


Pop-up City Report.



A field report on temporary, intentional cities: what they are, where they produce value, and where they predictably fail

Contents.

Executive Summary	3
-------------------	---

Part I · Orientation	
----------------------	--

Ch .01	14
What pop-up cities are	

Ch .02	16
Who this report is for	

Ch .03	18
The landscape	

Ch .04	22
How pop-up cities compare to adjacent forms	

Ch .05	26
Key findings	

Part II · Mechanics	
---------------------	--

Ch .06	35
Membership and selection	

Ch .07	39
Governance and legitimacy	

Ch .08	43
Incentives and contribution	

Ch .09	48
Enforcement and operations	

Ch .10	52
Economics	

Part III · Outcomes	
---------------------	--

Ch .11	58
Outcome measurement	

Ch .12	61
Three ways pop-up cities fail	

Part IV · Practice	
--------------------	--

Ch .13	65
Design choices that matter most	

Ch .14	71
Minimum viable institutions	

Appendices	
------------	--

App .A	76
Focal case profiles	

App .B	89
Metrics dictionary	

App .C	91
Failure mode reference	

App .D	92
Terms used in this report	

App .E	93
Methodology	

App .F	98
Case directory	

Executive summary.

A pop-up city is a temporary physical gathering in which people with shared interests or values build a short-lived, city-like environment: shared spaces, social norms, cultural activities, services, and coordination structures. In practice the form is typically a curated cohort of hundreds to low thousands of people, together for 2 to 12 weeks. 107 have been catalogued as of May 21, 2026, and around 50 self-describe ambitions beyond hospitality: governance experiments, research environments, prototypes for new kinds of institutions. Since the 2023 founding case the format has scaled ahead of the diligence frameworks funders can consult, the playbooks operators can share, and the protocols host jurisdictions can apply. This report addresses that gap. A pop-up city compresses membership, governance, safety, legitimacy, economic design, and host relations into a cycle short enough to observe whole.

The evidence base

Interview and survey corpora closed April 30, 2026; desk sources coded through July 2026. Case directory updated May 21, 2026.

Unique coded interviews	92
Valid survey responses (analytic corpus)	86
Coded desk research sources	337
Theme memos (one per hypothesis)	22
Communities and organisations represented through interviews	54
Catalogued pop-up city cases	107

Primary corpus. 20 cases studied closely (12 consented focal cases and 8 comparators).

Hypothesis register. 22 numbered hypotheses (H1 to H22); 24 register entries, two of them split into halves that earned separate verdicts. Across the 24 entries: 11 supported, 11 qualified, 1 contested, and 1, on local integration, under assessment.¹

Research period. November 2025 to July 2026.

Self-identification, not endorsement: of the 107 catalogued cases, 102 meet the report’s working definition and 5 are boundary cases carried as scale or edge benchmarks. Full methodology disclosures in Appendix E.

Reading conventions

How to read claims. This report makes evidence-weighted claims, not population-wide statistical claims. “The field” means the documented ecosystem visible through the sources studied. Prevalence language is directional unless a direct count is given.

Confidence markers.

-  **High** – 3+ independent sources from 2+ source streams.
-  **Medium** – 2 independent sources, or proxy measurement, or 40 to 60% case coverage.
-  **Low** – 1 source or inference-heavy; retained where central to a finding and flagged.

Hypothesis stamps. Chapter openers carry a stamp showing the chapter’s central hypotheses and their register status, where a chapter tests specific hypotheses; the register is set out in full in Appendix E.

Anonymisation. Persons and organisations appear under coded identifiers or role-based language. Authors of published works cited as desk research are named normally. Appendix A profiles and the Appendix F directory use real names by consent and by design. All amounts are stated in USD.

01

What a pop-up city is, and what it is not

Pop-up cities sit in a crowded neighbourhood of forms: conferences, retreats, co-living operators, intentional communities, charter cities and special economic zones. The working definition separates them by four conditions that must hold at once: daily co-presence for the full duration, a fixed start and a fixed end, selection-generated trust from an explicit entry filter, and a shared intention to work on something substantive together. Each adjacent form fails a different condition, and the boundary matters practically: the risks a host jurisdiction should price, the economics a funder should expect, and the design decisions an operator faces all change at these edges.

02

The landscape

Across the 107 catalogued cases and the 20 studied closely, what holds constant is a shared design space, not a shared model. Scale runs from roughly 30 participants to tens of thousands at the benchmark comparator; durations cluster at 2 to 12 weeks, and 15 of the 20 primary-corpus cases ran shorter than 8 weeks. The report sorts the space with a four-type lens (Sprint, Launchpad, Season, Proto-jurisdiction), ordered by duration and institutional ambition; the lens is a simplification, and the honest number of recurring configurations is probably higher. Longevity is rare: 2 parent entities report 5+ years of operation without external-grant or per-event-sponsorship dependence, and both run small permanent spaces in place of flagship events.

03

Who this report is for

Each audience arrives with a different question. Operators and organisers: which design decisions matter most, and how are the predictable failures avoided? Investors and funders: how are these ventures evaluated, and what separates survivors from failures? Host jurisdictions and institutions: does hosting create value, and on what terms?




Exhibit 1

The form and its neighbours

Four conditions must hold at once. Each adjacent form fails a different one, which is what separates a pop-up city from its neighbours.



A central ‘pop-up city’ region carries the four conditions (co-presence; time-bounded intensity; selection-generated trust; shared intention), ringed by five adjacent forms, each labelled with the one condition it fails.

Working definition; Ch.1 and 4. Directory n=107 (self-identification, not endorsement).



04

The strongest findings

Seven findings carry the strongest convergence across independent sources. Ch.5 walks each one.

1



Legitimacy rests on a record no one keeps.

Participants consistently value trust formation, ambient collaboration,² serendipity and the long arc of work begun in-event; operators track attendance and budgets. The field claims legitimacy on competence and track record, yet the track record is unrecorded, so no participant, funder or host can audit it. Funder fatigue in the cases studied traces to the same gap.

2



Patron capital underwrites the form as practised.

Most models consume more than they generate, and mission-aligned patron capital typically closes the difference; in the mid-scale cases studied the backer is often a single committed funder, though more than one funding lineage now operates in the space and several events run on other models entirely (premium pricing, earned revenue, grants). Where patron money arrives without structure it also concentrates governance and leaves continuity resting on individual funders' decisions, so the operative levers are funding structure, obligations and allocation rather than volume, and structured models, milestone-based foundation funding among them, are already appearing.

3



Founder-led operation is the default; the experiments are bounded.

Founder authority persists at every observed scale and duration, participatory elements survive only where nothing expensive depends on them, and every documented transition runs toward centralised operation, none the other way.

4



The trust machine and the host problem are the same machine.

The filters that generate selection-based trust are exactly what the host community fails. Exclusion is the selection model's default output, and the counter-cases in the evidence all redesigned the operating model itself.

The strongest findings

5

Knowledge travels with people; tools get orphaned.

Institutional innovations transfer where personnel move and a receiving structure exists; no artifact studied has a documented cross-community adoption record, which places the bottleneck in maintenance incentives rather than idea scarcity.

6

The capacity floor is high, the ceiling is real, and both are priced at current technology.

Minimum viable institutional capacity, especially for safety and host relations, exceeds most founding teams' plans. Coordination costs rise superlinearly under participatory governance; the ceiling appears to sit somewhere in the low hundreds of participants (a working hypothesis anchored by a single comparison, not a benchmark). Both figures are priced at current coordination technology; whether AI tooling moves them is an open question.

7

The ecosystem is pre-tipping-point, partly by its own selectivity.

Publicly available institutional funding dedicated to the space is limited, the talent pipeline recycles a small pool of organisers, and demand is narrower than the vision implies. The network-state framing attracts mission-aligned capital and tool-builders while repelling the institutional investors and professional operators whose resources compound; the observed fixes are reframings rather than fundraising pushes.



05

How pop-up cities fail

Three failure modes stack by level in the cases studied.

Event level

Operational overwhelm. Small founding teams cannot sustainably cover the operational surface the format demands, and the most productive contributors burn out first.

Organisation level

Institutional knowledge loss. Playbooks go unwritten, retrospectives are skipped, and organisers turn over between events.

Ecosystem level


Subsidy dependency. Continued reliance on patron funding in a field still ahead of its tipping point.

The connecting thread is Finding 5: the information that would prevent each failure usually already exists in the ecosystem and fails to be captured, maintained, or handed over. Leading indicators are recognisable in advance, though the evidence does not support a specific lead time.

Exhibit 2

The reach of patron money

The same patron money makes events possible, concentrates governance, and leaves continuity resting on individual funders' decisions – three effects at once.



An actor node 'Patron' feeds a central 'Event / community' node through three simultaneous effects: a solid money arrow (subsidy → possibility) and two dashed arrows (allocation authority → concentration; funding continuity dependence). Below, four attempted exits – none demonstrated at flagship scale.

Interview and survey corpus, informant-deduplicated; Ch.8 and 10.



What this means for each audience

For operators and organisers

Ch.13's 18 design decisions are worth walking through before the first event, and naming an owner for each of the seven institutional functions in Ch.14 reduces the chance that a function falls to whoever is available in a crisis. The cheapest legitimacy intervention available is record-keeping: published cost breakdowns, retention tracking, outcome attribution.

For investors and funders

In the cases studied, three leverage points stand out: pre-committed capital pools; talent-retention infrastructure, so that experienced organisers can professionalise instead of burning out; and narrative artefacts clear enough for mainstream institutions to evaluate. Flagship-scale events tend to remain dependent on continued funding, and without a clear transition mechanism, funding one is closer to ongoing patronage than to investment. Structure matters more than volume: subsidy with attached obligations mitigates the distortions unstructured giving creates. One practical test: whether the operator tracks the Tier 0 metrics. Their absence from a pitch is itself diligence information.

For host jurisdictions and institutions

Direct economic impact may be modest, so the case for hosting rests elsewhere. The contributions that tend to matter (new organisations, new policy frameworks, positions in global knowledge networks) operate on a longer timeline, and they arrive only with genuine local integration. Logistics-only hosting is workable but raises the chances of regulatory mismatch, continuity gaps across electoral cycles, and friction over how value flows to local constituents. What separates the risky cases is structure, not vocabulary: the documented host-rejection cases sit almost entirely in adjacent forms carrying design-time markers (land tenure beyond ordinary lease, legal derogation, displacement of residents' rights) that pop-up cities as studied here mostly lack. The practical implications: lower hosting barriers, local-integration infrastructure, multi-year evaluation horizons.


Institutional engagement at agency level is already in the record

In November 2025, alongside Forma's pop-up village in Tbilisi, Georgia's Innovation and Technology Agency (GITA) signed a memorandum of understanding with the Solana Foundation, with both sides framing the arrangement in development-economics terms. The MoU's contents are not public; the significance is the fact and level of the partnership. One development from June 2026, outside the analysed evidence, points the same way: Gelephu Mindfulness City, a sovereign-initiated new-city project conceived by Bhutan as a national development instrument, is not a pop-up city, but it suggests the institutional interface is forming faster than the field's own tooling for it.

Exhibit 3

Using the report by decision, not by chapter

Three decision-keyed swimlanes route each audience by the question they face, not by reading order.



Report structure; decision framing from Ch.13, 8 and 10.

07

Why pop-up cities matter now

Pop-up cities are still the only setting where membership rules, governance mechanisms, economic design and host relations are being varied and observed on cycles of weeks rather than decades: the whole institutional arc fits inside one observation window.

The practical value is timing: the field is pre-tipping-point, so the design choices, measurement habits and host-relations norms being set now are set before lock-in, while they are still cheap to change.

08

What this report does not claim

The report does not claim pop-up cities work; it documents where they produce value and where they predictably fail. It does not claim the coordination ceiling is permanent, that artifact-mediated transfer functions, or that host rejection cannot occur (one reported pop-up case is noted at A.2; others sit in adjacent forms). It does not quantify economic impact; the available figures are directional, and several rest on single sources disclosed as such.

