Richmond Civic Hackathon: Replication Toolkit

A focused playbook for another city

Ford Prior, Crystal Harvey, Michael Kolbe, Heather Lyne, Sinclair Jenks, Ashley Ray, Mike Huddleston, Will Melton, Gray Crenshaw, Christian Markow, Vivian Vendt

Flying Pig Labs / hack4rva · City of Richmond

ford@flyingpiglabs.com

Version 1.0 · May 2026 · github.com/hack4rva

In partnership with

Virginia Commonwealth University · AI Ready RVA · PlanRVA
Lost Office Collaborative · The Community Foundation for a Greater Richmond

Abstract. Most civic hackathons concentrate effort in the weekend itself and produce roughly 5% project survival five months later. The Richmond Civic Hackathon stretches the model across three stages — preparation, event, and post-event diffusion — to raise that rate through stronger pre-event work and a more intentional implementation pathway. This toolkit documents the operating sequence, what worked, and what did not. A companion planning plugin — fill-in templates and example artifacts that operationalize each section of the toolkit for another city — is published at github.com/hack4rva/annual-civic-hackathon-toolkit. Throughout, an italic *Templates* → callout at the end of each section links to the matching template directory in that companion repo.

Event: March 27–29, 2026 · VCU Snead Hall · Richmond, Virginia
Anchor: Mayor Danny Avula's Mayoral Action Plan (seven civic pillars)

Contents

Contents i

1 Core Thesis 1
- 1.1 The Survival Benchmark . . . 2
- 1.2 Operating Architecture . . . 2
- 1.3 The Seven Pillars . . . 4

2 Pre-Event Infrastructure 8
- 2.1 Origin . . . 8
- 2.2 The Happy Hour . . . 9
- 2.3 Engaging City Hall . . . 9
- 2.4 Track Governance: The 3C Committee . . . 10
- 2.5 The Off-Site: Lost Office Collaborative . . . 10
- 2.6 Pre-Event Research Repositories . . . 10

3 The Event 13
- 3.1 Three-Day Structure . . . 13
- 3.2 Platform and Tooling . . . 15
- 3.3 Award Winners . . . 17
- 3.4 The Judging System . . . 18

4 Post-Event Evaluation 22
- 4.1 The 73-Idea Corpus . . . 22
- 4.2 Two Evaluation Layers . . . 23
- 4.3 Implementation Reality Rubric . . . 23
- 4.4 AI-Assisted Review . . . 25
- 4.5 The Public Marketplace . . . 26

5 Continuation Engine 28
- 5.1 Cadence over bursts . . . 28
- 5.2 Day 30: The Phase 2 Launch . . . 28
- 5.3 The Two-Path Post-Event Funnel . . . 30
- 5.4 Core roles and responsibilities . . . 31

6 Replication Blueprint 33
- 6.1 The Four Organizing Tenets . . . 33
- 6.2 The Replication Operating Sequence . . . 34
- 6.3 Communications Infrastructure . . . 35

Closing **37**

A Source Research **38**

Chapter 1

Core Thesis

Design goal: raise P(implementation) — the probability that a participant's weekend work reaches real-world use in city government.

Prior art and origin

A **civic hackathon** is a short, structured prototyping event where teams build software against problems sourced from a city government or public-sector partner. The Richmond Civic Hackathon space itself has thin prior art — a handful of one-off events, no sustained program — but the broader genre is well established. The most-cited benchmark, Washington D.C.'s *Apps for Democracy* (2008), produced 47 applications in 30 days and is referenced again in the survival discussion below.

Two preconditions let organizers skip re-pitching the concept:

- **Hackathons are a familiar format in Richmond.** Local developers, designers, students, and civic-tech volunteers have attended hackathons before. The term does not need to be defined for participants, sponsors, or venues.
- **A mayor's office liaison was already in place.** Michael Kolbe, embedded in the Mayor's Office of Strategic Initiatives, understood the format and translated it inside City Hall, which removed the friction a first-of-its-kind event usually creates with city government.

The rest of this toolkit assumes a similar starting point: a city where the format is recognizable and where at least one person inside city government is willing to vouch for the work. If neither of those conditions exists, an organizer's first job is to create them; the weekend is the second job.

The three-stage model

The event itself is rarely the hard part. Most hackathons concentrate effort on the middle stage — a compressed weekend of team formation, prototyping, and demos. That stage matters but is not enough on its own; civic work depends heavily on what happens before and after.

The Richmond model stretches across three stages — **Preparation**, **Hackathon**, **Diffusion** — to raise the number of prototypes that survive the weekend through stronger preparation and a more intentional implementation pathway.

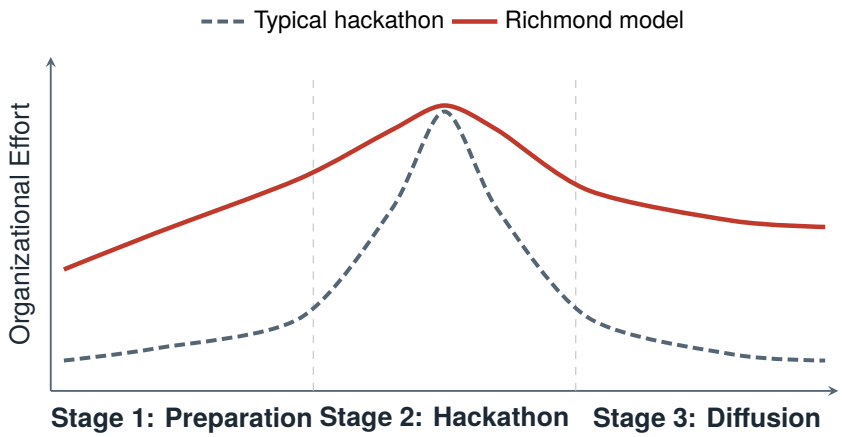


Figure 1.1: Typical hackathons concentrate effort in the weekend. The Richmond model invests across all three stages.

1.1 The Survival Benchmark

The standard continuation benchmark for civic hackathons is roughly 5% of projects still active five months after the event. Washington D.C.'s *Apps for Democracy* produced 47 applications worth an estimated $2.3M in development value; the District adopted none of them.

Applied to a 47-project cohort:

$$47 \times 0.05 = 2.35 \quad \text{(baseline)} \qquad 6 \div 47 \approx 12.8\% \quad \text{(double baseline)}$$

Set the first-event bar accordingly: three surviving projects to meet the benchmark, six to double it. The hackathon should not be judged only by the number of people in the room or the polish of the demos — those are Stage 2 metrics. The test is whether preparation makes the weekend useful, and whether post-event diffusion helps the strongest ideas survive.

1.2 Operating Architecture

Treat the hackathon as a pipeline, not an event. Three operations determine whether anything survives the weekend: **define the inputs**, **structure the transformation**, and **funnel the outputs**. Everything else is implementation detail.

Define the inputs

The hackathon is anchored to an existing public strategy (in Richmond, the Mayoral Action Plan), not to open-ended ideation. Teams are not asked to invent civic problems. They are asked to work within a shared civic framework with real public-sector priorities. Problem statements are written by steering committees before the event; participants cannot substitute their own. That constraint keeps demos aligned to civic work instead of drifting into generic product ideas. Alongside the problem statements, the inputs include vetted data sources, named civic partners, and a city ally per pillar — the material a team needs to start building on day one.

Structure the transformation

The weekend works only because the organizers invest heavily before it. Pre-work includes partner coordination, problem framing, research, sponsor outreach, mentor recruitment, participant communications, venue planning, and pillar-specific repositories. Pre-work exists so teams spend the weekend building, not figuring out what problem they are solving. Real data is required; synthetic or fabricated inputs are deployment blockers, not shortcuts. The transformation step turns problem statements + data + people into demos in 48 hours.

Funnel the outputs

The hackathon is not the finish line. The hardest work happens after the weekend: reviewing ideas, matching projects to owners, identifying implementation paths, and helping prototypes survive contact with real institutions. The event has to produce more than demos: artifacts, relationships, and enough structure for a second phase of work.

Three sub-lessons sit underneath this principle:

- **Solutions need owners.** The post-event pathway breaks down without a named person — inside a City department, nonprofit, or partner organization — willing to take a prototype forward.
- **Desirability is not survivability.** An idea can be good and still fail every test of whether it can ship. The post-event evaluation makes that distinction visible without punishing good ideas.
- **Narrow beats broad.** The ten most shippable ideas in the Richmond cohort averaged 2.3 data sources each; the ten least shippable averaged 4.7. Every additional integration point is a potential blocker.

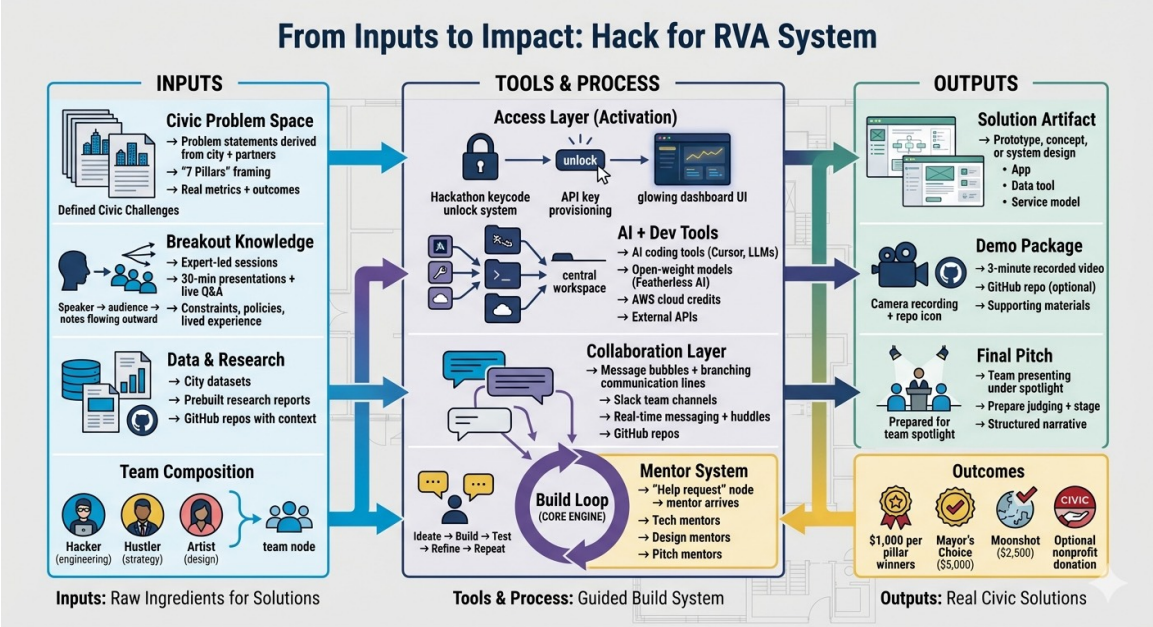


Figure 1.2: The architecture in motion. Each of the three stages produces concrete artifacts that become inputs to the next. Stage 1 defines the inputs — problem statements, pillar research, partner committees. Stage 2 structures the transformation — demos, team relationships, live award winners. Stage 3 funnels the outputs — an evaluated idea corpus, a City Champion Path, and civic working sessions for ideas that need a different home. Only the last stage produces prototypes that survive contact with real institutions.

