promising, Tier 3 monitor, special use do not advance. Neabsco-Potomac, third, is the high-review garage.

Screenshot from the repository's README assets (docs/assets/readme), cropped to the chart.

2. Where Each Score Comes From



Where Each Score Comes From

Each site's readiness score split into its four weighted parts, in points. The parts add up to the score.
See every weight.

Rank	Site	Demand (45)	Feasibility (30)	Safety (15)	Equity (10)	Score	Rank Range	Chance of First
1	Fairfax County Government Center FFX004	25.3	12.6	13.5	2.0	53.4	1 to 4	70.8%
2	City of Fairfax Regional Library FFX002	17.4	14.3	13.5	6.6	51.8	1 to 3	19.0%
3	Neabsco-Potomac Commuter Garage PWC002	21.2	10.9	11.3	7.7	51.1	1 to 6	10.2%
4	Bull Run Library PWC006	18.0	12.6	13.5	6.1	50.2	3 to 4	0.0%
5	Central Library PWC007	18.4	12.6	13.5	4.8	49.3	4 to 6	0.0%
6	Fairfax City Hall or municipal building FFX001	16.3	12.6	13.5	6.3	48.7	4 to 7	0.0%
7	Chinn Park Regional Library PWC004	17.5	13.1	13.5	3.4	47.5	6 to 7	0.0%
8	PWC McCoart Government Center PWC001	18.9	10.9	13.5	2.1	45.4	8 to 9	0.0%
9	Pohick Regional Library FFX003	13.9	13.5	13.5	4.0	44.9	8 to 9	0.0%
10	Potomac Library PWC008	13.8	12.6	13.5	1.9	41.8	10	0.0%
11	Haymarket Gainesville Library PWC005	9.1	12.6	13.5	3.8	39.0	11	0.0%
12	Prince William County Public Safety Training Center PWC003	4.8 *	12.6	13.5	1.9	32.8	12	0.0%

Close at the top. The 1.6-point gap between FFX004 and FFX002 is within what modest changes to the weights can reverse. How stable the ranks are.

Feasibility is mostly a placeholder. Spare electrical capacity is unverified at all 12 sites, so 7 sites get exactly the same feasibility points (12.6).

Demand marked * is counted at 65 percent because public access at that site (PWC003) is limited.

1. Chance of first place across 2,000 reweightings: **70.8, 19.0 and 10.2 percent**. Read it as a share of plausible weightings, not the chance a site is the right place to build.

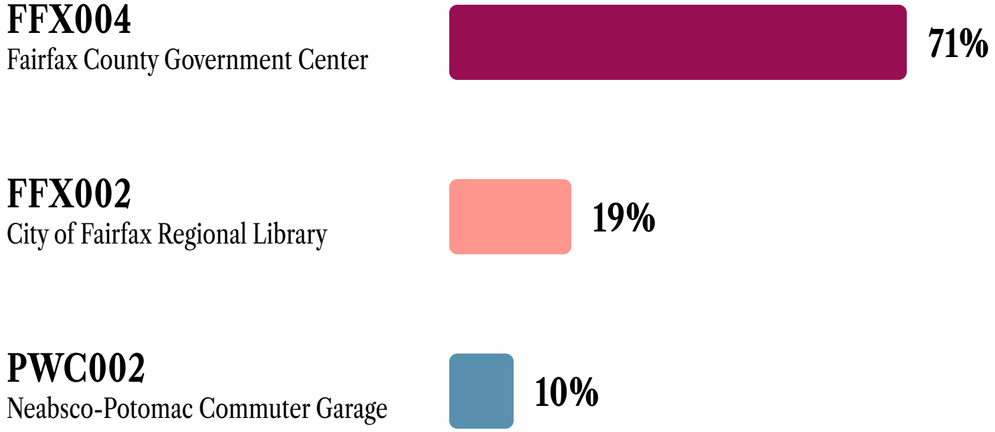
2. Seven sites get exactly the same feasibility points, **12.6**. Spare capacity is unverified everywhere, so it cannot tell them apart.

3. The page says so itself, in the note under the table.

Screenshot from the repository's README assets, cropped. Demand marked * counts at 65 percent at PWC003, where public access is limited. PWC means a Prince William County site; FFX, Fairfax.

3. How Stable Is First Place?

Chance of first place, across 2,000 reweightings



Weights are drawn from a Dirichlet distribution centred on 45, 30, 15 and 10, with a concentration of 50 and a fixed seed, 20261018. Every draw sums to 100 points. Only these three sites ever come first, and FFX002 is in the top three in every draw.

Read "chance" as a share of plausible weightings, not the chance that a site is the right place to build.

Sensitivity

53.4 against 51.8
a 1.6-point gap at the top

Equity from 10 to about 13 points
the smallest single change that puts FFX002 first