It does not yet carry the host community's own voice at evidential strength, and the survey base of 86 valid responses against a 200 to 300 target is a disclosed shortfall, not a foundation for population claims. The legal and regulatory landscape, festival cities, mass religious gatherings, refugee and disaster settlements, commercial brand activations, and the environmental sustainability of build and teardown are out of scope.

Notes

1. The count includes one hypothesis whose upgrade to supported was certified in July 2026 after primary-source verification of its two comparator-level anchors.
2. "Ambient collaboration" is the report's editorial term for the unscheduled productive overlap that selected, co-present participants describe; it is not a phrase drawn from any interview.

The findings are walked in Ch.5; design decisions in Ch.13; the institutional checklist in Ch.14. Full methodology in Appendix E.

Orientation.

What pop-up cities are, where they sit among adjacent forms, and who this report is for.

Ch . 01	14
What pop-up cities are	

Ch . 02	16
Who this report is for	

Ch . 03	18
The landscape	

Ch . 04	22
How pop-up cities compare to adjacent forms	

Ch . 05	26
Key findings	

What pop-up cities are.

A pop-up city is a temporary physical gathering in which a selected group builds a short-lived, city-like environment, and four conditions define it at once: daily co-presence, a fixed start and a fixed end, selection-generated trust from an explicit entry filter, and a shared intention to build something substantive together. Exhibit 1.1 scores the neighbouring forms condition by condition.

The conditions are observational rather than stipulated: across the 20 cases studied closely, practitioners converge on similar designs without a shared written definition. The label travels further than the form; the case directory records self-identification, not endorsement, and Exhibit 1.1 shows where two catalogued label-users fall short, and the name itself predates the form.¹

One boundary is already moving: some operators now run small permanent spaces alongside or instead of events, and duration is starting to behave as a spectrum rather than a switch. The time-bounded condition holds for the period studied; a future edition may need to treat temporariness as a dimension.

Ch.3 maps the space the definition encloses; Ch.4 runs the conditions as the operational test against the institutional forms readers already know.

Out of scope

Festival cities (programmed consumption, not shared work); mass religious gatherings (devotional, unselected); refugee and disaster settlements (involuntary); commercial brand activations (no cohort); tactical urbanism interventions (no residential cohort or shared intention).

Notes

1. On the term's lineage: Wikipedia, [Pop-up retail](#); [Tactical urbanism](#).

Exhibit 1.1. Comparison matrix

The four conditions, form by form

Each form scored against the four defining conditions. Only the pop-up city carries all four at once; the boundary rows are catalogued cases that use the label without meeting the test: a sponsor-funded event operating as a large conference with community features, and a permanent co-working building marketed as a pop-up environment. Filled square = condition present, half-filled = partial, empty = absent; glyphs score conditions, not confidence.


	Co-presence	Time-bounded	Selection-trust	Shared intention
Pop-up city				
Boundary: sponsor-funded mega-conference				
Boundary: permanent co-working venue				
Conference				
Co-living operator				
Special economic zone				
Intentional community				
Co-working space				
Retreat				

Source: working definition applied per Ch.4; boundary cases from the case directory (self-identification, not endorsement).



Who this report is for.

Three audiences read the same evidence and need different things from it, so this chapter routes each to what it is deciding and where the material sits.



Operators and organisers

Operators face dozens of design decisions under time pressure with a small team, and in the cases studied the decisions that matter most are economic and governance-mode, often made by default: how entry is priced and provisioned (free entry with full provision quietly drifts participants into consumer mode), and whether participatory governance is designed in from the start (where it is not, authority centralises once anything expensive is at stake). Ch.13 walks the 18 design decisions; Ch.8 the incentives; Ch.14 the minimum measurement set.

Investors and funders

The public diligence frameworks do not yet cover the category well, so the baseline to price in is Ch.10's: no flagship or large-format case studied is financially self-sustaining absent some external cushion, with small permanent spaces the scale-capped partial exception. Ch.10 carries the economics; Ch.12 the failure analysis.


Host governments and institutions

The operator-side evidence points to institutional rather than direct economic value, on a longer timeline than a single event, arriving more reliably where hosting is co-produced with local partners rather than logistics-only. Ch.10 carries the host-side evidence; Ch.13 (Exhibit 13.1) the host-facing design decisions.

Exhibit 2.1. Audience pathway map

Audience reading pathway


Reading order by audience. All three routes share the same spine: the executive summary, the key findings, and the methodology appendix.



Source: report structure.

The landscape.

Four shapes keep appearing when the field is sorted by length and ambition.




The field is wider than any single model: as of the May 21, 2026 directory update, 107 pop-up cities are catalogued, of which 20 were studied closely. Scale runs from roughly 30 participants to the benchmark comparator, Burning Man, at tens of thousands; setting the benchmark aside, cases run from roughly 30 to 5,000 participants, most over 2 to 12 weeks, with one year-long residency at the long end. Two dimensions organise the cases studied better than geography, theme or funding source: duration and institutional ambition, which together track governance complexity, cost structure, membership architecture and failure mode.

Exhibit 3.1 sorts the space on those axes into four recurring configurations, Sprint, Launchpad, Season and Proto-jurisdiction, and two honesty notes attach. The lens classifies configurations rather than organisations: a single parent can run a short flagship in one geography and a longer residency in another in the same year, and cases drift across types as capacity accumulates. And four is a reading aid, not a measured result: the evidence keeps surfacing candidates the lens does not separate, and a systematic count would probably land higher.¹

Exhibit 3.1. 2x2 matrix

The landscape, by duration and ambition

Four recurring configurations the cases studied keep returning to: a reading aid, not a taxonomy.



Source: case database, May 21, 2026 (n=107; 102 mapped + 5 boundary). Dots show the 12 consented focal cases (Appendix A) placed by archetype and duration.

Exhibit 3.2. Comparison matrix

The four configurations compared

Qualitative signature of each configuration. Duration and scale live on the Exhibit 3.1 map, not here.

	Sprint short, low ambition	Launchpad short, high ambition	Season long, low-to-mid ambition	Proto-jurisdiction long, high ambition
Purpose	Social experiment; network seeding; cultural production	Curated output: projects, entities, governance pilots	Community formation through extended co-presence	Governance system design and testing; institution-building
Membership	Open application; low curation acceptable	Application plus curation; intentional composition	Values-aligned; open or invite	Highly curated; founder plus application; high filter
Governance	Founder unilateral or core-team	Core-team with some delegation	Founder-led or emergent	Delegated, participatory or sortition; formal governance is the product
Revenue	Tickets and sponsorship	Fees and grants; sponsorship less common	Owner equity or sponsorship, by variant	Subscriptions plus land or property equity plus grants
Primary output	Networks and cultural moments	Intellectual work, proofs of concept, new entities	Persistent community; cultural production	Governance protocols, legal entities, replicable institutional infrastructure
Signature risk	No persistence mechanism beyond informal networks	Institutional ambitions outrun the time and trust available	Coordination costs creep upward without institutional scaffolding	Claims authority without mandate or enforcement capacity

Configurations describe events, not parent organisations; a single parent can run more than one configuration concurrently.

Source: primary-corpus case analysis (n=20). Qualitative fields only; numeric bands are not printed.



What the landscape shows

Most activity clusters at the short end: 15 of the 20 cases studied ran in episodes shorter than 8 weeks, so most of the field's operating experience sits in the two short-duration quadrants. The types also behave as stages: the common path starts short and low-ambition and drifts toward longer duration and higher ambition as capacity accumulates, and cases that invert the path, launching with high institutional ambition before operational foundations exist, appear failure-prone.

Durability, where it exists, lives at parent level: 2 parent entities report 5+ years of operation without external-grant or per-event-sponsorship dependence, and both anchor short events to a small permanent venue, self-funded through owner equity and membership fees (directional rather than a benchmark).

Composition varies less than structure: across all four types, participants are predominantly tech-focused, young and Western-leaning in cultural orientation, and the host community sits largely outside the integration the form naturally produces.² The demand-side consequence is Finding 7 in Ch.5; the host-side machinery is Finding 4 and Ch.6.

Ch.4 sets this map against the institutional forms readers already know.

Notes

1. Per the hypothesis register's archetype entry, the count is to be derived by dimensional scoring rather than asserted; one desk comparator reaches 15 types, counted on organisations rather than configurations.

2. Convergent across informant positions, spanning organiser, participant and local-host vantages. The observer base is itself demographically narrow, and the single dedicated local-host account carries a disclosed conflict of interest (Appendix E); age-band and share estimates from that account are single-source and not printed.

How pop-up cities compare to adjacent forms.

Fellowships and festivals come close; none passes the whole test.

Ch.1's four conditions carry the comparative test, and Exhibit 1.1 scores the neighbouring forms; this chapter adds what the grid cannot hold. The nearest relatives are residential fellowships and accelerator cohorts, which satisfy all four conditions in narrowed form: a fellowship is an institution hosting a cohort, where a pop-up city is a cohort briefly acting as an institution; an accelerator points the shared intention at commercial product development rather than institutional experimentation. Special economic zones and charter cities sit furthest away in operation yet closest in ambition, which is why they dominate the comparator record below.

Hypotheses

- H16a **Qualified**
- H16b **Supported**
- H20 **Supported**

01

The intellectual landscape

The field has no single canonical literature: seven positions recur across its published self-description and the cases studied, operators most often invoke the d/acc and network-state lineages, and none of the seven treats the form exhaustively; all are objects of analysis rather than commitments of the report. Exhibit 4.1 orders them by what temporariness is for: the gathering as the point at one end, a rehearsal for permanent arrangements at the other.

The same case reads differently by position: one funding pattern can be healthy R&D infrastructure under d/acc, incomplete recruitment under the network-state frame, and extractive rather than overlay under the commons frame. Ch.12 §12.3 carries a counter-view from within the field: Buterin's December 2025 diagnosis of pop-ups becoming shorter, smaller and less experimental.¹

Exhibit 4.1. Comparison matrix

Seven positions on what pop-up cities are for

Ordered top-to-bottom by what temporariness is for: the gathering as the point, through to a rehearsal for permanent arrangements. Positions describe ideological centre of mass; individuals, cases and funders may sit in several at once.

Position	Set out in (key works)	What pop-up cities are for
d/acc	Buterin, My Techno-Optimism (2023); d/acc: one year later (2025)	Field sites for governance and defensive-technology R&D, with democratic governance preserved
Plurality	Tang and Weyl, Plurality (2024); RadicalxChange essays	Testbeds for plural-voting, quadratic-funding and liquid-democracy mechanisms
Network state	Srinivasan, The Network State (2022)	Recruitment and cohort formation toward sovereignty recognition
Developmentalist / charter cities	Romer, charter-cities papers; Charter Cities Institute; Próspera writings	Rules-based zones for economic development, often co-designed with a host state
Commons governance / Coordi-nations	di Filippi, Coordi-nations; Bauwens, Network Sovereignties in the Context of Macrohistorical Patterns (2024); Foster and Iaione, Ostrom in the City (2017)	Temporary overlays on state functions and commons prototypes; surplus reinvested rather than captured
Libertarian / exit	Davidson and Rees-Mogg, The Sovereign Individual (1997); Friedman, seasteading literature; Free Cities Foundation	Small-footprint demonstrations of alternative rule-sets; trade scale for autonomy
Prefigurative community	Bey, T.A.Z. (1991); Burning Man Project, “Ten Principles” (Harvey, 2004); Chen, Enabling Creative Chaos (2009)	Live the alternative for a bounded period; practice teaches what permanent versions could be

Source: published works as cited and the field's self-description in the cases studied (n=7 positions). Prefigurative row rests on desk references, not coded interview sources.