Templates → `templates/01-civic-alignment/`

1.3 The Seven Pillars

The event anchors to seven civic pillars. Each pillar published two **practical** problem statements (real, scoped, data-backed) and two **blue-sky** problem statements (broader, exploratory framings written by the steering committee but not assigned to a specific data source). That is 14 + 14 = 28 problem statements total. Teams choose one; they cannot substitute their own.

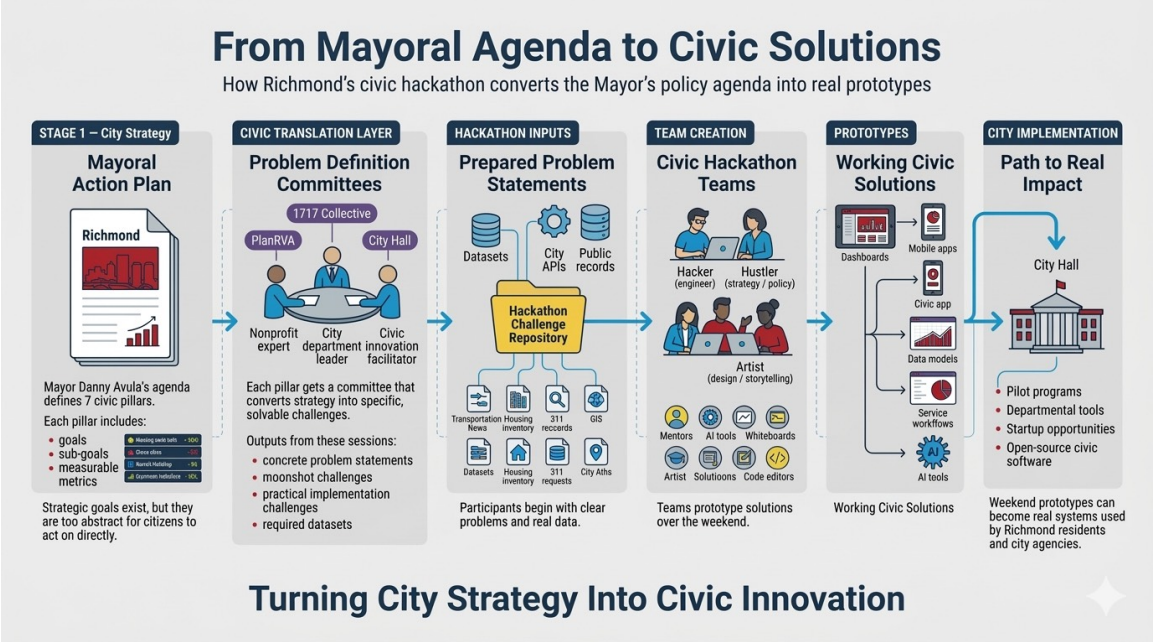


Figure 1.3: From the Mayoral Action Plan to civic solutions. Public strategy → seven pillars → 28 problem statements (14 practical + 14 blue-sky) → weekend projects + AI-expanded ideas → post-event implementation candidates. The pipeline is what keeps demos aligned to priorities the City already published.

For each practical problem statement, the steering committee identifies the data the weekend needs, the civic partners willing to provide it, and the city ally who can unblock access or sign off on post-event work. The named partners below were engaged during pre-event recruiting — not aspirational contacts.

Pillar	Practical Problem Statement	Data Needed	Civic Partner / City Ally
A Thriving City Hall	Resident Service Navigation	rva.gov content map; RVA311 / AvePoint service catalog	Mayor's Office (Michael Kolbe) / Pete Briel (311 Director)
A Thriving City Hall	Procurement Risk & Opportunity Review	City Contracts (`xqn7-jvv2`); OnBase / RAPIDS contract PDFs; state & federal procurement feeds	UNOS (Ankit Mathur) / Doug Gernat (Sr. Dep. Dir. Strategy)
Thriving Neighborhoods	Neighborhood Development Discovery	Legistar OData; ArcGIS Dev Tracker; GeoHub parcels and zoning	PlanRVA (Martha Shickle) / Brian Mercer (Planning & Development Review)
Thriving Neighborhoods	Affordable Housing Compliance Monitoring	HUD covenant records; HCD program files; RRHA unit data	PHA + HOME (Annika Schunn) / Sharon Ebert (DCAO Community Development)
Thriving Families	Youth Employment Pathways	Employer / internship listings (mostly missing); RPS career services data	Girls For A Change (Angela Patton) / Eva Colen + Jacqueline Hale (NCS)
Thriving Families	Maternal Health Data Coordination	VDH, RCHD, Census, and CDC indicators; agency reporting templates	Richmond Ed Fund (Taikein Cooper) / Amy Popovich (DCAO Human Services)
A Thriving Economy	MBE Contracting Discovery	B2GNow vendor directory; OMBD certifications; City Contracts dataset	OMBD / People's Advantage FCU
A Thriving Economy	Small Business Navigation	SCC entity registrations; City permits (EnerGov); business-license workflow	1717 Collective (Heather Lyne) / Sharon Ebert (DCAO Community Development)
Inclusive Communities	Immigrant Service Discovery	Trusted-org list (OIRE); language-access content; service inventories	ReEstablish Richmond (Kate Ayers, Helen Rai) / Aneesah Smith (NCS)
Inclusive Communities	Cross-Agency Service Navigation	Housing, workforce, and reentry intake records; case-management exports	YMCA of Greater Richmond (Will Thornton) / Karla Almendarez-Ramos (NCS)
Built Environment	Infrastructure Project Visibility	GeoHub project layers; CIP roadmap; VDOT feeds	Partnership for Smarter Growth (Richard Hankins) / Al Wiggins (DCAO Operations)
Built Environment	Fleet Operations & Street Service Tracking	DPW Fleet GPS (not exposed); snow-tier maps; 311 service records	Groundwork RVA (Nathan Burrell) / Al Wiggins + Daniel Klein
A City That Tells Stories	Arts & Cultural Event Discovery	CultureWorks calendar; gallery and venue feeds; neighborhood email lists	CultureWorks / Chaya Braxton (Office of Strategic Communications)
A City That Tells Stories	Resident Stories as Civic Insight	No structured corpus exists today; Valentine Museum oral histories as anchor	The Valentine (Meg Hughes) / Sylvio Lynch III (History & Culture Commission)

Table 1.1: The 14 practical problem statements with their required data and the civic partner or city ally who owned access. Blue-sky problem statements (14 more, two per pillar) are listed in the per-pillar research repositories under `pillar-repos/`.

End to end

The full architecture, read top to bottom, is: **public strategy → seven pillars → 28 problem statements → weekend builds and AI-expanded ideas → evaluated idea corpus → ideas with named owners → prototypes that survive contact with real institutions → production use inside city government, partner orgs, or community channels.** The funnel narrows at every stage on purpose. Stage 1 deliberately over-produces problem statements; Stage 2 over-produces demos; Stage 3 strips both down to what an institution can absorb. The replication test for any other city is whether each of

those arrows survives without hand-waving — especially the last one, where most civic hackathons stop.

Templates → `templates/05-track-strategy/`

Chapter 2

Pre-Event Infrastructure

Pre-event work is the leverage point: the weekend is short, and preparation is what determines whether teams can use it. The Richmond event ran a roughly three-month pre-event window; six months would be better.

2.1 Origin

Richmond's hackathon was launched in late 2025 by Ford Prior and Michael Kolbe — both Richmonders, both with prior hackathon attempts behind them, both already operating in or with the Mayor's office under Mayor Danny Avula's MAP — who pulled in Crystal Harvey within the first week and ran the five-step sequence below in roughly twelve weeks.

The five-step pre-event sequence

1. **Find the vision and bring in the Mayor.** Anchor the event to an existing public strategy (the seven MAP pillars) so the Mayor's office had a reason to lean in rather than a request to evaluate. The City liaison did not need to be recruited: Michael Kolbe was already inside the Mayor's office, which compressed weeks of relationship-building into the kickoff conversation.
2. **Build the core team.** Ford called Crystal Harvey first (immediate yes); from there: Heather Lyne (entrepreneurial ecosystems, support lounges), Sinclair Jenks (marketing / communications), Will Melton (partnerships / community), Mike Huddleston (sponsorships / corporate), and the wider crew listed on the title page. Each owned a domain rather than a checklist.
3. **Set the three success goals.** Every later decision was traced back to these:

   a) Generate positive press for civic innovation, during or after.

   b) Give participants an experience worth telling stories about.

   c) Make the Mayor's office want to come back next year and invest more.

   Treat the three goals as operating constraints, not marketing copy. "Make the Mayor's office want to come back" filters out demos that are impressive but disconnected from MAP priorities. "Participants tell stories" justifies real spend on venue, food, kickoff, and awards rather than treating them as overhead.
4. **Call friends in high places — the VIP Happy Hour.** Attached to a larger conference Ford was already organizing in December, the happy hour announced the event to the Mayor, his

staff, and a group of local executives in one room and collected the first real opt-in list. Roughly 80 of the people in that room later showed up to the event itself. Detail in §2.2.

5. **Stand up the event in parallel.** Once the Mayor was in and the team was set, four threads ran simultaneously:
   - **Venue and logistics** — everyone swarmed; Heather Lyne floated across whatever needed picking up.
   - **City Hall and community engagement** — led by Michael Kolbe, with relationships built ahead of the off-site (§2.5).
   - **Fiscal sponsor + scholarship and sponsorship flow** — Crystal Harvey set up the financial pipes so dollars could land before they were spent.
   - **Everything else not otherwise assigned** — Ford, acting as the partial-time generalist who unblocked whichever thread was at risk that week.

Replication lesson: write down the small number of outcomes the team cares about, and use them as a filter for every later decision. The Richmond goals are not the right goals for every city. The discipline of having only three is.

Templates → `templates/01-civic-alignment/`

2.2 The Happy Hour

Format: December VIP happy hour, piggybacked on a larger conference Ford was already organizing. **Audience:** the Mayor, his staff, and roughly 100 local City, nonprofit, and private-sector executives in one room. **Ask:** sign up, give your email, get involved. **Yield:** roughly 80 of those people later appeared as VIP attendees of the event itself.

The happy hour is not a social side-event; it is the partnership strategy. Getting City, nonprofit, and private-sector executives into the same room before any operational work starts compresses months of one-on-one relationship-building into a single evening.

Replication lesson: do this early, do it before the public announcement, and treat it as the operational kickoff for the partnership graph the event will run on.

2.3 Engaging City Hall

Opening posture: exploratory, not transactional. "How might groups of innovators help with the problems the City is already trying to solve?" lands; "can we run a hackathon in your city?" does not.

Framing that worked: *"We want to help deliver your plan."* Every element of the event then traces back to a specific MAP pillar.

Minimum viable City engagement: a mayor or city manager willing to lend their name to the problem statements. **Full version:** a city that co-authors them. Co-authoring takes time. Plan for six months; ten weeks is too short for serious alignment.

Owner in Richmond: Michael Kolbe, who held existing relationships across DCAOs and department leads and could open the right doors without a cold-introduction backlog.

Templates → `templates/03-stakeholder-map/`

2.4 Track Governance: The 3C Committee

Model: a three-seat committee per pillar, recruited before the off-site so each pillar enters Lost Office (§2.5) with named people in named roles.