Five fewer points of Demand
FFX004 stays first, 53.1 to 53.0

Capacity alone, from none to plenty
9 of the 12 sites could rank first

Every unconfirmed input free to move
any of the 12 could rank first

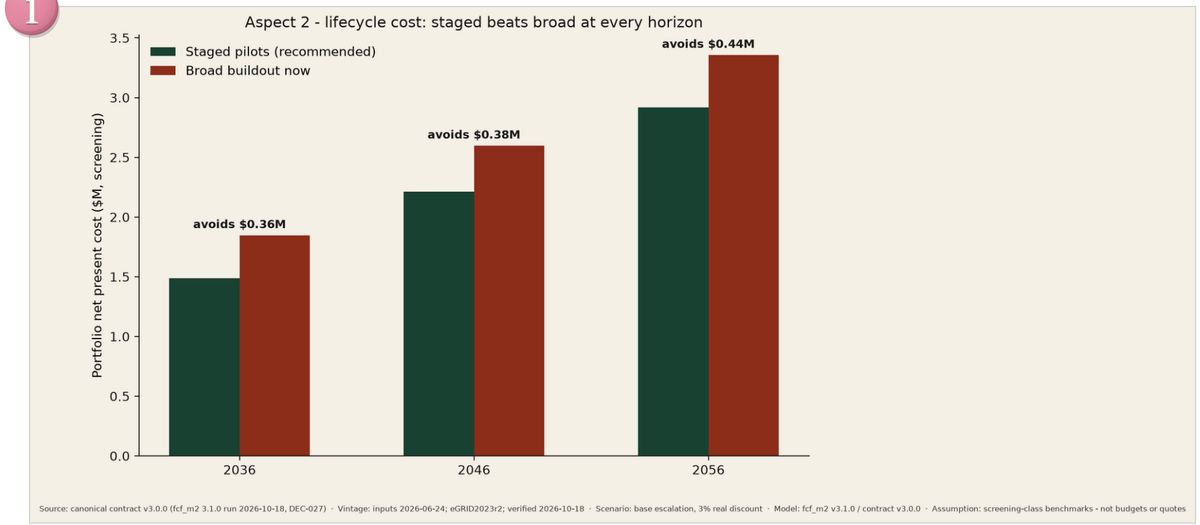
Seven sites share one feasibility subscore
41.92, because capacity earns no credit anywhere

: Findings of the rank-stability check, judged 18 October 2026

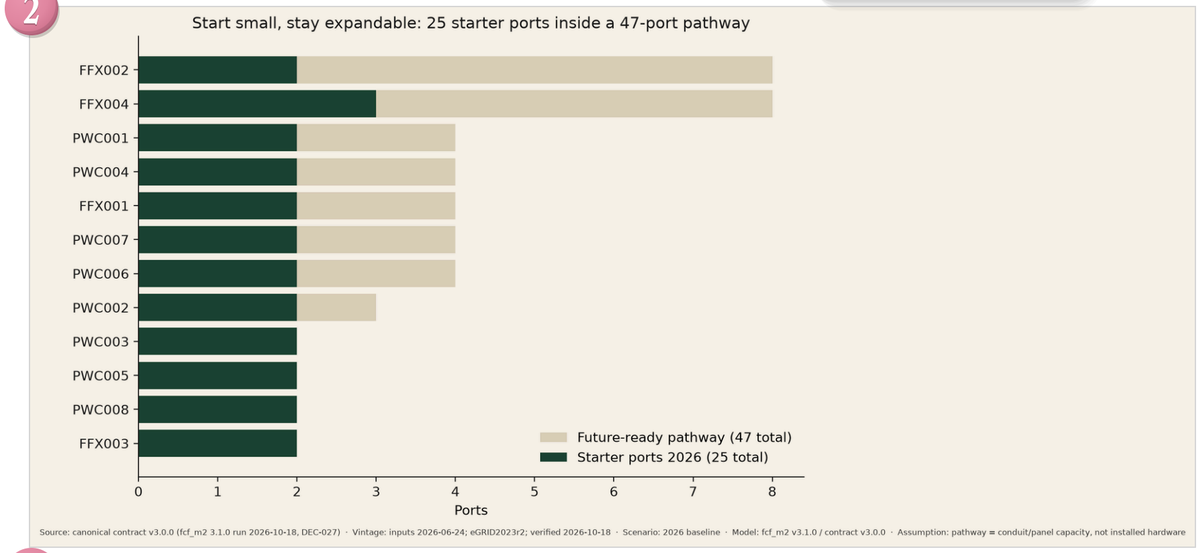
4. Cost, Ports and Emissions, in the Repository's Charts