02

Doctrine and practice in the network-state lineage

The field's most influential doctrine decomposes informatively on contact with operations. The Network State (2022) prescribes two sequences: a geographic one (start decentralised and online, grow many nodes, acquire territory later) and a governance one (founder authority from the outset).² In practice the lineage's most prominent operations run as single founder-led residential hubs, with satellite nodes planned to follow the hub, not precede it.³ Read by axis, the geographic doctrine reversed on contact with operations, while the governance doctrine had nothing to reverse. That distinction says more than 'the plan failed' or 'the plan is on course': the practised model converges on a residential hub, a form the Ch.3 map already covers, and the governance continuity anchors from the doctrine side what Ch.7 documents from operations.

03

An open question in the d/acc lineage

Buterin's 2025 revision makes democratic governance a binding criterion of d/acc, alongside the defensive and decentralised criteria it already carried.⁴ Across d/acc-associated venues in the cases studied, governance at pop-up scale runs through a small operational core with limited structured resident voice: the venues have not yet democratised, and the published commitment has not yet specified what democratic governance at this scale and duration would look like.

04

What the comparator record shows hosts


Nearly every documented host rejection or clawback event in the evidence sits in adjacent forms carrying design-time structural markers; one pop-up city proper now has a documented rejection, a reported shutdown order (A.2); Exhibit 4.2 renders the test.⁵ One hedge belongs here: the near-absence of rejection is not host approval, since no host official has been interviewed anywhere in the evidence base and a host that quietly declined a permit would be invisible to it.

The four-condition test remains the first-pass filter for the label; the seven positions are lenses, and the chapters that follow read differently depending on which one a reader brings.

Exhibit 4.2. Decision tree

Structure, not vocabulary, carries the host risk

The ex-ante test separating the documented host-rejection record from the pop-up record. Framing and vocabulary sit outside the test, carrying reputational rather than structural risk.



Source: hypothesis-register host-value criterion and the comparator desk record. Key comparator cases primary-verified July 2026.

Notes

- 1. Buterin, Let a Thousand Societies Bloom (December 17, 2025).
- 2. Srinivasan, The Network State (2022); on governance: "A state, like a company, needs a leader. Especially early on."
- 3. Organiser accounts and the public case record (Appendix A); single-lineage evidence centred on one operator core, disclosed as such.
- 4. Buterin, My Techno-Optimism (2023); d/acc: one year later (2025).
- 5. Desk-research record of adjacent-form cases, two conceded by advocates of the models concerned; two underlying documents await primary verification, so the cases are described generically. The repeal's outcome was dampening, not termination. Surveyed operators report no adversarial host interactions; the figures are withheld pending verification and not printed.

Key findings.


Seven findings carry the strongest convergence across independent sources.

The research tested 22 working hypotheses against the evidence base, and after the July 2026 re-adjudication the register carries 24 entries, two hypotheses having split into paired halves with separate verdicts: 11 are supported, 11 qualified, 1 contested, and 1, on local integration, remains under assessment.¹ The seven findings below are drawn from that register; read together, they connect into one system, and Exhibit 5.1 draws the dependencies.

Exhibit 5.1. Flow diagram

How the findings connect

The findings are not a list: each documents a dependency on the others. Arrows read as “produces / enables”; F7 is where the chain lands.



Source: synthesis of the hypothesis register; informant-deduplicated bases per finding. Confidence marks per finding.



Finding 01

Legitimacy rests on a record no one keeps

The outcomes participants value most are consistently the outcomes operators do not measure: participants report trust formation, ambient collaboration, serendipity and the long arc of work begun in-event; operators track attendance and budgets. The instruments exist and go unused; Ch.11 carries the record, including a funder-built benefit-cost instrument that no operator adopted, while retention stays trivially trackable and mostly untracked.² The gap sits in deployment, not in measurement science.

Legitimacy is claimed on competence and track record, and not on participatory process, by the operators' own account, but a track record that is never recorded is a narrative, and a narrative can be punctured faster than a record. The same gap reappears on the funding side as donor fatigue: subsidy flows without accountability structure, and when a funder eventually asks what the money did and cannot find out, the capital exits. The evidence contains one end-to-end worked example, a field-level funding mechanism that ended when the ecosystem's anchor patron declined a co-funding match in writing.³

Record-keeping, before any governance reform, is the most direct legitimacy intervention available to operators, and an unusually cheap one (published cost breakdowns, retention tracking, outcome attribution). The report cannot show unrecorded competence claims are false, only that nobody can audit them.

Finding 02

Patron capital underwrites the form as practised

The evidence supports a stronger reading of patron capital than a financing gap waiting to be closed: mission-aligned patron capital is what makes the form as currently practised possible. Most models consume more than they generate, premium pricing erodes the contribution culture the form needs, and in the mid-scale cases studied the party absorbing the difference is typically a single committed backer, though the funding base is broader than one source: more than one patron lineage now operates in the space, and several events run on other models entirely (premium pricing, earned revenue, grants).

Where it arrives without structure, the same money does three things at once, which is why fundraising fixes miss the point: it makes events possible; it concentrates governance, because personal financial capital settles allocation questions no participatory mechanism is ever asked to decide; and it leaves continuity resting on individual funders' decisions, because unstructured patron money carries no accountability obligations and does not compound into institutional capacity. The operative levers are funding structure, obligations and allocation design, not volume, and structured models, milestone-based foundation funding among them, are already appearing in the record.⁴

The pattern extends into participant behaviour: where entry is free and provision is full, contributors drift into consumer mode; one organiser at a zero-ticket event predicted the drift and then watched it arrive; Ch.6 sets out the three factors carrying it (price, provision, payment direction). One small zero-fee, no-provision case produced a producer culture, so free entry appears to be one input, not a sufficient cause.⁵

Two boundaries belong in the headline. The largest case studied is a comparator running predominantly on earned revenue, so the pattern describes the pop-up-proper mid-scale cohort rather than the largest-format flagships; and the four attempted exits from patron dependence (decommodified models, incubator-equity, endowment designs, self-funded small permanent spaces) are genuine design work, none yet demonstrated at flagship scale. The finding is about what the money constitutes, not whether it was generous; several interviewees credit it with the field's existence and quality.

Finding 03

Founder-led operation is the default; the experiments are bounded

The field describes itself as a governance R&D environment, yet in operation it runs on founder authority at every observed scale and duration, including the largest permanent node in the cases studied. The tension reads as misdescription rather than hypocrisy: the R&D framing is real, but what is mostly being tested is variations of founder authority.

The mechanism is liability, and Ch.7 carries the mechanics: decisions with real stakes migrate to whoever can be held responsible, in every documented transition toward centralised operation, never back. The founding doctrine itself prescribes founder authority from day one.⁶ On-chain governance underperforms where scoped to collective decisions and holds up where scoped to narrow instruments; Ch.7 carries that record.⁷

Designed participation works where it is architected before dysfunction appears (sortition designs, co-creation architectures, vouching with public-ledger accountability, trust-based governance at village scale), and its absence surfaces as concentrated authority with costs, not breakdown. The practical reading for operators, argued in Ch.7: pick the governance mode on purpose and tell people, since a disclosed hierarchy carries costs a community can price, while one that contradicts the stated design erodes the legitimacy Finding 1 shows the field cannot afford to spend.


Finding 04

The trust machine and the host problem are the same machine

Selection and host relations are one mechanism, because the filter portfolio that produces selection-based trust (application curation, price, visa and geography, referral networks) is exactly the machinery the surrounding community fails: the price filter on a local income, the network filter where discovery is referral-gated, sometimes a literal access filter at a venue closed to local residents. In the cases studied, exclusion is the selection model's default output, and not a lapse in host-relations practice.

Host-relations advice framed as etiquette therefore misses the lever, which is economic and design-level. Every documented counter-case redesigned the operating model itself; Ch.6 carries the three moves (locally priced access, municipal co-design, local-economy integration as a design principle).⁸

The confidence mark is held down deliberately, because the exclusion accounts converge across organiser, participant and local-host vantages, but the cluster is demographically narrow, the causal join from filter portfolio to host exclusion is inferred in most cases, no host official has been interviewed in the evidence base, and the single dedicated local-host account carries a documented conflict of interest (disclosed in Appendix E). The hypothesis-level version (exclusion is the default unless the operating model prices locals in by design) is registered and awaits scoring. The failure pattern reads as capacity gap and default economics, not extraction by design.



Finding 05

Knowledge travels with people; tools get orphaned

The transfer bottleneck in this field is artifact decay, not idea scarcity, for knowledge demonstrably moves when people move: returning participants, migrating organisers, standing organiser calls across communities. Personnel-mediated transfer is supported, with one condition attached: it fails without a receiving structure in the adopting institution.⁹

The artifact channel is the register's one contested entry. The support side is entirely design-stage or aspirational while the counter side is entirely outcome-level: a shared community platform no longer maintained, a governance design surviving only in participant memory, and a practitioner's verdict, after three years of the flagship lineage's tool-building, that no tool a community can actually use exists.¹⁰ No artifact studied has a documented cross-community adoption record. The asymmetry is itself a finding about the field's tooling. The ecosystem funds artifact construction and does not fund artifact maintenance, because relationships and careers reward carrying knowledge while no one's model pays for upkeep once the event that produced a tool ends.

Funding maintenance and receiving structures before more tool-building is the precise implication. One boundary: the observation window is roughly 2023 to 2026, and tools with longer maintenance horizons may simply not have matured inside it.


Finding 06

The capacity floor is high, the ceiling is real, and both are priced at current technology

Minimum viable institutional capacity is higher than most organisers expect, particularly for safety and host-jurisdiction relations; the threshold scales with size, scope and design rather than sitting at one number. Above it, coordination costs rise superlinearly under participatory governance and roughly linearly under delegated governance, which is why Finding 3's centralisation drift is rational, not lazy.

The ceiling for participatory models appears to sit somewhere in the low hundreds of participants, a working hypothesis anchored by a single comparison (one stable case below that range, one case above it where participatory governance did not hold), not a benchmark.¹¹ One qualification is disclosed once here instead of per claim. Both the floor and the ceiling are measured at current coordination technology.

Two practitioner voices argue that AI-delegated coordination will move them; the argument is design-thesis grade, no deployment outcome exists anywhere in the evidence, both voices are tool-side actors with an interest in the thesis, and the field's revealed preference so far runs the other way (the largest permanent node's response to scale was a bespoke community platform occupying roughly a tenth of its staff). The report holds this as an open question.



Finding 07

The ecosystem is pre-tipping-point, partly by its own selectivity

The ecosystem lacks the capital density, talent density and narrative category that a self-reinforcing adoption curve requires, and the re-analysis adds a fourth component: a demand-side ceiling, in a participant base that is demographically narrow and shows signs of saturation. Publicly available institutional funding dedicated to the space is limited; the talent pipeline recycles a small pool of organisers; the category remains hard for mainstream institutions to evaluate, fund or regulate.¹²

The mechanism is partly self-imposed, which is what makes the finding actionable. The network-state framing operates as a selection filter over incoming resources: it attracts mission-aligned crypto capital, patron sponsorship and tool-builders, and it repels institutional investors, non-crypto operators and the professional talent that compounds into durable capacity; what it admits tends not to accumulate, and what could move the field past its ceiling is what the vocabulary screens out. The observed fixes fit the diagnosis: protocol language adopted explicitly because institutions can evaluate it, “network city” naming, one first-documented dedicated fund that failed because limited partners were unconvinced of the asset class. Part of the field is exiting the vocabulary instead of managing it; a filter being abandoned still confirms the mechanism.

The direction of causality is inferred rather than demonstrated (no source observes the same project under both framings), and the researcher's own position on the field's trajectory is disclosed in Appendix E; hence the medium mark. For operators, framing is a design lever with measurable resource consequences. For funders, the constraint on this field is partly a story problem, which capital alone does not fix.