- **City** — City Hall or agency representative who can speak to operational reality.
- **Community** — nonprofit or civic partner with domain expertise and resident-facing reach.
- **Capacity** — private-sector operator, builder, funder, or implementation partner who can translate ideas into cost, staffing, delivery path, and post-event ownership.

Recruiting split: Heather Lyne recruited 2–3 nonprofits per pillar (Community); Michael Kolbe recruited 2–3 City representatives per pillar including DCAOs and department leads (City). Splitting the work along existing relationships takes recruiting from months of cold outreach to 2–3 weeks per pillar.

Capacity is the seat that doesn't fill itself; recruit it deliberately, because what an empty Capacity seat costs is implementation realism, not budget.

Replication pattern: recruit all three seats before the off-site. Treat Capacity as required. The off-site is where committees *converge*; it is not where they get assembled.

Templates → `templates/04-governance-model/`

2.5 The Off-Site: Lost Office Collaborative

When: once the 3C committees (§2.4) were stood up. **Where:** Lost Office Collaborative (`lostoffice.co`) in Richmond. **Facilitator:** Christian Markow, co-founder of the studio. **Why a neutral off-site:** City Hall conference rooms invite position-defending; nonprofit boardrooms invite mission-defending. A neutral creative space made it easier for City officials, nonprofit directors, and civic partners to think openly about underlying problems instead of defending existing programs. **Christian's job:** push past familiar talking points and into specific, scoped problem statements the City could point teams at.

Outputs per pillar (one session):

- **Two practical problem statements** — bounded enough to be a weekend prototype, tied closely to a MAP priority. These became the 14 published problem statements.
- **Two blue-sky prompts** — longer-term, less structured, more exploratory. Used post-event as seed material for the AI-generated idea expansions (the 14 BLUESKY ERRC ideas in the post-event corpus).

Replication pattern: when converting public priorities into hackathon problem statements, host the working sessions in a neutral creative space and use a facilitator whose job is to produce concrete outputs, not to summarize the conversation. Run the pipeline pillar → metric → dataset → lever → statement → stakeholder validation; the output is fewer statements, each with a named dataset and a named lever.

2.6 Pre-Event Research Repositories

For each pillar, the organizing team produces a research repository that teams can clone at kickoff. Each repo contains the pillar's problem statements, supporting research with citations, dataset and API references, prior-art notes, and structural-wall warnings (see Table 2.1).

The research pipeline is AI-assisted but human-curated. Parallel AI Systems' deep-research workflows run a consistent set of PhD-level research questions against every problem statement,

returning structured reports with citations. City reviewers vet the output for factual alignment before it becomes ground truth for participants. AI does synthesis; humans approve.

Hardest infrastructure lesson: **map the structural walls before the event, not after.** A structural wall is a vendor-owned system or closed data source that consistently blocks ideas. The specific walls differ by city. The constraint of mapping them in advance does not. Teams that know about the walls plan around them; teams that don't hit them on Saturday afternoon.

Wall	Vendor / owner	What it blocks	Workaround used in Richmond
311 / AvePoint	AvePoint on Dynamics 365	Public routing API; live ticket status; resident feedback loop	Build standalone triage / wizard tools; ship without a 311 write path
OnBase / RAPIDS	Oracle ERP + OnBase ECM	Contract PDFs and procurement workflow data	Use Socrata City Contracts (`xqn7-jvv2`) and Payment Register (`5y73-enav`); skip the PDFs entirely
EnerGov	Tyler Technologies (replacement RFP pending)	Permit data; building / land-use workflows	Read public ArcGIS layers (Dev Tracker, parcels); manual ingest only where required
B2GNow	OMBD vendor directory	MBE vendor lookups behind the highest-ownership Economy ideas	Use Socrata Payment Register; vendor-side public data; manual MBE match against City contracts
DPW Fleet GPS	DPW (telemetry not exposed)	Real-time fleet tracking, snow / street-service progress	Use 311 service-request history + published snow-tier maps as a lagging proxy

Table 2.1: The five structural walls in Richmond's stack, named once so participants and post-event reviewers can recognise them on sight. Every demo that brushed against one of these systems took the workaround in the right-hand column. Replication test: a city without its own version of this table is one weekend away from teams hitting the walls live during the build.

Recommended repo structure (any working layout is fine; the contents matter more):

- `problem-statements/` — the two pillar statements, plus blue-sky prompts
- `research/` — batch-research outputs with citations
- `datasets/` — dataset references, API endpoints, sample queries
- `structural-walls/` — known vendor or data blockers and workarounds
- `prior-art/` — existing solutions, civic precedents
- `starter-code/` — a minimal scaffold teams can clone

Make the repos opinionated about AI usage

The working assumption is that most teams code with an AI assistant (Cursor, Claude Code, Copilot, ChatGPT, or similar). Pillar repos are *opinionated* about that, not ambivalent — so the agent picks up the same context, conventions, and guard-rails that a human reading the repo would.

Two files do most of the work:

- `AGENTS.md` at each pillar-repo root. This is a de-facto convention readable by Cursor, Claude Code, OpenAI Codex, and other AI coding agents. It declares: what data is canonical for this pillar, which structural walls to respect, what *not* to invent (no synthetic data, no fabricated APIs), and what tone to use when drafting documentation. Teams using an agent inherit

those constraints automatically; teams not using one read the same file as a human-readable contributor guide.

- `.cursor/skills/` (and the equivalent `~/.claude/skills/` layer for Claude Code) — small procedural **Skills** that wrap repeatable agent workflows. Examples that shipped or were drafted for the Richmond event: `onboard-late-idea` (ingest a Devpost link or post-event submission and produce a canonical idea page end-to-end), `route-hackathon-feedback` (route retro notes into the right GitHub issues and replication-toolkit sections), `extract-judging-data` (refresh the master judging xlsx into normalized CSVs and SQLite), `generate-rvahacks-data-module` (wire a new per-idea data dimension into the website). Each Skill is a short Markdown file the agent reads on demand instead of being re-prompted from scratch.

What this enables: a participant arriving on day one with an AI assistant can ask "what are the structural walls for this pillar?" or "what dataset do I use for affordable-housing compliance?" and get a consistent, citable answer drawn from the repo — not a plausible-sounding hallucination. Teams that use the agents onboard faster to their pillar's data, citations, and known blockers, and produce fewer demos that break on contact with reality at judging.

Replication pattern: ship every pillar repo with a real `AGENTS.md` and at least one or two Skills covering the workflows you most want done consistently. Do not assume teams will figure out your conventions by inspection; write them down for the agent and the human in one place.

Templates → `templates/06-problem-statement-development/`, `templates/07-data-readiness/`

Chapter 3

The Event

Most hackathons already run the weekend reasonably well. This chapter documents the decisions that distinguished Richmond's.

3.1 Three-Day Structure

Mentor and volunteer staffing

Both mentoring and volunteering are structured into named lanes, not left to "someone-will-figure-it-out" coverage.

Mentor lanes. Mentors are organized into three specialty lanes — **Hacker / Tech**, **Design**, and **Pitch / Demo**. Aim for every Hacker/Tech Mentor slot to be fully staffed across Friday and Saturday, with Friday evening achieving full coverage across all three lanes; mentor blocks should be filling before the event begins.

Volunteer lanes. Volunteer roles cover the operating layer of the weekend, in six functional groupings:

- **Arrival** — Breakfast Crew, Welcomer, Registration, Usher. Fully staffed for Friday's opening, giving participants a smooth check-in and orientation.
- **Room management** — Breakout Leads anchoring each breakout room.
- **Help desk** — on-call coverage for participant questions throughout the weekend.
- **Flexible support** — Floaters, available for unpredictable on-site needs.
- **Food** — meal coordination across the weekend.
- **Closing event** — Banquet Crew running the Sunday-evening close.

Specialty lanes for mentoring, functional groupings for volunteering. The split makes community participation concrete and prevents unowned tasks from accumulating on the lead organizer's plate. The shift framework lives in `rvahacks/src/data/grayTasksPlan.ts`; mentor and volunteer routing runs through the Slack workspace defined in `infra/hackathon-slack/`.

Registration mechanics

Two signup paths feed one Slack workspace; the team count comes from Slack, not from the registration form. Registration opens roughly eight weeks before the event. A **team lead** can register and add up to four teammates in a single form, or an individual can register solo and find a team during the kickoff (and earlier, in the `#find-a-team` Slack channel). Once a registrant confirms their email, the provisioning system in `infra/hackathon-slack/` auto-adds them to the appropriate participant channels. Reconciliation between the two paths happens automatically: team leads' rosters

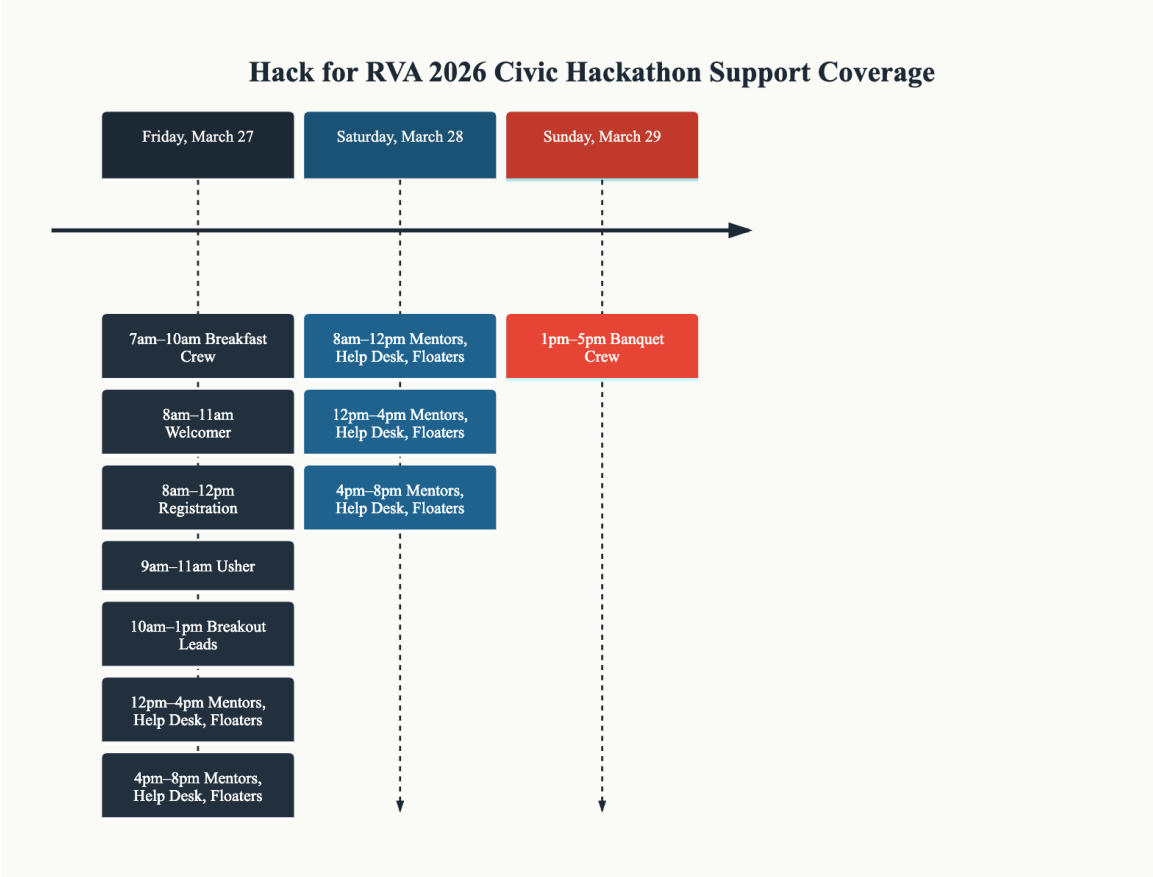


Figure 3.1: Support coverage across the weekend. Friday morning ran on five named arrival roles plus Breakout Leads. From Friday noon through Saturday evening, a backbone of mentors, help desk, and floaters held continuous coverage. The Banquet Crew anchored the Sunday close. The chart maps to the shift framework in `rvahacks/src/data/grayTasksPlan.ts`.

attach to a `team_id` in the registrant CSV; solo registrants merge into teams as they self-organize. The team count is recomputed from the active Slack roster, not from the form.