M = million US dollars

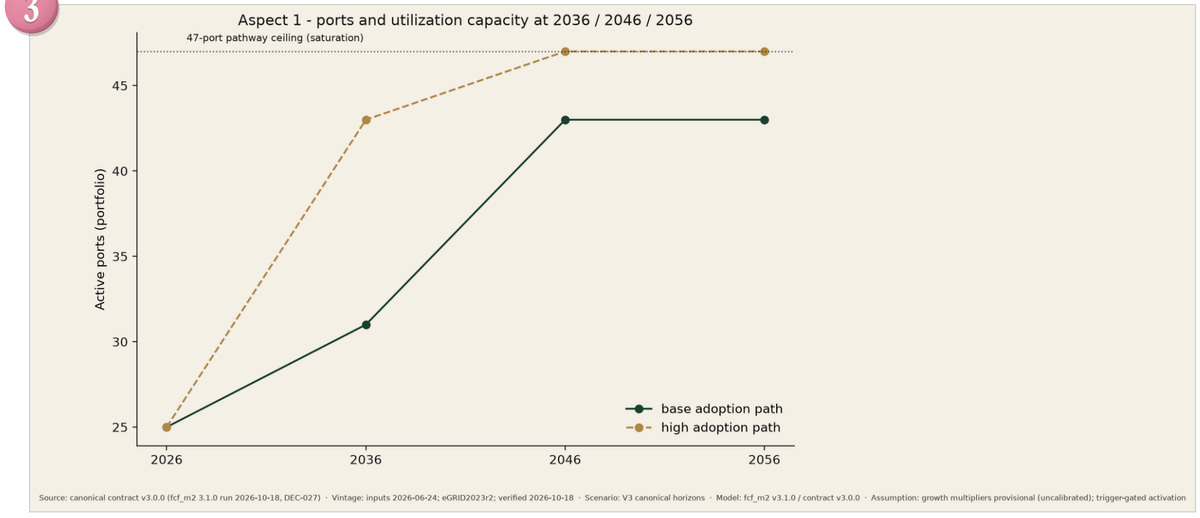
1



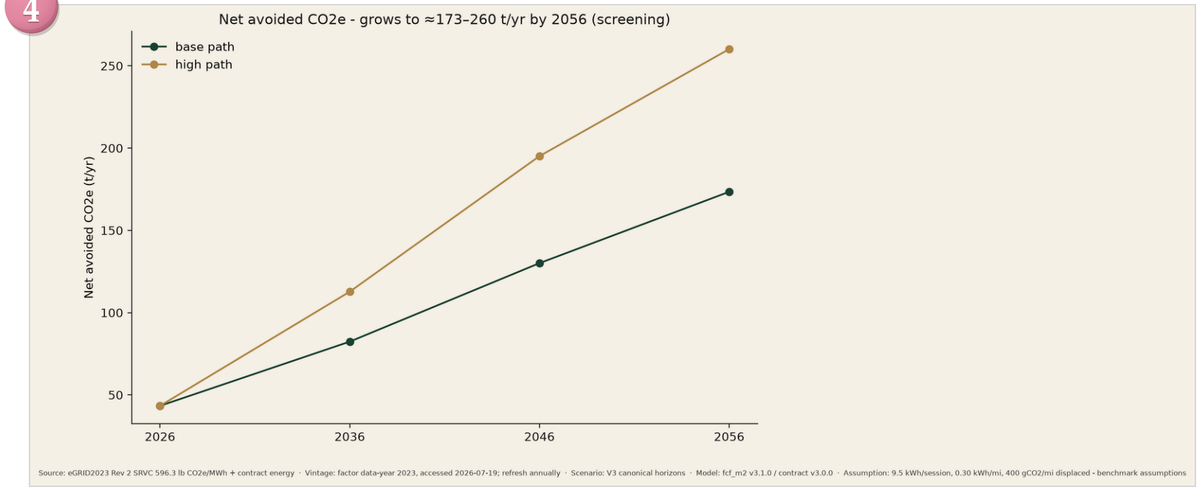
2



3



4



1. Staged pilots avoid $0.36M by 2036, $0.38M by 2046 and $0.44M by 2056. 2. 25 starter ports inside a 47-port pathway; FFX002 and FFX004 carry 8 each. 3. Active ports reach the 47-port ceiling on the high path by 2046. 4. Net avoided carbon dioxide equivalent grows to about 173 to 260 tonnes a year by 2056.

Charts from the repository's dashboard/static folder, generated from contract v3, in its own colours; each footer names source, vintage, scenario and assumptions. Costs are net present cost at a 3% real discount rate, screening-class.

E. Honest Data



A cleaner that fills gaps makes a file look better than it is, and nothing downstream can tell.

So EV-SIENT does the opposite: absent stays absent, a guess is never a match, and every confidence score comes only from typed provenance, never from prose.

Today all 12 sites score 33.8 for data confidence, which is Low. None reaches Medium.

1. Two Ways a Cleaner Lies

ID = identifier

a third-party file
any columns, any units
→ intake proposes
matches and findings
→ a person reviews
every finding, first
→ apply
accepted proposals only
→ the model runs
on what was accepted

1. It Fills Gaps

Put a regional average where a traffic count is missing and the run succeeds, but nothing downstream can tell that score from one built on real data.

So an absent field stays absent: it is scored as an absence, lowers confidence, and reaches provenance as missing.

2. It Guesses

A column matched below the confidence floor is reported as unmapped, not mapped hopefully.

An absence is scored and visible. A wrong mapping is a wrong number with no trace.

projected coordinates
refused, not reprojected

duplicate or missing site IDs
refused

Reprojecting by hand would move every site by an amount nobody could check. Two places scored as one site is a wrong answer, not a degraded one.

Measured, Then Fixed

The first version let accepting a mapping upgrade the source as well: a messy file scored 58.5, Medium, above the real study's 33.8, Low.

Now accepting clears only the mapping doubt. Ingested 0.0, accepted 18.0, real study 33.8, and a test fails if that order ever flips.

: Intake, in order

2. Missing Data Must Never Score Better

Absence earns no credit, in either direction

Where	Before	After
Absent yes or no flag	read as no; the inverted enclosed-garage and stormwater flags earned their maximum	**no credit, either way**
Absent install cost, service distance, trenching or existing ports	0, so (1 − 0) gave full marks	**no credit**
Flag outside 0 and 1, such as 7 or “0.5”	7 counted seven times its weight; “0.5” and 0.5 scored differently	**refused, naming the field**
Lot cap with no parking count	6 ports, three times a measured 40-space lot	**2, the minimum**
Scaling across sites with a value absent	entered the range as 0 and moved every other site	**range over recorded values only**
Enclosed garage not recorded	treated as not enclosed	**high-review candidate tier**

Found in a review, reproduced by running the code before anything changed, and resolved on 17 October 2026 under decision record DEC-011. A test file now pins the fix. None of it changes a published number: all 12 real sites have every flag and field filled in.

Typed Data Confidence

Source class sets the base: verified utility or facility 100, official source 90, public web 60, derived public data 55, planning assumption 35, synthetic 30, unconfirmed 15, missing 0.

Review adjusts it: reviewed +5, unreviewed −5, a source older than two years −10.

Field groups are weighted, electrical service heaviest at 25%.