Open questions the register carries forward

Three questions remain open, and the report flags them and leaves them open. First, local integration: whether host-community exclusion is the default absent deliberate design awaits formal scoring (the Finding 4 hedge stands); closing that gap is the clearest research priority the report can name. Second, portable identity and reputation: informal mechanisms still carry selection at today's scale, and Ch.6 holds the open question of when cross-event volume outgrows them. Third, the AI cost-curve question from Finding 6. Each is a place where the next edition of this report could differ from this one.

Notes

1. Confidence marks follow Sidebar B in the Executive Summary, rounded-square scheme: HIGH (#0E00C2; 3+ independent sources from 2+ source streams), MEDIUM (#1862F5), LOW (#1B9CF3). Informant bases in these notes are deduplicated at informant level; exact counts are maintained in the reviewer materials.

2. Measurement-gap evidence: a broad multi-source coded base, plus paired operator and attendee survey items whose figures are withheld pending verification and stated without numbers. The deployed-but-unadopted instrument rests on one funder-side verbatim account, strong for the instrument's existence and deployment. Counter-examples: one residency published a line-item cost breakdown; at least one organisation collects cohesion data (not yet verified).

3. Documented in a single funder-side practitioner's verbatim account, including a written refusal; each link is separately corroborated across the multi-informant subsidy and measurement evidence bases (reviewer materials). The associated dollar figure is single-source and held in the reviewer materials, unprinted. One major funder adapted instead of exiting, so the chain describes a pressure, not a completed exit.

4. Financial-sustainability and subsidy evidence: multi-informant deduplicated bases across operator, funder, organiser and doctrine sources; the field-maturity base is broader still (reviewer materials). Anchors: a single-patron account of open-ended monthly support the organiser himself called unsustainable, and one operator's designed defence (no single funder holds a majority, with a disciplined ticket-to-sponsorship ratio). Revenue-mix figures for one flagship organisation conflict between two informants and are excluded as unreconciled.

5. The zero-ticket causal account is organiser-first-person (opposed the decision, predicted the drift, observed it) but single-organisation; the payment-direction and provision refinements are each single-source. Stated as the strongest available reading of a small case set.

6. Srinivasan, The Network State (2022). Governance evidence: corroborated across a broad deduplicated informant set on founder-led defaults and legitimacy grounding (reviewer materials); a verbatim community post-mortem documents proposal fatigue, consent-by-exhaustion and proposal-writer capture at one flagship's on-chain governance experiment. An operator survey item on the basis of authority points the same way; its figures are withheld pending verification and not printed.

7. A multi-source deduplicated base including institutional post-mortems and desk research; the comparative claim still lacks an off-chain counterfactual, which is disclosed rather than resolved. Live experimental stacks with no outcome data are tallied on neither side.

8. Exclusion cluster: six informants deduplicated, spanning several events and organisations, including one bilateral pair describing venue access closed to local residents; the visa filter is documented in verbatim testimony as operating before any organiser curation. Counter-cases: a three-city municipal run in the host language with local decision-makers present (second-hand), one state-backed model, and two small design-principle cases with unmeasured outcomes. A survey item on discovery channels is consistent with network-gated entry; its figures are withheld pending verification.

9. Personnel-mediated transfer: several new consistent sources in the latest evidence pass across organiser, director and participant vantages; the receiving-structure condition is anchored by a funder-side account of transfer failing where no adopting structure existed.

10. Every coded artifact-side support is design-stage, every outcome-level entry a counter, all retrospective (reviewer materials). The three-year verdict is a verbatim but unconfirmed-speaker account from an ecosystem retrospective session, paraphrased for that reason. A reported deprecation of the ecosystem's flagship credential tool remains externally unverified and is not asserted; the same tool is endorsed as the exception within the most critical account, so the artifact record is thin rather than uniformly negative.

11. Capacity evidence: a broad deduplicated informant base; coordination-scaling evidence corroborated across a substantial informant set, with the superlinear pattern validated theoretically and across three institutional domains in desk research; the scaling exponent is unmeasured. The low-hundreds ceiling rests on one stable case in the tens-to-low-hundreds range against one far larger case where participatory governance did not hold, and is labelled a working hypothesis wherever it appears.

12. Field-maturity evidence: a broad deduplicated informant base plus desk checks documenting the absence of mainstream consulting coverage; the failed dedicated fund is a single investor's account. One frequently cited talent-density estimate from a recorded operator call is capped for a disclosed conflict of interest and is not used. The research lead's prior view on the field's tipping-point trajectory and the network-gated recruitment of the informant base are disclosed in Appendix E (positionality).

Mechanics.

How a pop-up city runs in practice: who gets in, who decides, what motivates contribution, how rules hold, and how the money works.

Ch. 06 35

Membership and selection

Ch. 07 39

Governance and legitimacy

Ch. 08 43

Incentives and contribution

Ch. 09 48


Enforcement and operations

Ch. 10 52

Economics

Membership and selection.

Hypotheses

H3 **Qualified**H4 **Qualified**H5 **Qualified**H22 **Assess**


Who gets in is the decision the rest of a pop-up city inherits: the entry filter sets baseline trust, the quality of unplanned collaboration, the range of the cohort, and how the venture reads from outside. In the cases studied, selection is not a single dial but a portfolio of filters, and the application form is the only one the operator consciously designs.

Before any application is read, visa and geography have already screened out much of the Global South,¹ discovery has run through referral networks ahead of public directories,² and the price has sorted by income. Even the open pole selects: gatherings with no application admit through the friend network, which buys cohesion at a homogeneity cost. Exhibit 6.1 orders the filters by when they operate and who controls them.

Exhibit 6.1. Flow diagram

Selection runs on filters the operator never designed

The filter portfolio in the order it operates. Which filter carries the selection matters as much as how tight it is.



Source: selection evidence base, 32 interview sources deduplicated at speaker level; visa filter from one verbatim account; discovery-channel survey figures await final review and are not plotted.

Tighter curation improves collaboration quality and narrows the community at the same time, and the strongest operators treat this as a conscious choice: ORG-009 accepts roughly 25% of around 1,000 applications a year and defends the narrowness as design,³ while one fellowship participant described a flagship cohort as an echo chamber missing the collaborators the work required. The axis of selection matters as much as its tightness: selecting on shared philosophy or demonstrated contribution tightens a community, while selecting on prestige admits people one practitioner called unhealthy contributors.⁴

What predicts both contribution quality and inclusion is less the number of tiers than the permeability of the boundary, and four mechanisms recur: work-exchange, sponsored fellowships, vouching in which the voucher bears reputational responsibility, and referral-bound growth from a hand-picked core.

Permeability alone does not deliver inclusion: a subsidised tier with no agency and no path to continued participation produces alienation, not diversity (“a monkey in a zoo”, in one organiser’s account of a scholarship cohort), and sponsorship that opens the financial gate can leave the network gate untouched, producing cohorts that are geographically diverse and network-homogeneous at once. Whether multi-tier models outperform single-tier ones remains open: no matched comparison exists, and the clearest single-tier case is itself growing an informal contribution gradient.

Price does two jobs, economic gatekeeping and commitment signalling, and the second survives without the first. Exhibit 6.2 sets out the observed configurations: free entry with full provision drifted participants into consumer mode at one heavily sponsored zero-ticket event, an outcome the organiser predicted and then watched happen; free entry with no provision sustained a producer culture at village scale; and where organisers paid contributors, the direction of payment created the right to demand work. Ch.8 carries the funding-structure mechanics.

Reputation currently travels with people, not through systems: returning participants run at roughly 40% across iterations at ORG-164 (one organiser-reported figure), and 86% of attendee-survey respondents had attended 2 or more pop-up city events, on a self-selected base of 68 respondents. At today’s scale of tens of active events, referral, vouching and alumni memory carry selection cheaply, which is why formal portable-identity tools show little adoption so far; whether that holds as cross-event volume outgrows the social graph is a forward-looking question the register holds open.

Exhibit 6.2. Comparison matrix

Price, provision and payment direction set the mode

Three variables rather than one appear to carry the contributor-to-consumer drift. Each row is a small number of documented cases; row 1 is a single organisation.

Configuration	Price	Provision	Observed participant mode
Zero-ticket, fully provided	Free	Full	Consumer mode; drift predicted and observed by the organiser
Zero-fee, self-provisioned	Free	None	Producer culture at village scale (20 to 30 people, referral-gated)
Organiser pays contributors	Reversed	Partial	Accountable production; payment direction confers the right to demand work
Premium-priced, fully provided	High	Full	Entitlement tension: payers expected to contribute feel entitled to consume

Source: organiser and participant accounts, one case per configuration except row 1 (single organisation, first-person organiser account; zero-ticket).



The filters have one more face, turned outward, because the machinery that produces trust inside the gathering is what the surrounding community meets from outside: prices set against international incomes, discovery through networks local residents are not in, and occasionally a literal boundary, as in paired accounts of venue access closed to local residents. In the cases studied the accounts point one way, and the counter-cases share a feature: each changed the operating model rather than the manners, through locally affordable base pricing, municipal co-design in the host language, or local-economy integration by design. The report reads this as default economics rather than intent, a design problem parts of the field are already pricing in; the stronger claim, that exclusion is the default unless the model prices local residents in by design, is registered and awaits formal scoring on an evidence base that is still narrow and contains no host-official voice.⁶ Ch.10 carries the host-side evidence.

One limitation frames everything above: the filters have never been tested separately. Most cases run several at once, so the field knows the portfolio exists but not yet which filter carries which cost.

Notes

1. One practitioner's verbatim account: "I think actually the main filter is the visa. If someone doesn't have [a visa]... they are always put on the side." The observed mitigation is administrative: early ticket releases and deadlines that leave time for visa processing.
2. A survey item on discovery channels is consistent with referral-gated entry; its figures are withheld pending verification and not printed.
3. Organiser-reported for one organisation: roughly 25% of around 1,000 applications a year, 50 to 80 participants per rolling cohort. The echo-chamber account is from one fellowship participant at a flagship organisation.
4. "When you rely on prestige as the thing you are selecting based on, you wind up with people there who are maybe not good, healthy contributors" (one practitioner, partial transcript).
6. The exclusion accounts rest on six sources spanning local residents, international attendees and practitioners across several organisations; the causal link from filter to exclusion is inferred in most of them. The one dedicated local-host account in the evidence base is a single voice carrying a documented conflict of interest, disclosed in Appendix E; no claim in this chapter rests on it alone.

Governance and legitimacy.

Wherever real stakes appear, authority moves to whoever can be held liable.

Pop-up cities need fast decisions on programming, logistics, conflict and sponsor commitments, and they attract participants with little patience for authority they did not choose. The field has largely resolved that tension in one direction: of the 15 operators surveyed, 5 ran founder-led governance, 6 core-team consensus, 3 a hybrid and 1 an assembly model,¹ and the interview record behind those counts runs the same way, with founder-led operation persisting at every observed scale and duration, including the largest permanent case studied. Ch.5 states the finding; this chapter adds the mechanics, because the field describes itself as a laboratory for decentralised governance while mostly testing variations of founder authority, and that gap reads as misdescription rather than hypocrisy.

The mechanism that keeps authority central is liability: whenever a decision carries real stakes (brand risk, sponsor money, safety, a deadline), authority is pulled back to whoever can be held responsible for it, and the participatory layer survives only where nothing expensive depends on it.² Every documented transition runs from stated participation towards centralised operation, and none runs the other way; that is an absence-of-counter-example claim, stated as such. The clearest worked example is ORG-044, which opened with an aspiration to community-based conflict resolution and reverted under coordination volume, having in its operational lead's words "morphed into a very traditional vertical NGO".³ Sponsors with material stakes who attend in person add an informal influence layer alongside the formal design. None of this makes the aspirations insincere: one flagship operator holds decentralisation as an explicit end-state, describes it as years away, and keeps experimenting towards it.