The paywall as no-show mitigation

Organizers do not want a cost barrier. Add one anyway — a small registration fee per participant — because the single largest source of investment-calibration error is unmitigated no-shows. A small volunteer-led organization cannot absorb the cost of catering, swag, mentor coverage, and credit allocations sized for an attendance number that turns out to be 30% lower than registered. Set the fee low enough not to exclude anyone for whom $X matters, with a no-cost scholarship pathway for anyone who needs it. Any city replicating the model with a better idea on this should publish it; Richmond treated the fee as the least-bad option available on short notice, not as the right answer.

Day	What Happens
Friday Mar 27	Mayor opening. Pick a pillar. Form a team. Pillar breakouts with the 3C committees. Knowledge transfers (data sources, structural walls). Hacking begins later the same day.
Saturday Mar 28	Open-format build sprint with optional mentor support across technical, pitch, and design lanes.
Sunday Mar 29	Final submissions (2-min video + links via Slack workflow), demos, live judging across three award tracks, family-friendly closing banquet.

Table 3.1: 38 teams submitted demos.

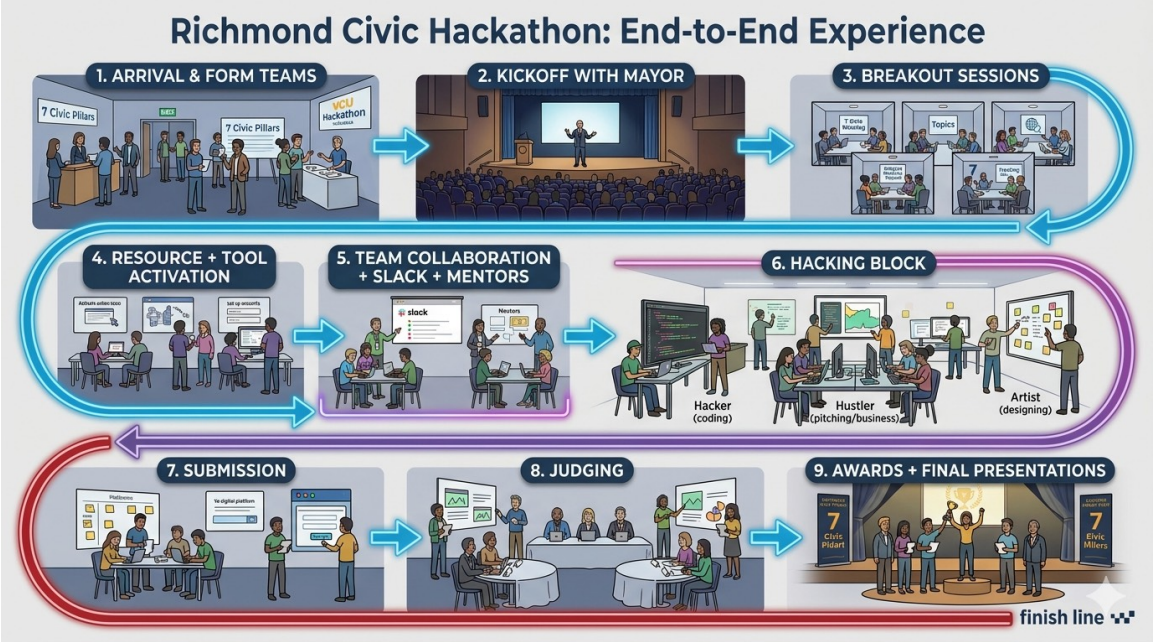


Figure 3.2: The end-to-end participant experience the model assumes. Pre-event onboarding, the weekend build sequence, and the post-event Champion / Working Session split shown as a single journey. Plan for this end-to-end from the start; treating onboarding, build, and post-event as three separate problems is what produces drop-off between phases.

Templates → `templates/12-event-operations-plan/`, `templates/10-participant-experience-plan/`

3.2 Platform and Tooling

The tooling stack is minimal. The operating discipline: organizers should know the chosen tools better than the participants do.

Need	Chosen	Cost	Complexity	Why	Other options
Central comms	Slack	sponsored	low	broadly adopted in the audiences the event recruited from; organizers already power users	Discord
Frontier AI for participants	Cursor ($50 / team credit)	sponsored	low	best-in-class agentic coding for non-experts; clone a repo and go	—
AI fallback for credit overruns	Featherless AI	sponsored credits	low	open-model fallback so teams could not stall mid-build; not selected, only reserved	—
Code + knowledge sharing	GitHub	free	low	per-pillar repos teams could clone; same place data lived	Word docs / per-team zip files
Submissions intake	Slack workflow → Google Sheet	free	low	participants already lived in Slack; no second tool to learn	Devpost
Cloud (sparse)	AWS $100 vouchers, on request	sponsored	medium	on-demand only when a team needed it; not broadcast	—

Table 3.2: Tooling decisions: one row per need, with the chosen tool, the alternative considered, and the reason for the choice. The Slack provisioning, channel routing, and submission collection logic are documented in `infra/hackathon-slack/`.

Slack workflows

Four Slack workflows carry the operational load during the event. Each is small, single-purpose, and runnable by an organizer who can read YAML.

1. **Team formation intake** (`#find-a-team`). Solo registrants post a structured message (skills, pillar interest, availability) via a workflow form; team leads watch the channel and recruit from it. The workflow turns the channel into a queryable list rather than a free-form chat.
2. **Mentor help claim** (`#help`, documented in `config/hackathon.yaml`). A team posts a help request; any mentor reacts with `:eyes:` to claim it; the team is notified when claimed and again when the mentor reacts with `:white_check_mark:` on resolution. Replaces a triage rotation with a self-serve queue.
3. **Final submission intake.** The Sunday submission form accepts a 2-minute video upload plus up to two links (live demo, GitHub repo, slides — whichever the team has). Submissions land in a Google Sheet that feeds the live judging spreadsheet directly; no Devpost, no second platform.
4. **Post-join routing** (Slack Bolt app, `router.py`). On the `team_join` event, a small Bolt app reads the new member's role record (hacker, mentor, judge, sponsor, volunteer, organizer) and adds them to the right channel set automatically. Replaces a human triaging the workspace as people arrive — which at ∼250 participants was not optional.

Replication note: workflows are cheap to build and easy to mis-design. The pattern that works: *one workflow per moment a participant has to do something specific*, with the workflow doing the routing/notification for them. Anything more elaborate becomes a debugging tax.

Anticipated onboarding hazards. Four failures recur: **(1) GitHub CAPTCHA on event day** — have a manual approval path or pre-create accounts at registration. **(2) Registration-form pillar-**

routing bugs — test the registration → Slack pipeline with real accounts a week before kickoff. **(3) Live channel provisioning** — run it before the kickoff session, not during it. **(4) Help-desk latency** — aim for sub-15-minute first response in `#help`, and seed `#introductions` with a workflow prompt.

Templates → `templates/13-technology-setup/`

3.3 Award Winners

Nine awards in total: one Mayor's Choice, one Moonshot, and seven Pillar Awards (one per MAP pillar). All scored on the same six-category rubric with different weightings.

Award	Winner	What It Was	Problem Statement
Mayor's Choice ($5K)	MIRA	Conversational AI routing resident service requests	Resident Service Navigation (City Hall)
Moonshot ($1K)	Ride Shift RVA	Equity-framed redesign of Richmond's transit network	Blue-sky / cross-pillar (transit equity — not on the 14 practical)
Pillar: City Hall	Text 311 (Byrd Park Geese)	Text-based 311 interface, no app required	Resident Service Navigation
Pillar: Neighborhoods	SIMBY	Multilingual development discovery (EN/ES/AR)	Neighborhood Development Discovery
Pillar: Families	Unboxed RVA	Project marketplace for paid youth work with local small businesses	Youth Employment Pathways
Pillar: Economy	TradePath RVA	Workforce and MBE navigation	MBE Contracting Discovery / Small Business Navigation
Pillar: Inclusive Communities	804Me	Physical hub plus anonymous web app for finding community resources	Immigrant Service Discovery
Pillar: Built Environment	StreetLights	Plain-English lookup of nearby infrastructure projects by address	Infrastructure Project Visibility
Pillar: Storytelling	Fulton Oral History	Interactive map surfacing the Fulton Oral History Project transcripts	Resident Stories as Civic Insight

Table 3.3: Award winners mapped to the practical problem statements they targeted. Mayor's Choice (MIRA) scored Tier B on the Implementation Reality rubric: the prize winner and the most shippable demo were not the same project. Source: `rvahacks/src/data/ideas.ts`.