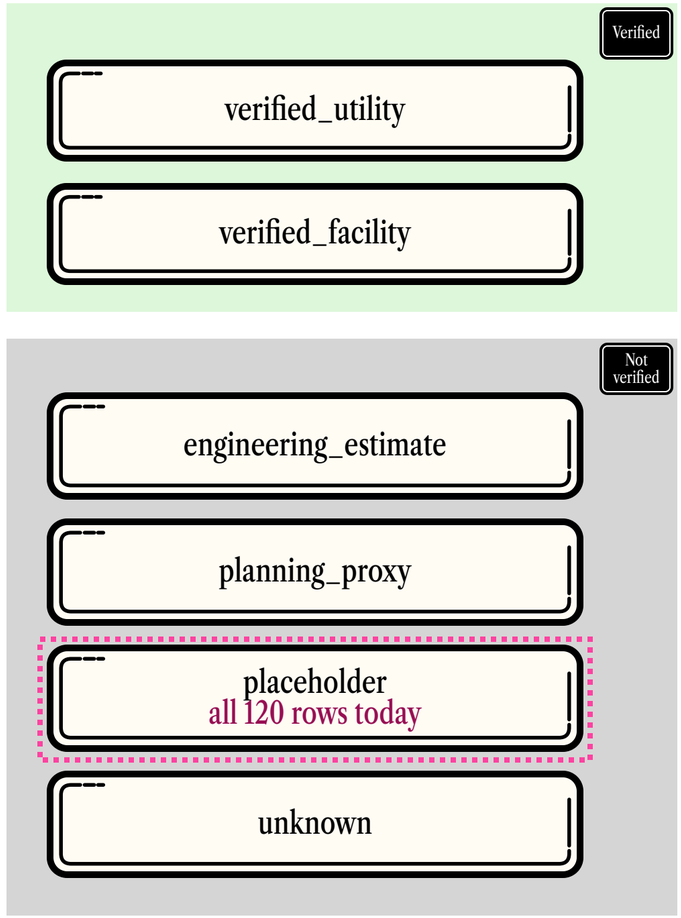
High at 75 or more, Medium at 50 or more, otherwise Low.

Editing prose in a notes column cannot change a score.

33.8 at all 12 sites
Low; none reaches Medium

3. Capacity Is a Placeholder, and the Code Refuses to Pretend

RT = reconciliation test



capacity_basis: a controlled vocabulary of six values. Anything else fails validation.

scoring
0 of 10.5 points

Spare capacity earns no credit at any site, so seven sites share one feasibility subscore, 41.92. That 35% of feasibility does nothing yet.

contract rule RT-07
fails validation

A placeholder row that claims high capacity confidence, or lacks its placeholder flag, fails. capacity_basis rides on all 120 rows.

MATLAB guard
fails closed

Refuses to emit a feasibility claim unless the basis is verified_utility or verified_facility. Confirmed firing on licensed MATLAB R2025b.

Octave proof
a verify.sh gate

Exercises the same guard against every non-verified basis where Octave is installed. It was not installed for this deck's run.

The most important unknown: nobody has confirmed that any of the 12 sites has the electrical capacity for chargers. Confirming it is the highest-value single data request in the project.