Hypotheses

- H6 Qualified
- H7 Qualified
- H8 Supported

Exhibit 7.1. Flow diagram

Governance drift runs one way

Drift from stated participation to centralised operation, and the one documented way designs hold. The absence of a reverse case is an absence claim, not proof of impossibility.



Source: governance evidence base, 19 interview sources deduplicated at speaker level; reversion case first-person from its operational lead.



Designed participation is the usable exception: where a mechanism is architected before any dysfunction appears, it holds, with costs, and where design is absent the result is hierarchy that works but concentrates load on a few people. The cleanest resolution the evidence supports is to choose a governance mode deliberately and say so, because centralisation that is chosen and disclosed spends far less legitimacy than centralisation that contradicts the stated design.

Legitimacy itself is paid in more currencies than process. Five recur in the cases studied: demonstrated competence and track record, founding authority as such, consent expressed through voluntary entry and exit, coherence between the story told and the thing built, and volume of contribution and recruitment. Competence is the most heavily drawn on, and not the only one; consent-through-exit does real work in a form anyone can leave in an afternoon. Two cautions attach: these currencies are partly steerable by the leadership they legitimise, a clause held at lower confidence,⁴ and a track record that goes unrecorded is impossible to audit, which is why the cheapest legitimacy intervention available to operators is the record-keeping case Ch.5 makes, with governance reform second.

Exhibit 7.2. Comparison matrix

Five currencies of legitimacy, and how each erodes

Legitimacy in the cases studied is earned through non-process currencies; competence leads but does not stand alone.

Currency	How it is earned	How it erodes
Competence and track record	Delivered events; visible operational skill	Unrecorded outcomes leave it unauditable; punctured faster than a record
Founding authority	Having started the thing	Stagnation across editions; critical voices go quiet
Consent through entry and exit	Voluntary joining; a real option to leave	Weakens where exit is costly mid-event; counts only those the filters admitted
Narrative-moral coherence	The story matching the thing built	Capital and talent arrive while community substance hollows
Contribution and recruitment volume	Participants building and bringing others	Volume mistaken for endorsement of how decisions are made

Source: legitimacy evidence base, 24 interview sources deduplicated at speaker level, plus published doctrine; an operator survey item points the same way but its figures are withheld pending verification.



On-chain governance needs careful scoping because the field's experiments there are misread in both directions. For collective decisions (token voting, delegation, treasury allocation), the mechanisms introduced more failure modes than they resolved in the cases studied, chiefly proposal fatigue, consent by exhaustion, and capture of outcomes by whoever drafts the proposals. Narrowly scoped instrumental uses (credentialing, access control, quadratic-funding payouts) show mixed-to-workable records and sit outside that claim; the most pointed critic in the evidence still endorses the credentialing tool. One disclosure travels with the finding: no matched off-chain comparison exists, so the record shows that on-chain governance did not solve these problems, not that off-chain governance would have. The fairest reading is of a technology still maturing, and the live experimental stacks, none with outcome data yet, are the designated test of this finding for the next edition.⁵

Six warning signs recur before governance failure, and they are cheap to check: multiple editions with no documented change in how anything is decided; participants who say they cannot raise dissent without risking their place; town halls and votes described by the people in them as theatre; a stated mechanism abandoned at the first hard decision; everything routing through one person, which is rarely felt below around 50 participants and becomes the main constraint above 150; and a well-built public coordination scaffold that nobody populates,⁶ which for a field positioning itself as governance R&D means the experiment ran without leaving a record anyone can learn from. The consolidated design decisions sit in Ch.13, Exhibit 13.1.

Notes

1. Operator survey, 15 respondents; a small base, reported as counts rather than percentages for that reason.
2. The rational-choice version, from one practitioner: "No one wants to delegate serious decisions to decentralised communities when valuable brands are at stake."
3. First-person account from the event's operational delegate; single organisation. The proximate cause in that account is coordination volume rather than event duration.
4. One practitioner account ("leadership has the ability to shift which things are considered high status", partial transcript) plus the founding doctrine, which prescribes founder authority from day one (Srinivasan, *The Network State*, 2022).
5. The failure-mode record includes a verbatim community post-mortem of one flagship's on-chain experiment and institutional post-mortems from desk research; the live stacks are tallied on neither side until outcomes exist.
6. Observed across several public coordination pages of one early case, well built but barely populated; page-level counts are held in the research record rather than printed. The 50 and 150 figures are directional thresholds, not measured constants.

Incentives and contribution.

Mission, belonging and access do what salaries usually do, until they run out.

Pop-up cities are built by people who are not paid market rate for the work, and the incentive stack that makes this possible runs through intrinsic motivation (the mission and the work itself), social belonging, access to a specific cohort, professional credential, and only then direct financial reward.¹ The strongest pattern in the evidence is that the intrinsic and social layers sustain contribution more reliably than financial incentives do, with one boundary: financial incentives can buy a burst of participation, and the documented large-scale token-distribution experiments show exactly that burst followed by churn, so the claim covers sustained contribution rather than turnout.²

Hypotheses

H9 **Qualified**


H10 **Supported**

The same reliance runs the burnout pipeline: the contributors most committed to the work absorb the most work at the least compensation for the longest time, because intrinsic motivation quietly absorbs the compensation shortfall instead of signalling it, until it is depleted. One organiser's account of the role is continuous availability with no off switch,³ and a long-tenure practitioner describes the arc in three phases, enchantment, frustration, focus: when motivation runs down without financial reinforcement, the most experienced practitioners redirect their energy to commercially viable work, and the operational knowledge they carry leaves with them (Ch.5 Finding 5 traces where it goes, which is mostly nowhere). The warning sign is qualitative but hard to miss: a lead organiser carrying a sustained, uncompensated workload well beyond a normal full-time week. Exhibit 8.1 draws the pipeline, together with the counter-designs now appearing against it.

Exhibit 8.1. Flow diagram


Intrinsic motivation absorbs the shortfall until it does not

The pipeline from intrinsic reliance to knowledge loss. The counter-designs are real and recent; none yet demonstrates measured retention.



Source: incentive and knowledge-loss evidence bases, deduplicated at speaker level; the three-phase arc is one long-tenure practitioner's framing.





The field is beginning to design against the pipeline instead of absorbing it: producer roles, engineered status and the first organiser-layer fee splits recur in the cases studied, none yet with outcome data, and the design implication is to layer financial reinforcement on top of social and credential recognition before core contributors hit depletion rather than after.


Money shapes contribution from the other direction too, and here the evidence points at structure rather than volume. Unstructured subsidy, meaning grants or open-ended sponsorship with no accountability mechanism attached, distorts participant behaviour toward consumer modes and leaves the event's continuity resting on its funders' decisions, while subsidy with attached obligations converts its recipients into accountable producers, and the one documented milestone-based sponsorship arrangement delivered against every milestone agreed.⁴ Ch.6 carries the participant-side configurations (price, provision and payment direction, Exhibit 6.2); this chapter carries the structure principle and its two other surfaces.

The same cause distorts wherever the money touches, which Exhibit 8.2 maps. On the sponsor side, an organiser at one flagship organisation describes sponsors requesting more access and visibility than the community's design could absorb: where sponsor revenue is necessary for viability, sponsors acquire leverage over community design decisions.⁵ On the organiser side, allocation mechanisms that track social connectivity route resources toward well-networked organisers and leave intrinsically motivated outsiders under-resourced and, in the account that documents the chain, burnt out.⁶ The clearest designed defence at operator level is a funding discipline, not a funding source: never let a single funder hold the majority of revenue, and take sponsorship as support rather than as investment carrying rigid expectations.

Exhibit 8.2. System map

What unstructured funding touches

Where funding carries no obligations, distortion appears on every surface the money touches; the documented mitigations work through structure rather than exhortation.



The structured alternative: obligations attached – milestone sponsorship, payment direction that confers the right to demand work (Ch.6), and a cap on any single funder.

Source: subsidy-structure evidence base, deduplicated at speaker level; sponsor-side vector rests on one organiser's first-person account at a single organisation. ■

Funders eventually meet the same structure from the buying side: when subsidy carried no obligations and the funder later asks what the money achieved, the record usually does not exist, which is the measurement gap of Ch.5 Finding 1 experienced as a purchasing decision. The observed funder response reads as practice maturing, not capital retreating, with the largest ecosystem funder in the evidence now sequencing trust before money and requiring a sustainability plan before funding; Ch.10 carries the funding-side evidence in full.

For operators the implication is the same at both ends of the relationship: attach obligations to the subsidy you take, keep the record that lets the person paying see what the money bought, and track the gap between what the core team carries and what it is paid.

Notes

1. The five-layer ordering is constructed by the research team from the interview evidence, ranked by how often each layer appears as the primary driver; it is a framework, not a measured ranking.
2. The burst pattern is documented in published analyses of a large 2023 token airdrop and its governance aftermath (Hall and Oak, 2023; the distributing project's own community retrospective, 2023). Disclosure: the "sustained" qualifier was added to the hypothesis after this burst evidence arrived; it scopes the claim rather than predicting it, and is recorded as such in the methodology appendix.
3. "You're on 24/7" (one organiser at a single event-series organisation). The three-phase arc that follows is one long-tenure practitioner's framing, presented as such.
4. The milestone-based case is a single sponsor-side account of one sponsorship programme across a small number of events; reported here as the documented existence proof for structured subsidy, not as a success rate.
5. One organiser's first-person account at a single organisation. Revenue-mix figures for the same organisation conflict between two accounts and are excluded as unreconciled; no revenue split is asserted here.
6. The allocation-capture chain is documented in one practitioner's first-person account of a community funding round; the associated dollar figures are informant-reported, unreconciled against public round records, and not printed.

Enforcement and operations.

Safety, conflict and host relations decide more editions than programming does.

Hypotheses

H11 Supported

H12 Supported

Pop-up cities are operationally medium-sized events carrying the governance ambitions of small cities. The difference between a good edition and a bad one is rarely the quality of the programming but whether anyone has thought through the medical emergency, the conflict that escalates past gentle correction, and the venue relationship that sours 2 weeks in. The capacity this requires spans hospitality operation, community management, facilities, safety and host relations at once, it is higher than most founding teams plan for, and the teams carrying it are small, often just a handful of people. Ch.5 Finding 6 states the finding; this chapter carries the mechanics of where the floor sits and what happens at the ceiling.

Behavioural enforcement in the cases studied runs on a 5-level ladder, drawn in Exhibit 9.1, and its two ends do nearly all the observed work. Selection is the real first line: visible norms mostly hold without any mechanism being activated, and most of the friction that does surface is resolved by a private word from an organiser with relational authority. What the ladder shows most clearly is the missing middle: no case studied has a documented, structured conflict-resolution process, and the space between an informal word and removal is where most interpersonal conflict actually lives. Operators increasingly say so themselves: one flagship organiser describes conflict handling as ad hoc, with a systematic framework planned but not yet built, and a practitioner's proposal for a cross-event warning system runs into the fact that the enforcement and safety infrastructure it would plug into does not yet exist.¹ The fair reading is that the field is early: the people inside it now name the gap themselves, and a named gap is the kind that gets resourced.

Exhibit 9.1. Layer diagram

The enforcement ladder and its missing middle

Five levels of behavioural enforcement. The ends carry the load, while the middle is where the cases studied

Level 5	Removal	Requires: perceived authority to make the call	the commonly documented final mechanism
Level 4	Conditional access restriction	Requires: staff who can monitor compliance	not documented as successfully implemented in any case studied; often skipped for level 5
Level 3	Formal warning	Requires: a documented code of conduct; willingness to record	present across most events; rarely used
Level 2	Informal organiser signal	Requires: an organiser with relational authority	most friction resolves here when that authority is present
Level 1	Social norm	Requires: selection that filtered for compatible norms	strongest first line observed; selection as the trust shortcut: participants assume others have passed the same bar
The missing middle No documented structured conflict-resolution process in any case studied (levels 3–4). Most interpersonal conflict lives between an informal word and removal.			