3.4 The Judging System

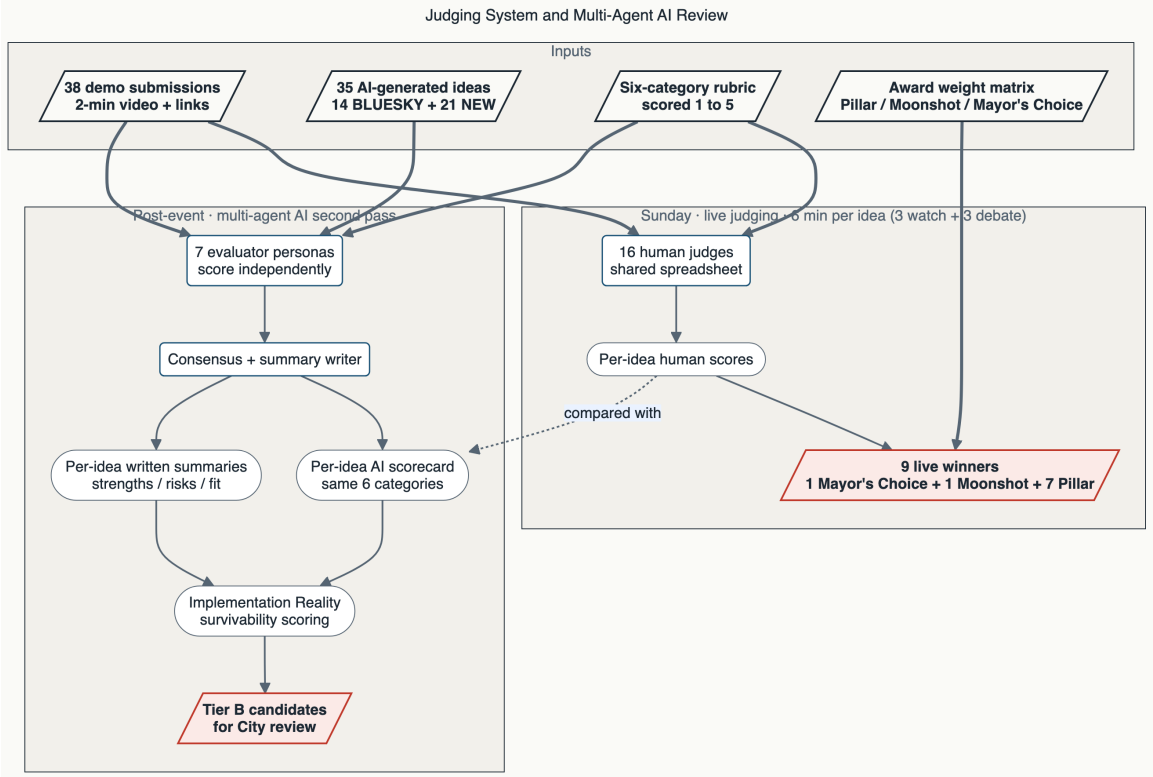


Figure 3.3: Inputs into two judging tracks. The 38 demos and 35 AI-generated ideas, the six-category rubric, and the per-award weight matrix feed two parallel evaluations. **Sunday live judging** (right): 16 human judges score in a shared spreadsheet at 6 minutes per idea (3 minutes watching the 2-minute video plus discussion, 3 minutes debating); the weight matrix produces 9 live winners. **Post-event multi-agent AI review** (left): seven evaluator personas score independently, a consensus + summary pass produces a per-idea AI scorecard and written summary, Implementation Reality scoring is applied, and Tier B candidates are surfaced for City review. The human and AI scores are compared, not merged.

The judging process solves two problems at once: produce live winners by Sunday close, and produce material the City can use later. The live process handles the first under time pressure; the post-event AI-assisted review (Chapter 4) handles the second.

Submission and live cadence

Submission format. Every team submits a single artifact package via a Slack workflow: a **2-minute video** (the team's demo, recorded however they want) plus up to **two links** — typically a live demo URL, a GitHub repo, or a slide deck. The Slack workflow writes each submission directly into the live judging spreadsheet. No Devpost, no second platform.

Live cadence. Sunday live judging runs at **6 minutes per idea**: 3 minutes for the video plus brief Q&A, 3 minutes for judges to debate and score in the shared spreadsheet. The clock is strict

because 38 demos in one afternoon does not allow drift. The trade-off: 6 minutes is enough to land an impression, not enough to adjudicate a hard call. The post-event AI review exists in part to catch what the live cadence cannot.

Framing for participants. Frame the judging — explicitly, multiple times, from the kickoff through the closing remarks — as *a weekend of positive innovation for the City, not a competition for prize money*. Cash awards matter, but the framing keeps teams focused on the work and makes the gap between live winners and Implementation-Reality top-tier projects easier to explain.

Rubric

Every demo is scored on six categories, 1–5. The Richmond panel had sixteen judges (about nine core organizers, the Mayor and aides, Thad Williamson from the City, plus civic partners and problem-statement contributors, some virtual) — a useful reference point for sizing a replicating panel. The six categories:

- **Impact** — does it meaningfully address a civic problem?
- **User Value** — does it serve a specific user well?
- **Feasibility** — could it realistically be piloted?
- **Innovation** — is the approach fresh or original?
- **Execution** — is there a working prototype, not just slides?
- **Equity & Inclusion** — does it consider who might be left out?

Sixteen judges × thirty-five teams × six categories produces 3,360 individual score cells. The full data lives in `data/`.

Three Awards, Three Weightings

The same six-category scores produce three different award totals via weighting:

- **Pillar Award** — balanced rubric, emphasis on Impact and User Value. Best fit to the pillar's specific problem statement.
- **Moonshot Award** — Innovation-weighted. Bold thinking with long-term potential.
- **Mayor's Choice** — Feasibility-weighted. Most plausibly implementable in the City near-term.

The exact weights are recorded in `data/award_weights.csv`. The design rationale: judges should not have to learn a new scoring system three times. They score once; the spreadsheet applies three weighting schemes.



Figure 3.4: One rubric, three lenses. The same six rubric scores receive different weights depending on the award: Mayor's Choice is Feasibility-heavy, Moonshot is Innovation-heavy, Pillar Awards are balanced. The right-hand sidebar shows the question each award answers. Values from `data/award_weights.csv`.

Prize money, demos, and the closing banquet

Cash awards. Mayor's Choice $5,000, Moonshot $1,000, plus the seven Pillar Awards. **Distribute and collect W-9s on-site** at the closing event so winners walk out with the paperwork to receive their checks rather than waiting for a follow-up email. Cities replicating the model should plan for this; chasing W-9s after the fact is a drag on the post-event weeks the organizers need for the continuation pipeline.

Live demos at the closing event. Even though scoring is locked in by the time the closing event starts, winning teams demo live for the room. The choices are already made; the demos are for the audience, the Mayor, the partners, and the press. Treating the closing as a *showcase* rather than a *decision moment* keeps the energy on the work rather than on suspense.

Family-friendly closing banquet. Close with a buffet-style banquet open to participants' families. Hackathons often end with everyone exhausted and scattering, which deletes the relationships the weekend just built. A banquet with families turns the closing into a community moment and gives partners, sponsors, and the Mayor's office a reason to stay in the room rather than slip out after the awards.

Known judging issues

Judging is easy to assemble late and easy to under-brief. The common failure pattern: judges arrive without having been briefed on the rubric, participants only encounter the criteria when scores are returned, and the live process becomes rubric-informed scoring in a shared spreadsheet under time pressure. Ranked-choice voting is worth considering as a tie-break layer on top of rubric scoring,

but it has to be described accurately up front rather than promised in the one-pager and then dropped at the door.

Announce the rubric at kickoff. Participants should hear the six categories, the 1–5 scale, and the three award lenses on Friday during the Mayor's opening, not infer them from the Slack channel mid-build. Treat the kickoff slot as a non-negotiable rubric briefing.

Voting is hard. Treat it that way. Live judging is a difficult balance of gut and detail under time pressure and deserves the same design attention as the build sprint. Build for **consensus** (judges disagreeing in the same direction is a signal, not noise), **consistency** (the same judge applying the same criterion the same way across every demo), and **transparency** (judges and teams should both be able to see how the score was produced). Collect feedback during and after, and adapt the format quickly if it is not serving the participants. There is no fully-worked playbook for this yet; one is owed.

Judging needs a single dedicated owner (a **Judging Lead**) who owns: judge recruitment, rubric briefing, publishing criteria to participants, designing the live flow, score management, and keeping the live awards distinct from the post-event implementation review.

Templates → `templates/14-judging-system/`

Chapter 4

Post-Event Evaluation

This is the part nobody else does. The post-event pipeline is what separates a well-run hackathon from one that produces sustained public value.

4.1 The 73-Idea Corpus

The hackathon produced 38 demos built by humans during the weekend. Those demos covered roughly half of the 28 published problem statements (14 practical + 14 blue-sky); the rest of the published statements got zero or near-zero human attention. To produce a complete corpus for City review — one where every published statement had at least one idea against it — organizers used AI tools after the event to generate additional ideas, calibrated to fill the gaps. The combined corpus reached **73 ideas**: 38 human + 35 AI.

For this toolkit, only the underlying distinction matters:

- **Human** — the 38 weekend demos
- **AI** — 35 post-event AI-generated ideas

Both sets were evaluated on the same rubrics. (The CSV files in `data/` use historical labels `DEMO`, `NEW`, and `BLUESKY`, mapping to Human and AI respectively.)

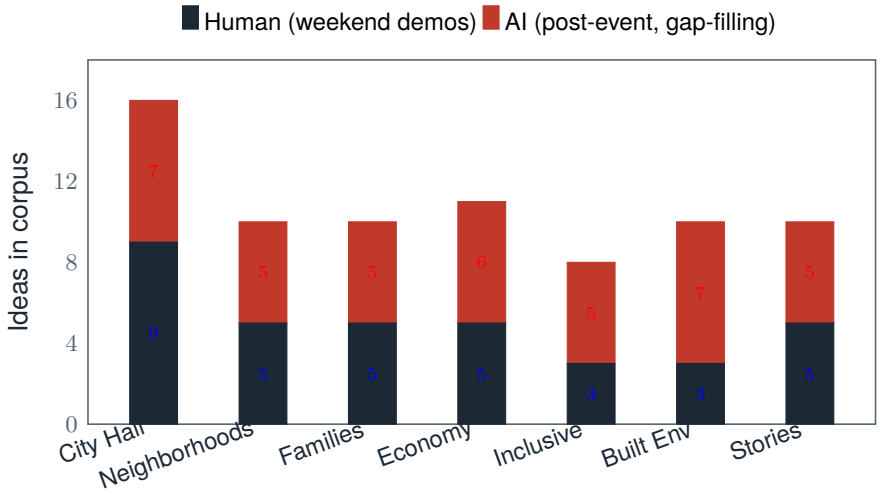


Figure 4.1: Ideas per pillar, human vs. AI. Pillars where the weekend attracted fewer teams (Inclusive Communities, Built Environment) received the largest AI top-up so every published problem statement ended up with at least one idea against it. Counts taken directly from `data/teams.csv` (human) and `data/ai_only_ideas.csv` (AI), with pillar labels normalized; the totals (35 + 40 = 75) are slightly higher than the canonical 73-idea corpus because the CSVs include a small number of naming variants and near-duplicates that `rvahacks/src/data/ideas.ts` resolves into 38 + 35. The shape — gap-filling concentrated on the under-attended pillars — is what matters here.

4.2 Two Evaluation Layers

Every idea is evaluated on two independent dimensions:

- **Desirability** — the live judging rubric (six categories, 1–5). Answers: *is this a good idea?* The post-event AI review re-scores every idea on this same rubric, effectively adding one more judge — a consistent one, applied uniformly across all 73 ideas.
- **Survivability** — the Implementation Reality rubric (five dimensions, 1–5, max 25). Answers: *can this actually ship?* Survivability is *generated post-event only*. Live judging is structurally bad at it: a 6-minute window cannot interrogate the dependency graph of a prototype, and survivability lives in those dependencies. The Implementation Reality pass is a hunt for the **black swan inside each prototype** — the single upstream blocker, vendor-locked system, or quiet maintenance cliff that would silently kill the idea between the demo and a pilot. That hunt is slow work, and it does not fit on stage.

The two dimensions are independent on purpose. An idea can score 95/105 on the Pillar Award and 8/25 on Implementation Reality. The point of running both is to avoid the failure mode where the prettiest demo wins the award but the most shippable idea quietly disappears (Figure 4.3 makes that independence visible across the corpus).

Templates → `templates/15-ai-assisted-evaluation/`

4.3 Implementation Reality Rubric

Five dimensions, each scored 1–5:

Dimension	Score 5 looks like	Why it matters
Ownership Clarity	Named department, mandate, capacity exist today	Without an owner, ideas become orphans
Data Readiness	All public, API-accessible, actively maintained	"If we just had the data" = most common last words
Tech Stack Fit	Only public/commodity infrastructure	Anything outside = 6–18 month procurement cycle
Integration Surface	Standalone, reads public APIs, no City system changes	Each integration point is a potential structural wall
Maintenance Burden	Automated, <2 hrs/month	Understaffed departments abandon tools quickly

Table 4.1: Implementation Reality dimensions. Maximum total: 25.