: Feasibility claims allowed : Feasibility claims refused

F. The Planner

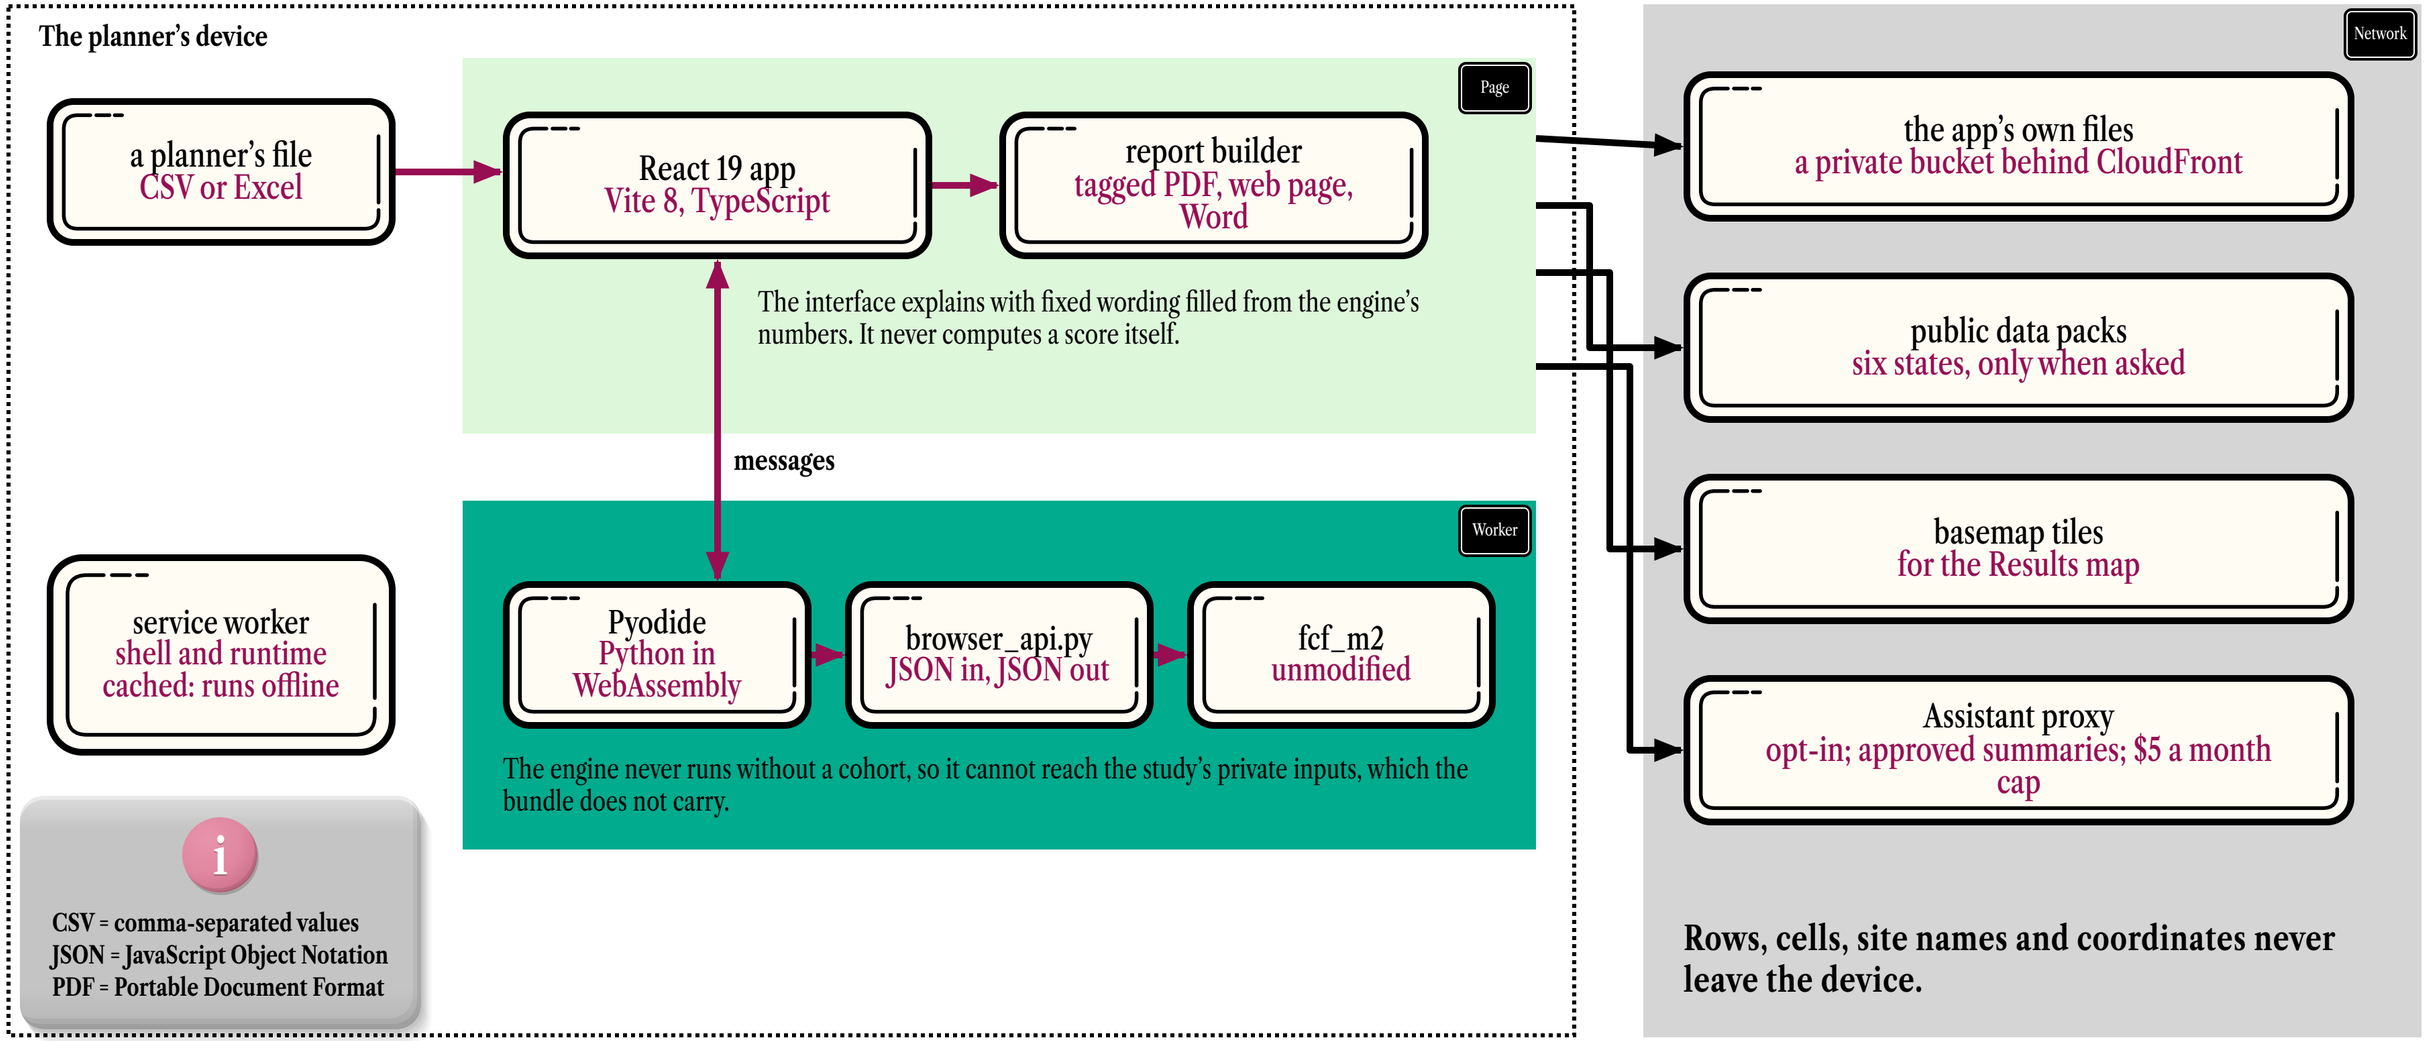


The planner at app.futurecity.andyyaro.com lets anyone rank their own candidate sites with the same model, with nothing to install.

The unmodified Python package runs in the browser, under Pyodide in a Web Worker. The planner's file never leaves the device.

The flow: Start, Upload, Map Columns, Check, Public Data, Run, Results, Explore, Your Index, Report.