Source: enforcement and capacity evidence base, operator and organiser accounts deduplicated at speaker level; the ladder is the report's synthesis, coverage notes are editorial judgements from case evidence.



What all of this protects is ambient collaboration:² the unscheduled productive overlap between people doing different work in the same place, which is the outcome participants most consistently describe valuing and much of the reason the format exists. Coordination overhead pulls against it and grows faster than headcount, substantially faster under participatory governance, which is why the ceiling for participatory models appears to sit somewhere in the low hundreds of participants, a working hypothesis anchored by a single comparison rather than a benchmark (one stable case below that range, one above it where participatory governance did not hold; Ch.7 carries that case).

Two designs manage scale without full reversion to founder authority: functional role assignment, which divides the operational surface into bounded domains each run by a self-managing crew, with only narrow coordination across crews; and tiered authority, which routes every decision to the smallest tier able to resolve it. One caveat travels with both: in at least one well-documented case the real coordination ran through chat back-channels that leave no durable record, so the formal design was not where the work happened (Ch.7's warning signs carry the general form). Exhibit 9.2 sets the documented scale thresholds against the capacity domains the evidence covers.

The operational knowledge an edition produces mostly leaves when its people do. In the cases studied, feedback is rarely collected, final retrospectives are skipped under end-of-event exhaustion, and handovers run on continuing personal contact, not on documents; Ch.5 Finding 5 states the pattern, and its operational implication is unglamorous: run the retrospective, write the runbook, name and brief a successor, and fund the maintenance of the tools that exist, not the construction of new ones. Counter-designs are appearing, and they are the most encouraging development in this part of the evidence: resident nuclei that hold culture and process across editions, systems open-sourced and handed to the next location, permanent hubs positioned as continuity infrastructure. None yet reports a measured retention outcome, which is the honest state of the art, no criticism of designs this young.

Exhibit 9.2. Spectrum

Operational capacity by scale band

Bands reflect documented thresholds; domains are shown only where evidence exists. This is not a measured capacity comparison across archetypes.

	Up to around 50 coordination ceiling rarely felt	Around 50 to 150 load rising; one stable flat-governance case at 50 to 80	Above 150 coordination becomes the main constraint; low-hundreds ceiling zone
Coordination and governance	cohesion carries most coordination; voluntary roles without stakes still fail	flat participatory norms documented as workable	functional role assignment and tiered authority; the one documented reversion moved to delegated authority
Safety and conflict resolution	under-provisioned at every documented scale; conflict handling ad hoc; no structured process (the missing middle, Exhibit 9.1)		
Host relations and local grounding	designed-in local relations documented at small scale, at low cost	staffed as logistics across the cases studied; one heavily documented local- grounding failure is a flagship-scale case	

Source: scale thresholds from the interview and case evidence (around 50 / above 150); the low-hundreds ceiling is a working hypothesis resting on a single comparison.



The leading indicators of operational failure, meanwhile, surface quietly and early, though on no timetable the evidence can fix, and they sit in the design itself: a task volume no team of that size can absorb (the load that breaks organisers is volume rather than difficulty), funding contingent on a single anchor event or patron, and a stated process set aside at the first hard decision.³ Ch.12 carries the failure modes in full.

The host interface is one of the least invested capacity areas in the cases studied: host relations tend to be staffed as logistics, meaning permits, venues and suppliers, when the evidence reads them as a relationship that compounds or sours over years. One local host's account distills what genuine integration requires into a three-layer model, genuine intention, resource allocation and structural facilitation, all three at once, and the same account describes being recruited as a cultural bridge and ending up used as a service provider, which is what the absence of the second and third layers feels like from the other side.⁴ Ch.10 carries the host-side evidence, including continuity across political cycles and the spectrum of engagement models; Ch.13 (Exhibit 13.1) consolidates the operations design decisions.

Notes

1. The ad-hoc conflict handling is the account of one co-founder at a flagship organisation; the cross-event warning proposal is one practitioner's, reported with their own assessment that current events lack the infrastructure to participate. A related safeguarding gap around under-age attendees appears in the evidence at low strength and is held in aggregate form for the research record rather than detailed here.
2. "Ambient collaboration" is the report's editorial term for the unscheduled productive overlap that selected, co-present participants describe; it is not a phrase drawn from any interview.
3. The volume-not-difficulty mechanism is one practitioner's account, consistent with the wider burnout evidence (Ch.8); the anchor-contingency indicator is documented first-person at one large sponsored event whose fundraising rode a co-located flagship conference.
4. The three-layer model and the service-provider account are a single local host's voice, carrying a documented conflict of interest disclosed in Appendix E; no claim in this chapter rests on it alone, and Ch.10 carries the fuller host-side treatment, including the independent accounts that corroborate the pattern.

Economics.

Geography sets what a week costs; patrons cover most of it.

Hypotheses

H14 **Qualified**

H15 **Supported**



Across the cases with usable cost data, spending per participant-week varies by a factor of 5 or more, rising above 6 once participant-side costs at the founding case are counted, and the variation traces to venue, accommodation and geography, with programming ambition contributing little of it: broadly the same format of workshops, shared meals and unstructured time is delivered at every price point in Exhibit 10.1.

Scale drives the absolute budget while geography drives the unit cost, which is why the largest budget in the table is not the most expensive place to be a participant, and jurisdictional compliance adds a third driver in some operating contexts (visa friction, local entity requirements, banking barriers for crypto-denominated funding). For an operator the ordering is the practical point: the biggest economic decisions are made before any programme is designed.

Exhibit 10.1. Comparison matrix

What 8 cases cost

Figures are directional rather than precise: organiser self-reports and operator-published reports, in approximate USD, not independently audited. Case labels A–H are local to this chapter and do not correspond to ORG codes used elsewhere; the cost-type grouping is distinct from the Ch.3 archetype map (Appendix E relates the two). Per-participant-week figures mix organiser-side and participant-side costing where noted; the mixed basis is the standing measurement gap. Local-currency figures are converted at event-period monthly average rates.

	Cost type	Participants	Duration	Total budget	Per-participant-week	Confidence
A	Flagship (large-format)	800–1,000	~4 weeks	\$750K–\$2M	~\$190–\$625	High
B	Flagship (premium-location)	190	3.5 weeks	~\$665K	~\$1,000	High
C	Residency	60	3 weeks	\$75K	~\$417	High
D	Residency	rolling	rolling	\$1,500/month all-inclusive	~\$375	High
E	Permanent-anchor	50–70	4-month season	~\$165K invested to date	~\$170	Medium
F	Large-scale urban	~5,000 at peak	1 month	\$1.5M (sponsorship; zero ticket)	~\$75	Medium
G	Flagship (founder subsidy)	not standardised	~2 months	largest single-backer contribution documented in this research	~\$1,125/week (participant-side)	High
H	Partial data	not disclosed	not disclosed	\$70K scholarships for 28 recipients	not disclosed	Low

Source: cost evidence base; organiser-reported and operator-published figures.



Underneath every cost structure sits the question of who pays, and Ch.5 Finding 2 states the report's answer: patron capital underwrites the form as practised. This chapter carries the case-level texture. Among the flagship and large-format cases studied, none reaches financial sustainability without premium pricing, founder subsidy, grants or sponsorship, usually in combination. The premium-priced residency sustains itself at \$1,500 a month by selecting participants who can afford it; one of the best-documented mid-scale cases earned approximately \$210,000 in tickets against roughly \$665,000 of expenses, and with roughly \$600,000 in grants, raised across the event and a planned permanent hub, its tickets and grants together exceeded expenses; one large urban event raised roughly \$1.5M in sponsorship with no ticket revenue at all, and its fundraising rode the timing of a co-located industry conference; the founding case ran on the largest single-backer contribution documented in this research.¹ The field's most prominent large-format organisation describes its recent editions as ticket-led, per one organiser account, and the report prints no revenue split for that case.² Even the strongest sustainability performances at scale resolve into premium pricing (the largest comparator studied runs predominantly on earned revenue and still posted a deficit of roughly \$10.6M in its 2024 financial year, when its top-of-market tickets did not sell as planned), which confirms the pattern from inside.

Two cases demonstrate a third path: small permanent spaces at 30 to 70 participants, sustained by operator cross-subsidy from a separate income source, with neither premium pricing nor large-scale external subsidy. The path is genuine and it is scale-capped; neither case has been attempted at substantially larger scale, and the cross-subsidy is itself a quiet form of founder support. Beyond it, the documented exits from the pricing-or-subsidy bind (decommodified models, incubator-equity, endowment designs) remain designs awaiting results. What is visibly changing is the structure of the subsidy: milestone-based sponsorship in the one documented case delivered against every milestone agreed, milestone-based foundation funding is appearing in the record, and institutional engagement now reaches agency level (the executive summary carries the documented example).³ Exhibit 10.2 sets out the five funding models observed and where each is exposed.

Exhibit 10.2. Comparison matrix

Five funding models and where each is exposed

The models are drawn from the funding arrangements documented in the cases studied; most cases combine two or more.

Funding model	Where it is exposed	Mitigation observed
Zero-ticket / full sponsorship	Consumer-mode participation; weak commitment signal from participants	Reintroduce even a nominal ticket price
Sponsorship-dominant (majority of revenue)	Informal sponsor influence on community design; continuity dependence on sponsor decisions	Milestone-based structured sponsorship; explicit cultural boundaries
VC / startup-funded	Risk-appetite mismatch; community members bear downside they did not consent to	Separate community and startup entities
Self-funded (zero external subsidy)	Revenue cap limits scale; owner bears disproportionate risk	Owned infrastructure lowers the cost base
Grant plus scholarship	Donor fatigue when outcomes go unmeasured (Ch.5 Finding 1); return-on-investment mismatch	Output-linked grants; structured accountability

Source: funding and subsidy evidence base; the accountability-structure principle is Ch.8's.



For the places that host them, pop-up cities are easy to miscount. Direct economic impact may be modest: the largest documented cluster, around 2,000 people across 10 or more concurrent villages at its peak (a local estimate, not a measured count), was put at less than 10% of the host city’s peak-season tourist volume by the same local host, and a format built around a few hundred selected residents was never going to move tourist-spending statistics.⁴

Host jurisdictions that evaluate on those statistics are therefore evaluating on the wrong ledger and the wrong clock at once: the contributions the cases studied actually produced (new organisations, new policy frameworks, positions in global knowledge networks) arrive years later, survive changes of government only where the relationship has been invested in across electoral cycles, and arrive at all only with genuine local integration.

One local host’s trajectory shows what the value can look like when it lands: hosting a flagship regional edition led to founding two host-side organisations, one of them in pre-launch under a government partnership during the research period.⁴

Whether the value lands is a design question, and Ch.6 carries the mechanism: the filters that produce trust inside the gathering exclude the surrounding community by default. The host-side accounts gathered here run the same way across six voices and four cases: local relationships handled as transactions, cultural-bridge and facilitation work absorbed without compensation or a line in the cost accounting, and, in a strong independent account, the founding case described by an attendee as an expat bubble with no grounding in the place.⁵ Where local integration is performed rather than substantive, host jurisdictions may be right to discount the institutional value claim; the Ch.6 counter-cases all redesigned the operating model itself, not the surface conduct. The pattern reads as Ch.6's default economics at work, and parts of the field are already pricing it in.

The economics of this chapter close on a naming question. An organisation that intends to outlive its founders' generosity has, at the scales studied, one documented model that needs neither premium pricing nor ongoing external subsidy, and it is small. For everything larger, funding a pop-up city without a transition mechanism is closer to ongoing patronage than to investment, which is worth saying without embarrassment: patronage built the field, and structured patronage is visibly maturing. What would change the picture is specific: a flagship-scale case reaching break-even net of grants, at prices that do not select for wealth, across consecutive editions. The report does not yet show one, and the field is young enough that this reads as early rather than failed.