The pattern across all five dimensions: the data an idea consumes is the most important architectural decision a team makes.

Tier bands: **S** (21–25) ship-it today; **A** (16–20) strong with 1–2 solvable blockers; **B** (11–15) promising, meaningful investment needed; **C** (6–10) aspirational, structural barriers; **D** (1–5) not viable without major change. The actual distribution across Richmond's 72-idea corpus is shown in Figure 4.2.

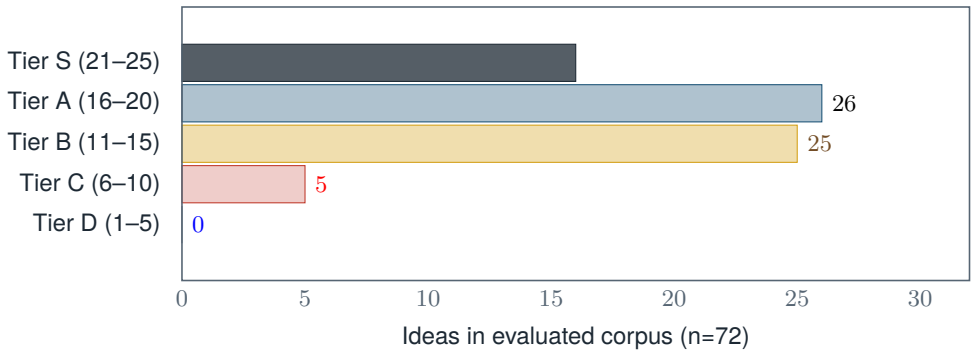


Figure 4.2: Implementation Reality distribution across the 72-idea evaluated corpus (1 idea was not scoreable). Tier S + A together is 42 ideas, or 58% of the corpus — the cohort the City could realistically review without further triage. Counts sourced from `rvahacks/src/data/ideas.ts`; ranked breakdown in `2026-annual-hackathon/ideas/IMPLEMENTATION_REALITY_SYNTHESIS.md`.

Tier B is not a consolation category. It is the second group of promising ideas worth attention — often less polished on stage but more feasible, better aligned to City priorities, or easier to pilot than some of the live winners.

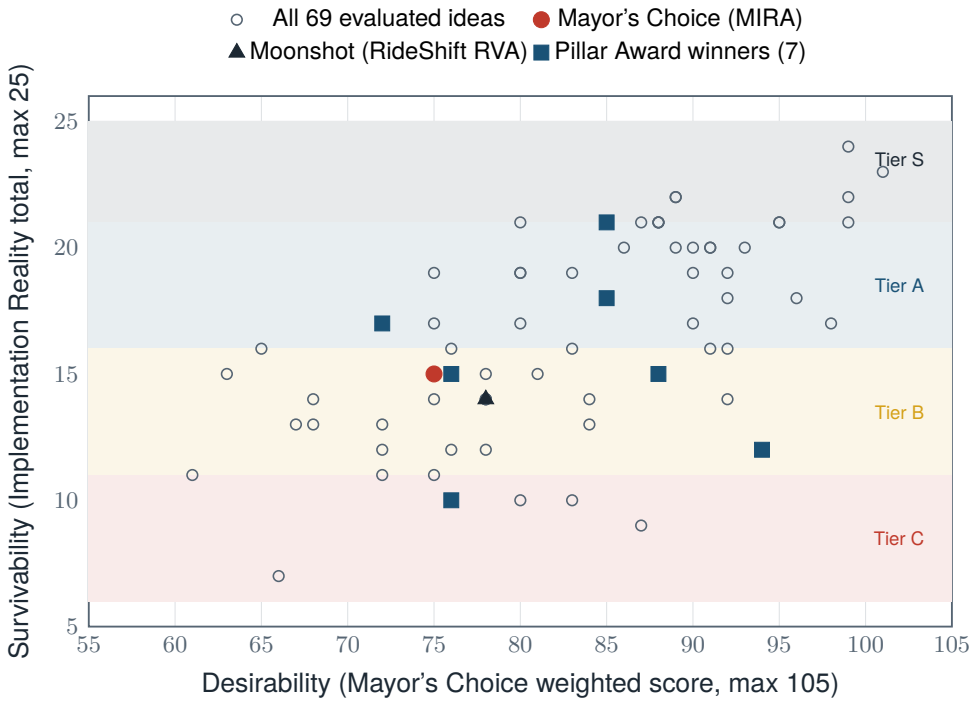


Figure 4.3: Desirability vs. survivability for the 69 evaluated ideas that carry both scores. **X-axis:** Mayor's Choice weighted score (applied uniformly across the corpus by the post-event AI pass). **Y-axis:** Implementation Reality total. The two dimensions are visibly independent: MIRA (Mayor's Choice winner) sits at (75, 15) — Tier B on survivability. StreetLights is the only live winner inside Tier S, at (85, 21). The cluster of high-desirability / low-survivability ideas in the lower right is exactly the failure mode the two-axis evaluation is designed to surface: ideas the live judges would have picked but that would not have survived contact with the City's data, vendors, or capacity. Source: `rvahacks/src/data/ideas.ts`.

4.4 AI-Assisted Review

What this layer buys. The status quo at most civic hackathons is that the live winners get all the post-event oxygen and the rest of the corpus quietly evaporates. AI-assisted review gives *every* idea proportionate attention — including the ones that did not make it on stage Sunday and the AI-generated ideas that filled gap pillars. That redistribution of attention is the practical value, not the per-score precision.

Methodology. The post-event review uses a multi-agent structure: seven evaluator personas score each idea on the same six categories used by human judges, producing a comparable AI evaluation alongside the human panel. The personas are defined in `2026-annual-hackathon/agents/CIVIC_EVAL_PERSONAS`; the scoring results live in `data/ai_evaluations.csv`; the data flow (inputs, parallel human + AI tracks, comparison rather than merge) is shown in the judging diagram in Chapter 3.

The structure follows the pattern documented in **MAFA** (*Multi-Agent Framework for Annotation*) by Hadidi et al., deployed at JP Morgan Chase's Consumer & Community Banking ML team for enterprise-scale annotation workloads.[1] Specialized agents apply structured reasoning; a judge-

[1] "MAFA: A Multi-Agent Framework for Annotation," arXiv:2505.13668. The civic-eval workload is conceptually similar: a fixed rubric, a large corpus, and a need for consistent application across items that no single human reviewer

based consensus mechanism merges their outputs; the same scaffold hosts different rubrics. The civic-eval workload here fits the same shape.

Models. All AI evaluations were run on a single model — Anthropic Claude **Opus 4.7** — to keep persona-to-persona variation a function of the prompts, not of inconsistent backbones. The seven personas differ in role, reading priorities, and scoring tendencies; they do not differ in underlying capability.

What an agent's output looks like in practice. The deeper research the AI review produced is most useful when it surfaces something a 6-minute live judging window could not. As one example, the lifecycle review for the RVA MAP Tracker demo (Pillar 1) opened with this honest assessment:

> *"Misaligned — strong concept, addresses government accountability (Blue Sky), not either targeted problem statement […] Both Thriving City Hall problem statements target specific operational pain. Neither asks for a progress tracking dashboard for the Mayor's Action Plan."*
>
> — `data/ai_rationale_packets/map-tracker.md`

A structurally good idea aimed at the wrong target is exactly what disappears under live time pressure. The human panel rewarded the demo on execution and impact (it was well-built); the AI lifecycle review captured a separate, also true, observation: it was solving a different problem than the one the pillar published. Both readings are useful. They are not in conflict.

Honest caveat. The framework is an analytical aid, not an authoritative decision-maker. Its purpose is to compare submissions consistently across a larger corpus than human judges had time for and to surface ideas the live process under-weighted. Treat the AI scores as a second opinion that is cheap to produce and uniform in its biases — not as ground truth.

One observed pattern: AI-generated ideas tended to score higher on Implementation Reality than human demos — not because AI is smarter, but because the AI ideas were designed *after* the tech-stack research was complete, with full knowledge of what data existed and what was locked. Teams built in 48 hours without that knowledge. The fix is better pre-event data briefings, not better AI.

4.5 The Public Marketplace

The post-event output is not a press release. It is a publicly browsable **marketplace of fully graded ideas** on the event website, with each of the 73 ideas surfaced as a structured listing. Per idea, the marketplace exposes:

- the demo link or repository (where one exists),
- a one-paragraph summary written from the lifecycle review,
- the human panel score (six categories, 1–5),
- the AI deep-dive grade (same six categories, plus the five Implementation Reality dimensions),
- the per-dimension scoring results and computed award totals,
- agent-written feedback on strengths, risks, and structural fit,
- suggested next steps and concrete action items for the team or a potential City owner.

The audience is dual: **the City** (procurement-aware, looking for the implementation-ready subset across the full corpus, not just the live winners) and **the participants** (every team finally sees how their work was evaluated and what would happen next). Most hackathons publish a winner list and a press release; publishing the entire evaluated corpus — including ideas that did not win and AI-generated ones that filled gap pillars — flips the post-event posture from celebration to triage. Source:

could process in the available time.

`rvahacks/src/data/ideas.ts` drives the website surfacing; the underlying scores and feedback live in `data/` and `data/ai_rationale_packets/`.

Replication note: the marketplace is the artifact that makes the rest of this toolkit honest. If the only post-event output is the nine award winners, the survivability work above has nowhere to land. Plan the publication surface alongside the evaluation; it is not a follow-up task.

Templates → `templates/17-open-source-handoff/`

Chapter 5

Continuation Engine

5.1 Cadence over bursts

High commit activity in the first week after a hackathon correlates with *lower* long-term project survival, not higher. Published research on open-source and hackathon projects has reported this pattern repeatedly. Survival tracks with steady cadence, clear milestones, and a room of people who notice if a team stops showing up.

The continuation engine is designed against that finding: cadence over bursts.

5.2 Day 30: The Phase 2 Launch

Phase 2 launches 30 days after the event, after a quiet period (rest, planning, thank-you notes). The organizing team meets with City Hall to kick off prototype evaluation, workshopping, and integration planning for the strongest ideas.

The Day-30 launch is preserved verbatim below — candid about which ideas hold up, explicit about what the second-layer review is for.

Phase 2 announcement (Day 30, April 2026).
"Phase two is officially underway. Today I sat down with Douglas A. Gernat and Michael Kolbe in City Hall to launch the next phase — prototype evals, workshopping, and integration planning required to move the raw value created during the hackathon into the hands of City Hall and, by extension, into the hands of citizens.
We are now entering a second layer of human review led by city planners, cybersecurity experts, and domain specialists. This group will synthesize the initial human scoring with the AI evaluation to produce a focused list of candidates ready to move toward implementation.
[…] No idea is sacred. Most ideas generated with limited context in a 48-hour sprint will not hold up over time. That is expected and part of the innovation process."

— Phase 2 announcement, April 2026

Three things matter about this checkpoint:

1. **Day 30 starts Phase 2, not implementation.** The realistic post-event timeline is measured in quarters, not weeks. Cities replicating the model should plan for the implementation timeline to be quarterly.