1. The Same Model, in the Planner's Browser

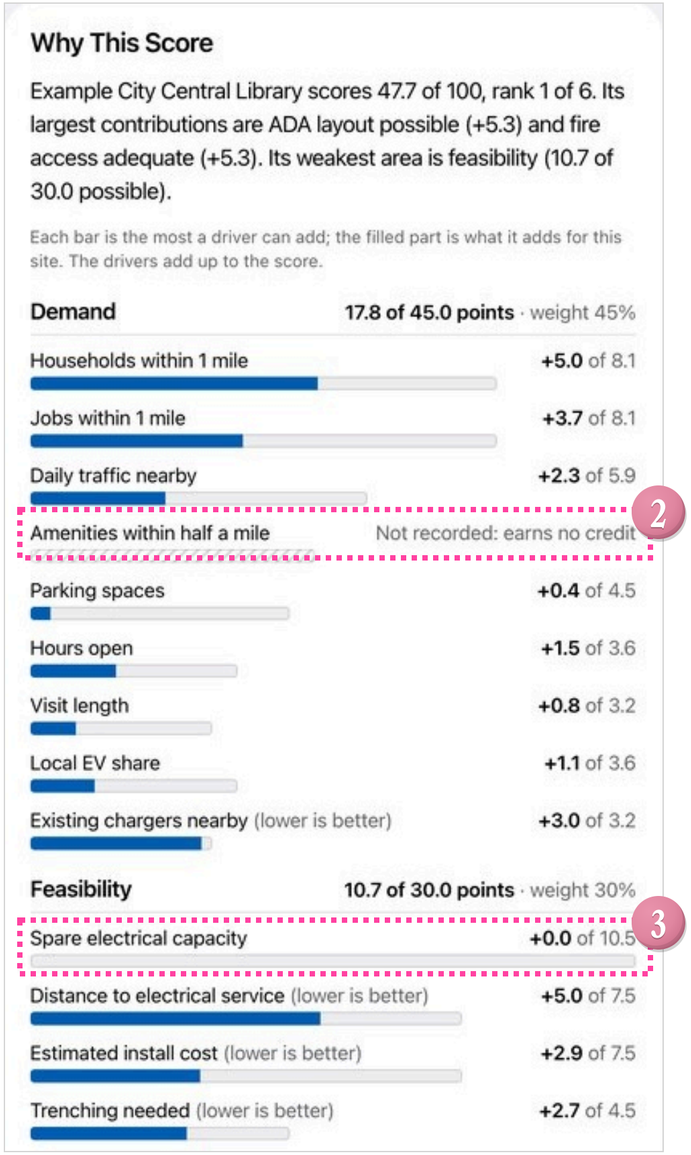


: The page : The model's worker : The network : On the device : Requests out

2. The Planner on Its Invented Sample

Rank	Site	Score	Tier	Data Confidence	Ports Now	Pathway	Compare
1	**Example City Central Library** EXC-01	47.7	Tier 2	Low	2	4	☐
2	**Example City North Branch Library** EXC-02	43.4	Tier 3	Low	2	2	☐
3	**Example City Transit Garage** EXC-05	42.5	Special Use	Low	2	2	☐
4	**Example City Hall Annex** EXC-03	42.2	Tier 3	Low	2	3	☐
5	**Example City Recreation Center** EXC-04	41.1	Tier 3	Low	2	2	☐
6	**Example City Riverside Park** EXC-06	26.6	Tier 3	Low	2	2	☐

Example City is an invented six-site sample that ships with the planner. Screenshot from the repository's README assets, cropped to the ranked table and the Why This Score panel.



1. Data confidence reads **Low** for every site, in the table itself, not in a footnote.
2. Amenities were not recorded, so that driver **earns no credit**: the gap is shown, not filled.
3. Spare capacity adds **+0.0 of 10.5**: with no verified capacity, the largest feasibility driver does nothing.

2. Checks Before Every Release

parity in CI	The browser engine and the native Python package give the same numbers.
layout matrix 11 sizes; 3 in CI	Every screen at 11 window sizes, light and dark: no sideways scroll, nothing covered, controls large enough, no console errors, no requests to other sites.
accessibility in CI	axe-core on every screen in all four themes, and the whole flow from Start to a downloaded report by keyboard alone.
real engines Chromium and WebKit	The real engine boots and runs the Example City sample in both browser engines.
offline in CI	The app installs; after one run the server is stopped and the whole flow runs again with the real engine.
security headers production policy	CI serves the build under the production Content Security Policy before testing it. One CloudFront Function is the only place that policy is defined.

Provenance, Made Linear

Opening any number's lineage once used a per-cell mask: exporting it for 600 sites under Pyodide took 466 seconds.

An index built once brought it to 9.0 seconds, with byte-identical output pinned by a test.

Accessibility, Honestly

A conformance report against WCAG 2.2, levels A and AA, self-assessed by the team.

No outside specialist reviewed it, no screen reader was used, and no person with a disability tested it.

i

CI = continuous integration
WCAG = Web Content Accessibility Guidelines

G. Beyond Python

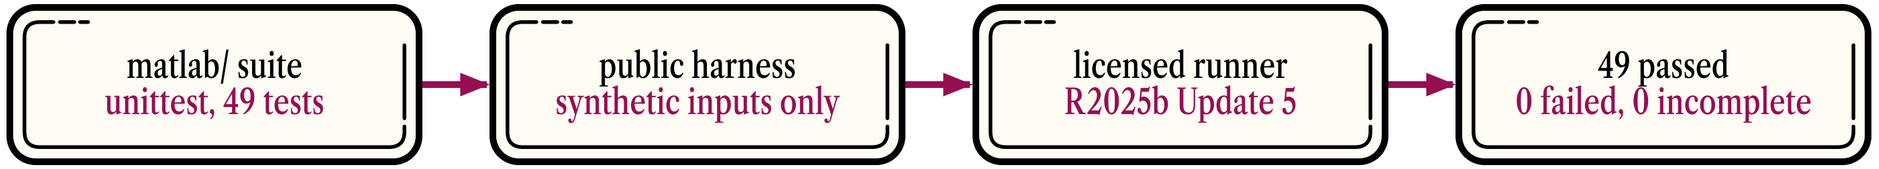


The model also runs as MATLAB code with a Simulink site-load model, executed on licensed MATLAB against a synthetic fixture.

A field layer lets someone standing at a site correct the study: a tapped near-field communication tag opens a form that produces a spreadsheet row for a person to review.

Both layers state plainly what they have not done.