Notes

1. The residency price is independently confirmed by an organiser account and a participant account; the mid-scale figures are converted from the operator's own published residency report at event-period average rates; the sponsorship-only figure is a first-person organiser account of funds raised, not audited cost; the founding-case subsidy is reported by a former core organiser and is stated here as magnitude class rather than a figure. The comparator's 2024 deficit is taken from its public FY2024 financial filing and its own published financial summary, which attribute the shortfall to higher-priced tickets not selling as planned.
2. Accounts of that organisation's revenue mix differ between a founder account and a core-team account. Following researcher review (2026-07-14), the report prints no percentage for this case and rests its revenue-mix claims on other cases' verified figures.
3. The milestone case is a sponsor-side account covering 3 events; the delivery record is stated as the sponsor reports it. The agency-level engagement is documented in the executive summary §6 and is not re-derived here.
4. The cluster headcount, the tourist-share estimate, the transactional-use account and the host-side founding story are the account of a single local host, who carries a documented conflict of interest disclosed in Appendix E: the critique of shallow integration is also the pitch for the same host's bridge-building ventures. The tourist-share figure is an estimate resting on this source alone and is presented as such; no other claim in this chapter rests on this account alone.
5. Verbatim: "a crazy expat bubble that had no grounding in the place", from a recorded interview with a multi-case attendee and critic of the founding case. The wider pattern aggregates six voices across four cases, including a paired local-resident and attendee account of venue access closed to local residents; Ch.6 and Finding 4 carry the mechanism, this chapter the host-side evidence.

Outcomes.

What the form produces, how to measure outcomes that arrive years later, and the three ways pop-up cities fail.

Ch. 11	58
Outcome measurement	

Ch. 12	61
Three ways pop-up cities fail	

Outcome measurement.

Organisers interviewed for this research could describe in detail what made their event work; structured measurement of those outcomes was uncommon. Ch.5 Finding 1 states the finding; this chapter carries the detail underneath: what participants value, why the standard tools miss it, and where the deployment gap sits. The absence of measurement is not itself evidence of value, only evidence that value claims cannot yet be evaluated reliably in either direction.

01

What participants value, and why the standard tools miss it

Four valued outcomes recur across the interview evidence: serendipitous encounters, accelerated relationship depth, subjective life density, and trajectory or identity transformation.¹ None appears in any standard event evaluation documented in the cases studied, and the mismatch is inherited, not an oversight: post-event surveys measure programming quality, logistics satisfaction and intention to return because they were designed for events where programming carries most of the value. In a pop-up city the reported value arrives largely in unstructured time, in forms a programming questionnaire has no cell for. Programming quality remains a useful indicator of community health and operator capability; it is simply not where most of the reported value lives.

02

The sponsor side measures what it buys

Sponsor and participant measurement operate on different axes at the same event. A corporate sponsor in the cases studied reported its objectives fully met across three consecutive sponsorships, measured on milestone deliverables, headcount targets and survey completion; participants at the same events valued relationship formation, intellectual access and serendipitous collaboration, none of which appears in the sponsor's framework. A corporate sponsor reported its objectives fully met across three consecutive sponsorships, measured on milestone deliverables, headcount targets and survey completion; participants at the same events valued relationship formation, intellectual access and serendipitous collaboration, none of which appears in the sponsor's framework.² The gap is participant-side rather than universal: where the valued outcome is defined concretely

Hypotheses

H17 Supported

H19a Supported

H19b Contested

enough, existing instruments capture it perfectly well. What goes unmeasured is the relational and long-horizon half.

03

The horizon dimension

Most of that unmeasured half is a time problem more than a people problem. Both concrete and relational outcomes have snapshot forms visible at event-end, which existing metrics can capture, and long-arc forms visible months to years later, which nothing currently deployed captures; the same participant typically holds both at once, so a measurement system needs to distinguish outcome horizons rather than segment populations. Exhibit 11.1 maps the four cells.

Exhibit 11.1. 2×2 matrix

Outcome types by horizon: what gets measured

A measurement system that captures only the left column undercounts long-arc value for every participant type.

	Snapshot (visible at event-end; mostly measurable)	Long-arc (visible months to years later; mostly not measured)
Concrete / artefact outcomes	Deliverables shipped at event-end; demo-day output count (Appendix B Tier 2)	Deliverable adoption, persistence, citation 6 months to 5 years later; institutional innovation transfer (Tier 3); counterfactual productivity change (direction contested)
Relational / experiential outcomes	New collaboration pairs at event-end; return intention; willingness-to-repurchase; ambient collaboration quality (proposed)	Trajectory transformation; trust durability at 1 year and beyond; identity shift visible months later; collaboration persistence at 3, 6 and 12 months; subjective life density retrospectively assessed

Source: outcome-measurement evidence base; metric tiers per Appendix B; horizon windows are proposed instrument windows, not measured findings.



04

The deployment gap, and what closing it looks like

The strongest pattern in the measurement evidence is not that instruments are missing; it is that they exist and go undeployed. A funder built a working benefit-to-cost instrument and deployed it across multiple pop-ups, and no operator adopted it; a practitioner-designed wellness framework exists and is not widely used; the simplest trackable measures remain mostly untracked while operator practice defaults to attendance and budgets.³ Where measurement does get deployed, timing undoes it: the surveys that exist run at the exit.

Three approaches would close most of the gap. Willingness-to-repurchase asks a revealed-preference question (would the participant pay the same price for a successor event, knowing what they now know?), proposed by one organiser and not yet validated at scale. Network density change compares relationship and identity measures before and after the event. Longitudinal follow-up at 3, 6 and 12 months tracks whether collaborations begun during the event persist, which is the only way the long-arc column of Exhibit 11.1 ever gets filled.⁵ A separate practice sits alongside these: one founding-case community closed with a collective-memory ceremony, which captures the experience for participants rather than measuring it.

05

Where this leaves the field

The evidence across at least five independent sources converges: pop-up cities produce many of their most valued outcomes in unstructured time, in forms that only become visible months or years later, and the tools inherited from conferences and festivals were never built to see them.⁶ The gap is wider for long-arc outcomes than for snapshot ones, and no project studied yet runs the longitudinal infrastructure that would close the long-arc half. The encouraging part is how little stands in the way: the missing piece is deployment discipline, not measurement science.

Notes

1. The four-outcome grouping is the report's synthesis, not a measured distribution. Illustrations: a fellow valued serendipitous collisions at meals that led to research collaborations not possible in normal institutional settings; a researcher-participant reported a trajectory-altering shift still visible two years later (the speaker is on the research team; Appendix E records the disclosure).

2. Sponsor self-report, one sponsor, framed as such; the three-sponsorship count is register-confirmed. Participant-side accounts are from the same events' interview evidence.

3. The benefit-to-cost instrument rests on one funder-side verbatim account, strong for existence and deployment, weaker for its ranking results. The metrics triad is one practitioner's proposal. One organisation reportedly collects cohesion data, not yet verified.

4. The arc account (a mid-stay survey would catch the trough, an exit survey the high, neither the arc) is one participant's, notes-derived. The endline-only pattern comes from the strongest operator-published impact report; none of its figures has been cleared for use.

5. The 3, 6 and 12 month windows are proposed instrument windows (Appendix B Tier 3), not measured findings.

6. At least five independent sources across participant types; the hypothesis is register-graded supported, with the boundary that concretely defined outcomes are already measurable.

Three ways pop-up cities fail.

The people closest to a failure usually saw it coming.

Failure in pop-up cities is more predictable than the attention around any single case suggests. Three failure modes recur across the cases studied often enough, and are heavily enough evidenced, to read as properties of the format, and not bad luck, each at a different level: the event, the organisation, the ecosystem.¹

What is specific to pop-up cities is less that they suffer burnout, turnover or thin capital (most young formats do) than that each failure is knowable in advance: the information that would prevent it usually already exists somewhere in the ecosystem and fails to be captured, maintained or handed over, which is Ch.5 Finding 5 operating at three scales at once. The leading indicators are recognisable ahead of the crisis, though the evidence does not support any specific lead time, and the field is beginning to design against all three.

Hypotheses

H13 Supported

H18 Qualified

H21 Supported

Exhibit 12.1. Layer diagram

Three failure modes, three levels

Three failure modes by level of analysis. Each is knowable in advance; the connecting thread is knowledge infrastructure, not effort.

Ecosystem Pre-tipping-point ceiling

capital fragmented, talent dispersing, no category institutions can evaluate events keep launching without the funds, credentials and narrative artefacts that would compound success

Organisation Institutional knowledge loss

no playbook, no retrospective, organiser turnover; knowledge leaves with people second editions start without the first edition's knowledge; failure modes repeat

Event Operational overwhelm

small founding teams cannot sustainably cover the operational surface the format demands the edition looks successful while quietly consuming its operators

One thread: knowledge that is not captured, maintained or handed over (Ch.5 Finding 5).

Source: failure-mode evidence base; hypothesis register status Qualified – the modes are heavily evidenced, their frequency ranking is not measured.



01

Operational overwhelm

The operational surface of a pop-up city (housing, safety, conflict, host relations, programming, money) is the surface of a small town compressed into weeks, and in the cases studied it is carried by founding teams of a handful of people, with no off-season to recover in. The result is a failure mode that hides inside success: an edition can look excellent while quietly consuming the people running it, and it surfaces not as a bad event but as a team that does not come back.² Ch.8 carries the mechanism; what belongs here is the diagnostic reading: the capacity the format demands simply exceeds what small founding teams have been funded to provide, so the gap sits in resourcing, not in effort. Operators themselves now name it plainly, and that is where designing it out begins. This is a resourcing gap rather than an effort gap, and operators themselves now name it plainly, which is the first condition for designing it out.

02

Institutional knowledge loss

The organisation-level failure is that second editions start from scratch. The ecosystem's founding event produced no documented operational record, and an academic researcher with independent field data could still not establish afterwards what had happened or how its funder measured results;³ almost no one from its core team continued in pop-up city work. The pattern repeats: an organiser at org-082 wanted to write a playbook and lacked the bandwidth, and a final retrospective was never conducted. The fair description is the register's own: historically unaddressed, now attracting explicit counter-designs whose effectiveness is not yet measured. Talent scatters faster than it accumulates, with no training or credentialing pathway for organisers, and the category remains hard for institutions to read: the field is absent from mainstream consulting coverage, as far as this research could find.⁴

03

The ecosystem ceiling

At the widest level, the evidence points to a field that does not yet compound; Ch.5 Finding 7 carries the mechanism (capital, talent, category recognition). Capital fragments instead of compounding: no dedicated fund, endowment or recurring institutional commitment appears in the evidence, much of what does flow is patron capital whose continuity rests on individual funders' decisions, and the one flagship attempt at a real-estate hub stalled when its plan assumed parallel deal closure and fundraising and neither reached completion. Talent scatters faster than it accumulates, and the category remains hard for institutions to read.

One counter-view comes from inside the field's own leadership. Buterin's December 2025 essay observes a “worrying pattern” of pop-ups “getting shorter, smaller, less experimental over time”, a regression reading rather than a pre-maturity one.⁵ The two readings can both hold, and they point at the same missing infrastructure either way; an organiser at org-082 reads the same period more simply, as a form still taking off rather than one people have tired of.⁶ Growth continues in the case count, and what remains unresolved is whether that growth is linear or self-reinforcing.

Read together, the three modes are one failure at different magnifications: knowledge that does not outlive the context that produced it. The fixes sit at different distances from an operator's hands. Institutional knowledge loss is the near one, within reach of any team that prioritises it: run the retrospective, write the playbook, brief a successor. Operational overwhelm is the far one, because sustainable teams need funding commitments today's fragmented capital base cannot make. And the ecosystem ceiling's narrative side is addressable with the field's own output, in reports and case studies that outsiders can evaluate; this report is meant as a contribution to that shelf.