2. **The City side of the meeting was procurement-aware.** City staff were already thinking in implementation-and-procurement terms, not in demo-day terms. A follow-up email from the City framed the Mayor's Choice winner as a candidate "vendor" solution to a Department of Public Safety need and asked how to bring pre-evaluated hackathon solutions into the City within existing procurement requirements. This conversation happens because the City co-authored the problem statements; it does not happen at most civic hackathons.
3. **The multi-agent AI review is part of the official package.** The announcement publicly confirms that the AI evaluation is reviewed alongside the human judging — which is the design intention behind running two rubrics in the first place.

What happens after the announcement: the City's intake process

Day 30 is a launch, not the only meeting. After the announcement, City leadership takes the full evaluated corpus into the City's internal **product-innovation intake process**. Align the timing with an active budget-cycle window, so short-listed ideas can be considered for budget allocation rather than queued for a later cycle.

The intake process belongs to the City, not to the hackathon organizers. City leadership walks Deputy CAOs and department heads through the corpus and produces a short list of ideas judged **ready to propose** — ideas with a credible department owner, a plausible budget hook, and a viable implementation path inside existing City systems.

For the short-listed ideas, the offer back to the makers is simple: come down to City Hall, or jump on the Zoom, and walk the team through the prototype. Optional, voluntary, and framed around helping the City get started — not around contracting. The working assumption is that some makers will want to help the City stand up their idea for the same reason they showed up for the hackathon: the work matters to them.

Considered and not chosen: instant high-value contract matching. One alternative the organizers weighed was using the demo plus the AI evaluation as a pre-qualification signal that could fast-track a maker into a paid contract for the same scope of work. The intuition was reasonable: some of the work is already done, the City has a need, the maker has shown they can build it. The organizers chose against it. Procurement, intellectual-property, and conflict-of-interest entanglements would have made it a multi-quarter legal exercise that exceeded available capacity, and the framing of the event would have shifted from civic innovation to vendor selection. A volunteer-led organization can host a hackathon; it cannot also operate as a procurement front end. Phase 2 therefore assumes voluntary collaboration, not contracting, as the default mode for getting an idea off the ground inside City Hall.

Ideas named in the announcement

The announcement is mostly about process, but it names a small set of specific items worth a one-line orientation:

- **Mira (AI Procurement Intelligence Dashboard)** — the Mayor's Choice winner. A follow-up email from the City framed Mira as a candidate vendor solution to a Department of Public Safety need and asked how to bring pre-evaluated hackathon solutions into the City within existing procurement requirements.
- **The multi-agent AI evaluation** — the post-event review described in Chapter 4, presented to the City alongside the human judging.

- **The second-layer human review** — city planners, cybersecurity experts, and domain specialists tasked with synthesizing the human and AI scores into a focused candidate list.
- **The "reworking funnel"** — the term used at the time for the Path 2 community working sessions described in the next section.

Templates → `templates/16-post-event-pipeline/`

5.3 The Two-Path Post-Event Funnel

Most hackathons end at the awards ceremony. The Richmond model adds a post-event funnel because the awards alone do not tell the City which ideas are worth investing in. The full 73-idea corpus is posted to a review website, and the City's product-innovation intake process (described above) produces the short list. The remaining ideas are routed into a community working-session format.

Path 1: City Champion Path

This is the path the City takes strategic leadership on. The Richmond model treats the post-event funnel as **opt-in for the City and open-source by default**: the organizers extend the courtesy of strategic leadership to the City on any idea coming out of the hackathon, and the City decides which ones it wants to drive. The organizers do not hand the City a list of "owned" ideas, because the ideas are not owned in the first place — they are the output of a public, sponsored event, and they belong to the public.

Where the City takes the lead, organizers stay in a support role: surfacing context, connecting the original builders, providing post-event materials. They do not drive the City's internal timeline. Outcomes are documented but not pressured. An idea can move, stall, or be set aside.

Considered and not chosen: assigning each idea a single owner up front. The original instinct was to assign each short-listed idea to a named department or individual at the moment of handoff. The organizers chose against it. Civic ideas of this kind benefit from **regional collaboration**, not exclusive ownership: if the City pilots a version and a community group also pilots a version, the city wins twice. The City is not competing with the community for these ideas, and locking ownership early would have introduced friction without buying anything. Strategic leadership is offered as a courtesy, not assigned as a contract.

This path also relieves pressure on the City to shepherd every idea. By taking strategic leadership on the ones it wants, the City implicitly gives the rest back to the community.

Alternative worth considering: pillar product owners with two participation tracks. Recruit a civic or startup-support owner per pillar who joins the 3C committee, is present during the build, and owns post-event coordination on the practical statements regardless of who wins. Participants choose between a **managed track** (work against a named statement with the owner present, in exchange for ongoing coordination support) or a **blue-sky track** (unmanaged, broader statements, creative latitude). The original Richmond event ran the blue-sky shape across all fourteen practical statements and produced zero named institutional champions at Day 30; pre-committing ownership through pillar product owners would have closed that gap by design.

Path 2: Civic Working-Session Path

For the ideas the City did not take strategic leadership on. This is not a consolation track. It is the path that keeps the rest of the corpus moving without depending on City capacity, and it is **a city-free**

space so civic groups can coordinate on what they want to do without waiting for an official City position.

In year one, the format is a single in-person session at 1717 Startup Virginia: seven rooms, one per pillar, each with a facilitator, with a rooftop happy hour after. Invitees include nonprofits, civic partners, residents, the original hackathon participants, and additional organizations relevant to each pillar. Inside each room the cadence is concrete:

- **Fifteen-minute walkthroughs.** Each idea (or each maker who shows up) gets a 15-minute slot to walk the room through the prototype, the path the team took to get there, and what they would do next.
- **Workshop dialogue.** Civic experts and community members in the room react: what they would actually use, what would have to change, what is unlikely to land, who else should know about it.
- **Introductions and partnerships.** Facilitators surface matches: a nonprofit that could host the prototype, a community group already running something similar, a working group that could absorb it. The session ends with a list of follow-ups, not assignments.

Three goals drive the format: enable cross-organization collaboration, **invite the people who missed the event** (every working session is also a recruiting moment for the next event), and **build on the momentum** created during the hackathon weekend instead of letting it dissipate.

Considered and not chosen: including the City in the working sessions. The organizers debated inviting City representatives into Path 2. They chose against it for two reasons. First, the two paths exist so the City can stay focused on the short list (Path 1) without being pulled back into every idea. Second, civic groups need a space where coordination does not have an implicit "the City might do this" hovering over the conversation — otherwise the room defers to an absent decision-maker instead of doing its own work. City staff are welcome to attend in an observer capacity but not in a decision-making one.

Two paths, two problems

The two paths solve different problems. The City Champion Path moves the small number of ideas the City has institutional energy for. The Civic Working-Session Path keeps the rest of the corpus in front of people who might be able to do something with it. Neither path alone is sufficient. A one-path model either dumps too much on the City (which most hackathons try to do and then abandon) or sends everything back to the community by default (which makes the City feel like it had no real role).

Replication lesson: design the post-event funnel before the event ends. Decide who reviews the corpus first, what the criteria are for moving an idea into the City path, and what the community working-session format looks like for the rest. Improvising this after the awards ceremony does not work.

Templates → `templates/16-post-event-pipeline/`

5.4 Core roles and responsibilities

The same small group that runs Phase 1 carries Phase 2, which is why the role mix matters more than for a one-shot event. The model relies on a set of named functions, each consistently filled by one person across both phases.

Technical Lead. Owns the technical spine of the event: repositories, AI tooling configuration, the post-event evaluation pipeline, and the public-facing site that surfaces the idea marketplace. Typically holds the lead-organizer role across both phases.

Operations Lead. Owns logistics, scheduling, venue, vendor coordination, and the day-of run-of-show. Carries the fiscal-sponsor relationship at the operations level and the cross-partner coordination that keeps the event running on the day.

Community Partnerships Lead. Recruits nonprofit and community partners, stands up the entrepreneurial support lounges, and holds the relationships that turn Phase 2 working sessions into rooms with the right people in them.

City Hall Liaison. Owns the City Hall side end-to-end: relationships across deputy CAOs and department heads, problem-statement co-authoring, the Day 30 meeting, the post-Day-30 product-innovation intake, and the department-head data contributions during judging week.

Fiscal Sponsor & Finance Lead. Holds the 501(c)(3) fiscal-sponsor relationship and runs the money side: sponsorship dollars in, prize money and event expenses out, and W-9s distributed to winners on site.

Volunteer Coordinator. Runs volunteer scheduling, on-site coverage, and the long tail of small moving pieces that no other role owns. Acts as the in-the-room utility player on event days.

Communications Lead. Owns outbound communications: announcements, recruitment, social, press, the Day 30 launch, and the post-event communications cadence.

Brand and Design Lead. Owns visual identity and collateral — web, print, signage, and slide decks — and the consistent look of everything the event puts in front of participants, sponsors, and the City.

Replication lesson: the lesson is not that another city needs the same eight people. It is that the same eight *functions* need named owners, and that gaps show up later as work the lead organizer ends up doing alone — which is the most common way a hackathon ends up as a one-shot event rather than the start of an annual one. Beyond these eight, name explicit owners for judging, participant communications, Slack and community operations, inbox triage, and participant onboarding / tech support.

Templates → `templates/03-stakeholder-map/`, `templates/11-mentor-volunteer-ops/`

Chapter 6

Replication Blueprint

Everything before this chapter describes what Richmond did. This chapter describes what to copy.

6.1 The Four Organizing Tenets

Four tenets break ties when the right call is not obvious. They are decision rules, not aspirational language — the choices in the earlier chapters fall out of them.

1. The team is guided by a government-owned strategic plan, and participants are guided by named problem statements that live downstream of that plan. The tie this resolved was a community group's late pitch to add an off-MAP problem statement during recruiting. The team declined and offered the group a partner role on an existing pillar instead, because statements live downstream of the strategic plan rather than alongside it.

2. The team provides pre-vetted research infrastructure friendly to non-technical teams, and routes teams to support when they hit technical walls. The tie this resolved was whether to ship raw datasets or pre-processed ones in the pillar repositories. The team shipped both, with the pre-processed version as the default and a documented schema, because non-technical teams should not have to debug raw data before they can begin working.

3. The team earns trust by facilitating effective outputs, and outputs must become valuable inputs into something else by design. The tie this resolved was whether the post-event surface should be a press release or a public marketplace of all 73 graded ideas. The team chose the marketplace, because a press release ends a conversation while a marketplace becomes input to the City's intake process and to the Phase 2 working sessions.

4. When the team hits a limit, it reiterates the mission, names the fear, asks for trust, and engages the network — because the team is well-connected enough to move the impediment. The tie this resolved was that the second-layer human review (city planners, cybersecurity experts, domain specialists) had no internal owner on the organizing team. The team named the staffing gap to the City Hall liaison, asked for trust, and recruited the panel through City Hall rather than absorbing the work inside the volunteer team; the same rule applies to asking for money, setting goals, and removing any specific impediment.

A replicating team should write its own tenets early — a small list of explicit decision rules on the wall before the hard calls arrive — not borrow these verbatim.

Templates → `templates/02-event-charter/`

6.2 The Replication Operating Sequence

The order Richmond built the model in. Not the only one that works.