1. MATLAB and Simulink: Executed on Synthetic Data



CSV = comma-separated values

What the Runs Establish

The suite runs on real MathWorks MATLAB, not only in review.

The fail-closed capacity guard fires there: no feasibility claim from placeholder capacity.

The Simulink model builds, and its 11 root blocks match the topology read from the builder by hand.

What They Do Not

Touch project data. The input is a synthetic fixture, so no output describes a real facility.

Simulate anything. The Simulink model has been built, never run, and no number derives from it.

Found by Running

The Simulink builder tested exist("new_system", "file"), which never reports built-ins, so the Simulink layer had been unreachable on every installation.

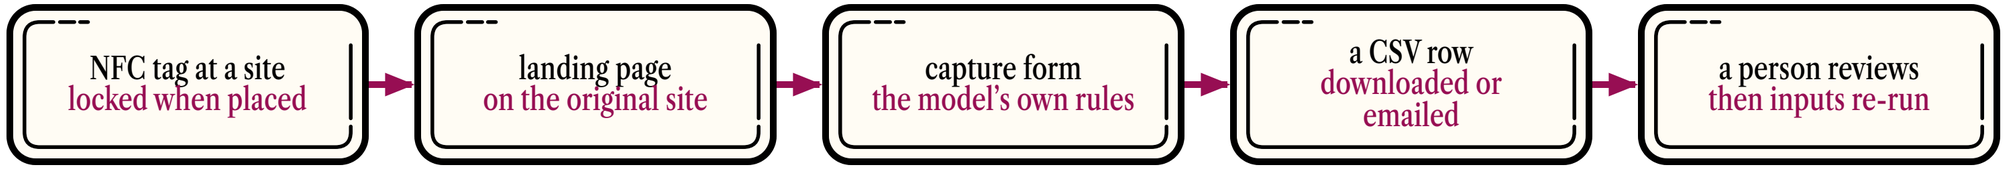
A banner saying no line had ever run stood from 19 to 27 July 2026. It is kept, marked superseded, not deleted.

A verify.sh gate checks that nothing bound for the public harness carries site-identifying data.

Run on a GitHub-hosted Ubuntu 24.04 runner through the public harness repository andyyaro/fcf-matlab-ci. Two more verify.sh gates tie the layers together: shared physical constants agree between Python and MATLAB, and MATLAB's canonical CSV reconciles with contract v3.

2. The NFC Field Layer

NFC = near-field communication
CSV = comma-separated values
ID = identifier
DEC = decision record



Decisions Behind It

Lock tags when placed, not when written (DEC-003). Locking stops a stranger rewriting a tag in public; bench tags stay re-writable.

Lock rather than password-protect (DEC-004). A password printed in a runbook protects against no one with physical access.

Never use the tag's unique ID as identity (DEC-005). It is trivially cloneable, so it authenticates nothing.

The form submits nowhere (DEC-006). No backend means no endpoint to abuse and no store to breach; a person reviews every row first.

Status

The pages and forms are live. The tags are not deployed: placing anything on a public facility needs the facility manager's consent, and that has not been sought.



A row that arrives by email is reviewed before it enters the inputs, and then the model re-runs. Nothing is ever edited in place.

: A field observation, in order

H. Proof and Limits



Every figure in this deck was checked against the repository on 2 November 2026, and the test suite, the contract validator and the verification script were re-run for it.

The last slides say what passed, what was skipped and why, and what the model cannot tell you.

1. Tests and Contracts, Re-run for This Deck

```
$ python -m pytest tests/ -q --cov=fcf_m2 --cov-report=term
.....................................sssssssssss.ssss.ssss..s.ss......ss [ 97%]
sss..s..s.....s.........s.                                              [100%]
============================== tests coverage ==============================
_________________ coverage: platform linux, python 3.13.18-final-0 _________________

Name                              Stmts   Miss  Cover
-----------------------------------------------------
[ 27 module rows left out here ]
src/fcf_m2/stability.py             260      9    97%
src/fcf_m2/standards.py              26      0   100%
src/fcf_m2/triggers.py               55      1    98%
src/fcf_m2/visuals.py                70      2    97%
src/fcf_m2/workflow.py              444    104    77%
-----------------------------------------------------
TOTAL                              6472    647    90%
1077 passed, 101 skipped in 141.87s (0:02:21)
```

```
$ python -m fcf_m2.workflow validate-contract data/contract/v3
Contract v3.0.0 valid: all rule checks, JSON Schema, and RT-10 hash chain PASS.
$ python -m fcf_m2.workflow validate-contract data/contract/v2
Contract v2.0.0 valid: all rule checks, JSON Schema, and RT-10 hash chain PASS.
$ python -m fcf_m2.workflow validate-contract data/contract/v1
Contract v1.0.0 valid: all rule checks, JSON Schema, and RT-10 hash chain PASS.

$ ./scripts/verify.sh
[ the 41 gate result lines are summarised on the next slide ]
PASS: 36   FAIL: 0   SKIP: 5
ALL GATES THAT CAN RUN HERE PASSED (5 need the study's private inputs)
```