1. **Secure a civic anchor inside an existing public strategy, before recruiting anyone else.** Tenet 1 above is operationalized here: find a senior person inside City Hall who can publicly endorse the event and put time on the Mayor's calendar, and anchor the event to an existing strategic plan, mayoral platform, comprehensive plan, or capital improvement plan. Richmond's anchor was already inside the Mayor's office, and the strategic plan was the Mayor's Action Plan. A city without that pre-existing relationship should plan for a longer engagement window. The combined step is one decision, not two: anchor and strategy travel together because each one legitimizes the other.
2. **Recruit a small cross-functional organizing team.** Cover: vision/narrative, event operations, City Hall coordination, partnerships, sponsorship/biz dev, finance, innovation facilitation, communications, marketing. Aim for 10+ years of experience per role. Be 100% explicit about commitment and expectations for each role at recruiting time — functional title, time commitment, what the role owns end-to-end, and what it does not. See §5.4 (Core roles and responsibilities) for the full role catalog.
3. **Build trust through informal convenings.** A VIP happy hour two to three months out, before any public announcement. Put City, nonprofit, and private-sector executives in one room and make a direct ask. This is the partnership strategy, not a social side-event.
4. **Convert civic priorities into pillar working groups.** Split the recruiting along existing relationships — one person owns the nonprofit/community side, another owns the City side. This makes recruiting fast.
5. **Host structured off-site problem-framing sessions.** Neutral creative space, not a government building. A facilitator whose job is to produce concrete outputs, not to summarize the conversation.
6. **Produce two practical problem statements plus blue-sky framing per pillar.** The practical statements fuel the weekend. The blue-sky prompts fuel post-event AI expansion.
7. **Validate problem statements with City and community stakeholders.** Informal review by City staff before the event confirms the research aligns with current realities. This is what makes it safe to treat the research as ground truth for participants.
8. **Build the operational infrastructure around the tracks.** Pillar Slack channels, pillar GitHub repos, sponsor and mentor assignments, post-event room structure. The pillars are the spine; everything else hangs off them.
9. **Stand up the minimum viable comms stack before recruiting at scale.** Public website, registration (route through a fiscal sponsor's existing infrastructure if possible), Slack provisioned from a single config file, broadcast email backup, role-based email aliases.
10. **Run the event.** The weekend is one of the *shorter* commitments in this model. Most of the work is before and after.
11. **Run a structured post-event evaluation.** Implementation Reality rubric on the full corpus. Produce a review-ready package the City can act on, not a second set of winners.
12. **Give City stakeholders a curated review path.** City Champion Path for the ten or so ideas the City claims; Civic Working-Session Path for the rest.

A replicating city that skips steps 1, 3, 5, 11, or 12 should expect the post-event drift that plagues most civic hackathons.

This is the sequence Richmond followed, not the only working sequence and not necessarily the perfect one. A future iteration of this toolkit could codify the load-bearing decisions inside each step as a matrix of architectural decision records (ADRs), so that each step carries forward not just *what* to do but *why* the alternative was rejected.

Templates → `templates/00-plugin-orientation/`

6.3 Communications Infrastructure

Comms is event infrastructure; stand it up *before* recruiting at scale.

Slack as the Operating System

One workspace covers participants, mentors, judges, sponsors, volunteers, and organizers. Run on Slack's **one-month free Business Plan trial** timed to overlap event week — the trial is the unlock for API access, which is what makes the team-formation, mentor-help, submission-intake, and post-join routing workflows possible; after it ends the workspace reverts to free-tier limits and the workflows persist.

Implementation lives in `infra/hackathon-slack/`. The replicable pieces:

Single-YAML configuration Channel names, topics, role mappings, post-join routing, and archive policy in one file the next city can fork.

Channel taxonomy Five categories: public defaults (`#announcements`, `#help`, `#find-a-team`, `#introductions`, `#volunteer-chat`); seven pillar channels; private role channels; generated team channels; sponsor/partner channels.

Post-join routing A bot assigns each new member to the right channels from their role record.

Mentor help workflow Two emoji as a lightweight ticket: eyes to claim, checkmark to resolve.

Team formation intake Slack workflow form → three-pass hacker/hustler/artist matching → team CSV that the provisioner consumes.

Submissions Slack workflow form → Google Sheet.

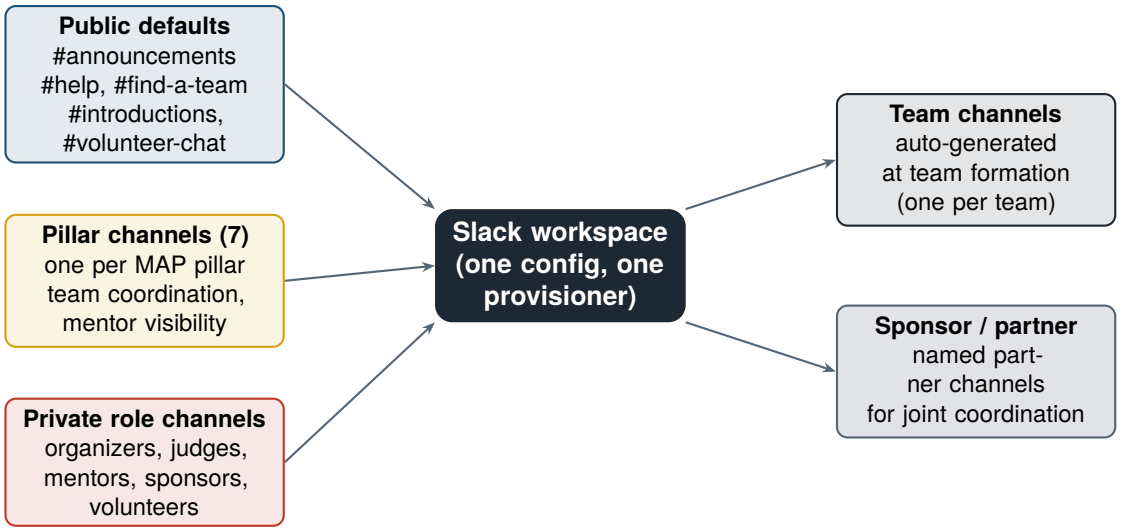


Figure 6.1: The five Slack channel categories. Public defaults broadcast and triage; pillar channels coordinate on civic-priority lines; private role channels separate operating concerns; team channels are auto-generated at team formation; sponsor and partner channels exist where named outside-organization coordination is needed.

Slack-First Onboarding

Registration confirmation routes every participant directly into Slack as one of the first onboarding steps; pre-event updates land in Slack first, with email as backup.

Role-Based Email Aliases

Simple forwarding aliases for stable destinations regardless of which organizer is most active that week: `judging@`, `sponsors@`, `participants@`, `mentors@`, `volunteers@`, `city@`, `press@`, `ops@`.

Minimum Viable Comms Stack

- Public website (any quick-build tool)
- Registration owned by one named person
- Slack workspace from a single config file
- Broadcast email (backup only)
- Role-based email aliases (the 8 above)
- Authoritative announcement channel
- Team-formation, mentor/help, and submission workflows (above)
- Judging criteria page linked from Slack, email, and website
- Post-event follow-up channel kept open for Day 30/90 and the next event

Templates → `templates/12-event-operations-plan/`, `templates/09-community-outreach-plan/`

Closing

The Richmond Civic Hackathon was an attempt to build a model that does not stop at the demo. The first year produced 38 human demos plus 35 AI-generated ideas, a complete 73-idea corpus that the City has begun reviewing, and a Day-30 Phase 2 launch with named City stakeholders. What it did *not* produce, as of April 2026, was a single named institutional champion across the fourteen problem statements. The infrastructure exists; the action it is meant to enable has not happened yet.

A replicating city should expect the same shape. The event will produce demos, the post-event evaluation will produce a ranked corpus, and the City Champion Path will produce a small set of ideas the City has agreed to evaluate. Whether any of them survive the implementation timeline is decided over quarters, not weeks, by named people doing unglamorous work in specific institutions.

The hackathon is the first commit, not the merge. The champion network, the evaluated idea corpus, and the partnerships that form in the years that follow are what determine whether any of this matters.

Appendix A

Source Research

The narrative claims in this toolkit are grounded in roughly 6,000 lines of cross-cutting research on how the Richmond Civic Hackathon was designed, run, and evaluated. The full index, organized by category, lives in `RESEARCH_INDEX.md` at the root of the `rvahacks` source repository. The table below maps each file to the chapter it most informed.

Chapter	Research file (path relative to repo root)	Category
Ch. 1	`2026-annual-hackathon/research/civic-innovation-value-chain-t`	Framework
Ch. 1	`2026-annual-hackathon/research/RVA_Civic_Innovation_Framework`	Framework
Ch. 1	`rvahacks/knowledge-base/01-planning-strategy/background-resea`	Strategy
Ch. 1	`rvahacks/knowledge-base/13-resources/metrics-supporting-hackat`	Strategy
Ch. 2	`2026-annual-hackathon/research/cats-institutional-framework-a`	Framework
Ch. 2	`rvahacks/knowledge-base/03-finance/richmond_civic_hackathon_sp`	Sponsorship
Ch. 2	`rvahacks/knowledge-base/03-finance/virginia_civic_hackathon_gr`	Sponsorship
Ch. 2	`rvahacks/knowledge-base/03-finance/hackathon_sponsorship_resea`	Sponsorship
Ch. 2	`rvahacks/knowledge-base/06-challenge-design/map-alignment.md`	Strategy
Ch. 2	`rvahacks/knowledge-base/02-team-governance/pillar-ecosystem.m`	Strategy
Ch. 2	`rvahacks/knowledge-base/02-team-governance/academia-k12.md`	Strategy
Ch. 3	`2026-annual-hackathon/research/cats-technology-architecture-a`	Framework
Ch. 4	`2026-annual-hackathon/ideas/IMPLEMENTATION_REALITY_RUBRIC.md`	Synthesis
Ch. 4	`2026-annual-hackathon/ideas/IMPLEMENTATION_REALITY_SYNTHESIS.`	Synthesis
Ch. 4	`2026-annual-hackathon/ideas/RICHMOND_TECH_STACK_2026.md`	Synthesis
Ch. 4	`2026-annual-hackathon/ideas/RESEARCH_QUERIES.md`	Synthesis
Ch. 4	`2026-annual-hackathon/research/parallel-ai-research-quickstart`	AI tooling
Ch. 4	`2026-annual-hackathon/agents/CIVIC_EVAL_PERSONAS.md`	AI tooling
Ch. 5	`2026-annual-hackathon/research/hackathon-to-public-value-richr`	Framework
All	`REPLICATION_QUESTIONS.md`	Source material

Table A.1: Source-research files mapped to the toolkit chapters they most informed. Full descriptions in `RESEARCH_INDEX.md`. Per-pillar implementation-reality reviews (one per pillar) feed `IMPLEMENTATION_REALITY_SYNTHESIS.md` and are indexed there.

Document: Richmond Civic Hackathon Replication Toolkit · Version 1.0 · May 2026
Prepared by: Flying Pig Labs / hack4rva
Companion plugin: github.com/hack4rva/annual-civic-hackathon-toolkit — fill-in templates and example artifacts that operationalize each chapter for another city.
Source research: See Appendix A and `RESEARCH_INDEX.md` in the source repository.

Source: github.com/hack4rva