This Run

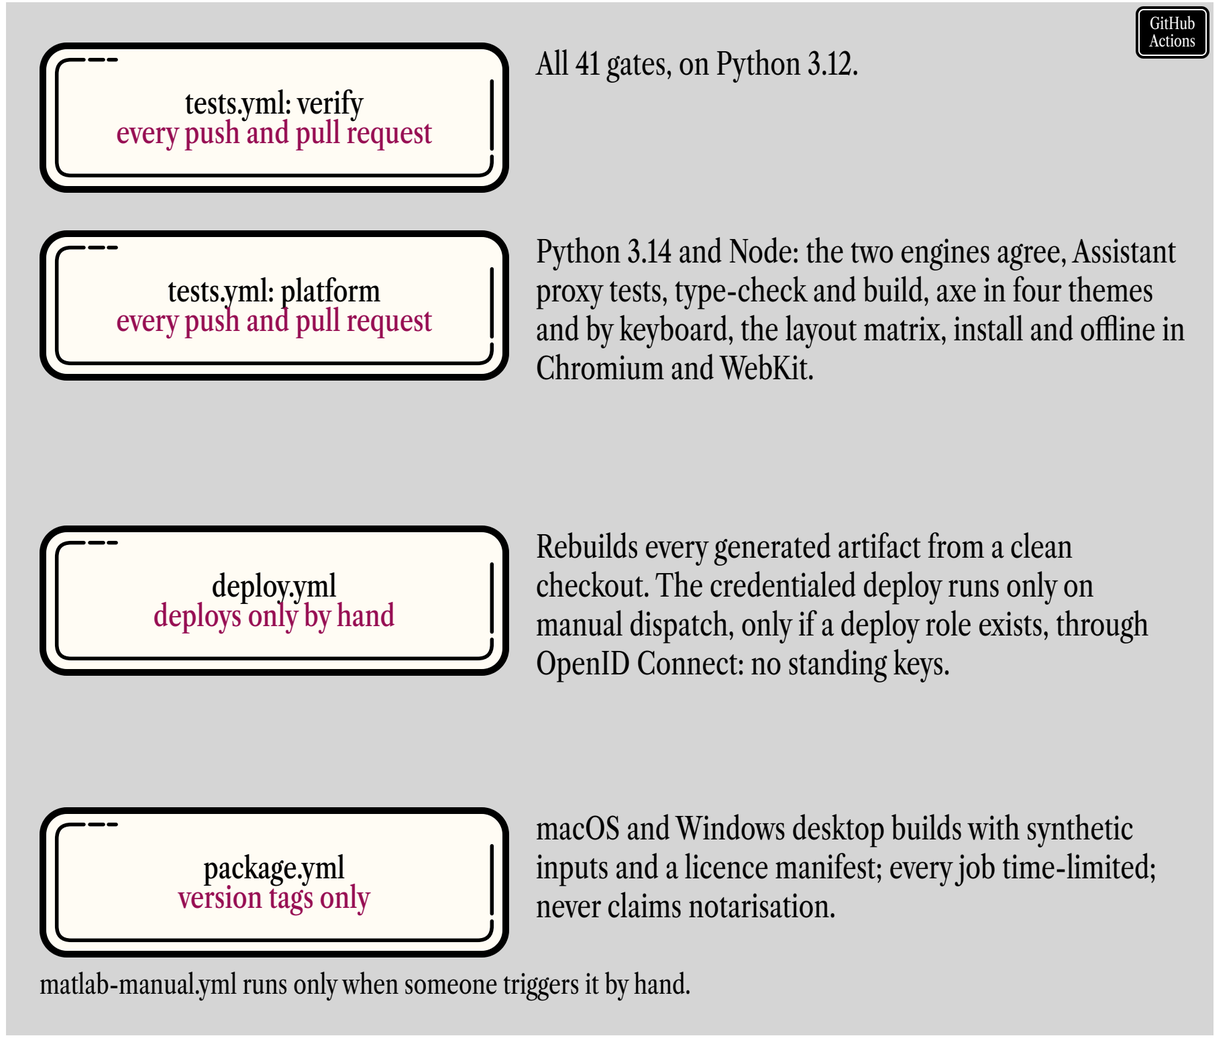
Tests passed	1,077
Tests skipped	101
Line coverage of fcf_m2	90%
Contracts valid	v3, v2 and v1
verify.sh gates passed	36 of 41
Gates failed	0
Gates skipped	5

Skips need the study's private inputs (decision record DEC-016) or tools not installed here. Run on 2 November 2026 with Python 3.13, from the uploaded copy of the repository, in its own environment; a local git commit lets the tests that list tracked files run. Coverage rows not shown are marked in the pane. One of the 36 passing gates, the Octave proof, passed without running because Octave is not installed here.

2. 41 Gates and Continuous Integration

Gate family	Gates	This run
platform	10	10 passed
contract	5	4 passed, 1 skipped
three-dimensional explorer	4	3 passed, 1 skipped
Python test suite and legacy mode	2	1 passed, 1 skipped
cross-model	2	2 passed
Tableau	2	2 passed
hygiene	2	2 passed
docs, dashboard, secrets, clone safety, MATLAB harness, field entry, supply chain, packaging, intake, browser, links	11	11 passed
Octave capacity guard	1	passed without running: no Octave here
synthetic inputs	1	skipped
website	1	skipped
All gates	**41**	**36 passed, 0 failed, 5 skipped**

A gate that needs the study's private inputs reports skip rather than failing. The browser gate ran its 44 Chromium tests: 44 passed.



: Workflows under .github/workflows

3. What It Can and Cannot Tell You

AHJ = authority having jurisdiction

It Can Tell You

How twelve specific sites compare, on modelled demand and screening-class cost.

Roughly how much overbuilding costs: the project's actual argument.

Where building the pathway now pays most.

Exactly how much is not known, and where.

Where every number came from: its run, its contract version and its source.

It Cannot Tell You

Whether a site is electrically feasible.

What anything will cost: these are not budgets, quotes or bids.

How many people will charge: demand is modelled and uncalibrated.

How long anyone will wait: waits are simulated, never observed.

Whether anyone has approved it: no client, utility or authority sign-off, and no funding.

Whether a trigger has been met: none has met real data.

Still open, from the repository's list of actions no software can close, and what the triggers and tags still wait on:

- utility and facility data
  FFX002, FFX004 and PWC006
- fire marshal and AHJ talks
  in both jurisdictions
- facility consent
  before any tag is placed
- six months of observed use
  before any trigger can run

What would change these answers: one thing, mostly. Someone confirming the electrical capacity at these sites releases the feasibility guard, puts the 35% capacity weight to work, and lifts data confidence off the floor.