The cases studied produce real trust, real collaboration and genuine institutional learning; what the field does not yet have is the infrastructure to preserve, measure and transfer that value reliably, and that reads as a description of its stage rather than a verdict on its people.

Notes

1. "Heavily evidenced" is the register's own grading: the three modes are well documented across independent sources, and their ranking against other possible failure modes is not, because no frequency count of failures exists. The register also tracks modes beyond these three: governance-scale mismatch (Ch.7), subsidy-dependency distortion (Ch.8), and venue-access loss in one project that wound down within weeks of losing access to its site. Host-jurisdiction rejection is documented in adjacent forms and in one pop-up city studied (A.2).

2. Corroboration at low strength: in one flagship operator's own small iteration survey, 3 of 7 responding organisers selected operational overload as the biggest risk to future editions (internal instrument, directional only).

3. The researcher's account, verbatim: "Nobody ever explained what actually happened and how the funder measured results." Independent field data, one voice; consistent with the absence of any operational record in the desk research.

4. A desk check of major strategy consultancies' published output, 2024 to 2026; not exhaustive, and used here as an absence indicator, not a measurement.

5. Buterin, Let a Thousand Societies Bloom (December 2025). Cited as a published work; the essay's broader taxonomy of network societies is beyond this chapter's scope.

6. The organiser's phrasing: "Pop-ups are still taking off. It's not that people got bored of pop-ups." One operator's read, offered against the essay's; both are interpretations of the same growth data.

Practice.

The design decisions that carry the most weight, and the minimum institutional functions an edition needs to stand up.

Ch. 13

65


Design choices that matter most

Ch. 14

71

Minimum viable institutions

Design choices that matter most.



Every case in this research faced broadly the same design decisions; what varied was whether they were made deliberately or by default. This chapter distils the evidence of Ch.6 to 10 into 18 decisions across six areas, set out in full in Exhibit 13.1: for each, the options the evidence supports, the trade-off underneath, and a guardrail worth watching once the event is running.¹ The intended use is practical: set aside a few unhurried hours before the first event to walk the 18 with the core team, record what was chosen and why, and revisit the guardrails mid-event; the warning indicators tend to show up early, on no timetable the evidence can put a number to.

Membership and selection · 1–4

Selection carries more of the outcome than any other area, and its central trade-off does not resolve: curation produces the trust the format depends on and narrows the community in the same motion, price filters people separately from any application process, and tier structures expand reach only where the boundaries between tiers stay permeable. The cleanest discipline observed is to treat the choice as a choice, made deliberately and revisited against who actually turns up.

Governance and legitimacy · 5–7

Founder-led operation suits short events and scales poorly as coordination volume rises, and legitimacy built on competence alone calcifies without a feedback channel. The expensive omission in the cases studied is escalation: where no chain is defined in advance, authority accumulates quietly with the founder, and the first crisis, which is the moment the chain is needed, is the worst moment to design it.

Incentives and funding · 8–10

The funding model sets the participation mode before any community value is stated, through price, provision and payment direction. Subsidy is better read through structure than through share: the bands in Exhibit 13.1 are scales for structuring the conversation rather than measured thresholds,² and the patterns associated with sponsor-influence risk in the cases studied were a single sponsor carrying a share above roughly 40% and contracts without milestone gates. Contribution runs on intrinsic and social recognition, but only to a depletion ceiling, which is why the guardrail on decision 10 is about reviewing core-team compensation rather than adding further recognition.

Enforcement and operations · 11–14

Centralised coordination hits a bandwidth ceiling that appears to sit in the low hundreds of participants (a working hypothesis anchored by a single comparison), and functional delegation is the evidenced middle path. Undocumented roles create responsibility gaps that surface only in crises, and an underpaid core carries a high burnout risk, because it tells the team the work merits pay while paying less than the market rate. The enforcement ladder is Ch.9's, unchanged; most friction resolves at the first two levels when selection has worked.

Economics · 15–16

Archetype and geography set cost per participant-week before programming ambition does, by a factor of 5 or more, and the sustainability strategy deserves a clear name at the design stage: premium pricing selects for wealth, external subsidy leaves continuity resting on funders' decisions, operator cross-subsidy requires a separate income source. The guardrail is blunt: whether the event would survive the withdrawal of its single largest funding source.

Host interface · 17–18

The host relationship compounds or sours over years, on the community's clock rather than the event's. Logistics-only hosting is workable but raises friction with both the jurisdiction and the local community, and government partnerships bring continuity exposure across electoral cycles. The test that costs nothing to ask: whether the local partner will still exist, and still want to work with you, in 2 years.³

Exhibit 13.1. Comparison matrix (part 1 of 3)

Membership and governance

#	Decision	Options	Core trade-off	Guardrail	Conf.
Membership and selection · Ch.6					
1	Selection mechanism	Application curation / self-selection / referral	Curation produces trust and narrows the community at the same time	Demographic composition across iterations	High
2	Economic accessibility	Market-rate / scholarships / contribution-based / zero-ticket	Price filters separately from application; token scholarships risk alienating recipients; zero-ticket erodes commitment	Scholarship recipients' engagement vs paying participants	High
3	Tier structure	Single-tier / multi-tier permeable / impermeable	Multi-tier expands reach only when boundaries are permeable	Cross-tier interaction rate	High
4	Reputation portability	None / informal social / formal credentialing	Reputation transfers socially rather than through tools	Returning-participant share across iterations	Medium
Governance and legitimacy · Ch.7					
5	Governance mode	Founder unilateral / core team / emergent / sortition / hybrid	Founder-led suits short events; scales poorly as coordination volume rises	Decisions deferred or escalated beyond the planned mode	High
6	Legitimacy basis	Track record only / plus co-creation / participatory mandate	Competence generates legitimacy faster than process; competence alone calcifies without feedback	Participant willingness to voice public criticism	Medium
7	Escalation protocol	None / pre-defined chain / community assembly	Undefined escalation is a silent accumulation of founder authority	Time-to-decision on the first crisis	Medium

Source: the report's synthesis of the Ch.6 to 10 evidence base; subsidy scales advisory. Options are those observed in the cases studied; guardrails are early indicators, without a stated lead time.



Exhibit 13.1. Comparison matrix (part 2 of 3)

Incentives and enforcement

#	Decision	Options	Core trade-off	Guardrail	Conf.
Incentives and funding · Ch.8, 10					
8	Funding model	Zero-ticket / ticket-primary / sponsorship-structured / subscription / self-funded	Revenue ceiling against participant-commitment signal	Zero-ticket without a scholarship mechanism showed selection-quality loss	High
9	Subsidy level	Under 30% / 30–50% / above 50% (advisory scales)	High unstructured subsidy brings consumer-mode and sponsor-influence risk	A single sponsor above roughly 40% of revenue, or contracts without milestone gates	Medium
10	Contribution recognition	Financial only / plus social credential / social only	Intrinsic recognition sustains contribution only to a depletion ceiling	Core-team compensation absent or not reviewed annually	High
Enforcement and operations · Ch.9					
11	Coordination mode	Centralised / functional delegation / tiered authority / full participatory	Centralised coordination hits a bandwidth ceiling in the low hundreds (working hypothesis); functional delegation is the middle path	Governance overhead as a share of total team hours	Medium
12	Role clarity	None / informal / documented role charters	Undocumented roles create responsibility gaps that surface in crises	Escalations where “whose job was this?” is the question	High
13	Staffing model	All-volunteer / underpaid core plus volunteers / market-rate	All-volunteer burns out; an underpaid core concedes the work merits pay while paying less than market rate	Core staff carrying sustained full-time-plus load without corresponding pay	High
14	Enforcement ladder	Social norm / informal signal / warning / conditional access / removal	Most friction resolves at levels 1 and 2 if selection has worked; levels 3+ need a documented code of conduct	Whether the code specifies consequences as well as expectations	Medium



Exhibit 13.1. Comparison matrix (part 3 of 3)

Economics and host interface

#	Decision	Options	Core trade-off	Guardrail	Conf.
Economics · Ch.10					
15	Archetype and geography	Flagship urban / rural retreat / permanent anchor / large-scale urban	Cost per participant-week varies by a factor of 5 or more, driven by venue and geography rather than programming	Cost per participant-week against the cost-type reference points (Exhibit 10.1)	High
16	Sustainability strategy	Premium pricing / external subsidy / operator cross-subsidy / emerging models	Premium pricing selects for wealth; subsidy leaves continuity resting on funders' decisions; cross-subsidy needs a separate income source	Whether the event would survive withdrawal of its single largest funding source	High
Host interface · Ch.9, 10					
17	Host-relations model	Logistics-only / transactional partner / co-executor / genuine co-producer	Logistics-only is workable but raises friction at both jurisdiction and local-community layers	Local-community complaints; local participant share; local partner continuity	Medium
18	Jurisdictional strategy	Private-sector partner / government co-production / no formal partner	Government transitions bring continuity exposure across electoral cycles; co-production requires multi-year investment	Whether the local partner will exist and still want to work with you in 2 years	Medium



Exhibit 13.2. Comparison matrix

Where the stakes concentrate, by cost type

The decisions most associated with strain at each cost type in the cases studied. The association is drawn from single-case evidence, not from a measured frequency ranking.

Cost type	Highest-stakes decisions	Why these
Flagship (400+, multi-week, rented venue)	#9 subsidy level · #13 staffing model · #16 sustainability strategy	Staffing and subsidy are where the documented flagship strains concentrate (small-core burnout; sponsor-influence and continuity exposure); sustainability is the unresolved bind.
Residency (50–80, 4–8 weeks, curated)	#1 selection mechanism · #10 contribution recognition	Small curated cohorts stand or fall on selection quality; burnout risk concentrates in a tiny team.
Permanent anchor (30–70, seasonal, owned space)	#15 archetype and geography · #12 role clarity	The low cost base rests on owned space, and on operational knowledge surviving seasonal turnover rather than living in individuals.
Large-scale urban (1,000+, short-duration)	#8 funding model · #11 coordination mode	Participatory coordination did not hold at this scale in the case studied, and zero-ticket at this scale produced consumer-mode participation.

Source: cost-type bands per Appendix E; decision-to-crisis associations are the report's synthesis over single-case memo evidence.




The chapter is designed to be used twice, as a pre-launch walkthrough and as an in-event tracking reference, and the guardrail column is the part most cases studied had not built. Nearly all of the 18 decisions get made somewhere, by someone, whether or not they are named; the cases that read best in this evidence are the ones that could say what they chose and why.

Notes

1. The 18 decisions, the six areas and the guardrail set are the report's own synthesis of the Ch.6 to 10 evidence rather than an industry-standard checklist; confidence marks are claim-level judgements.
2. The subsidy bands and the roughly-40% single-sponsor mark are decision-framework scales for structuring the design conversation, not measured distortion thresholds; the operative principle (accountability structure over subsidy share) is Ch.8's.
3. The 2-year horizon is an operationalisation from organiser accounts rather than a measured threshold: the case evidence supports the direction (host relationships compound or sour over multi-year horizons) but records no measured 2-year figure.

Minimum viable institutions.

Seven owners and six numbers keep an edition honest.



Ch.13 sets out the decisions; this chapter carries the two instruments that keep them honest once the event is running. The first is a checklist of seven institutional functions that need a named owner before the event begins (Exhibit 14.1), and the second is the minimum measurement set a small team can feasibly track while running one (Exhibit 14.2). Both exist because of Finding 6: the institutional capacity a pop-up city needs, especially for safety and host relations, exceeds what most founding teams plan for, and the shortfall tends to surface mid-event.

The checklist is built against a pattern that recurs across the cases studied: documents exist without deciders. A code of conduct that states expectations but not consequences, escalation, or who decides is the commonest form, and conflict resolution is the clearest gap, the missing middle of Ch.9's enforcement ladder.¹ The floor scales with size rather than disappearing at small scale: the smallest sustained case, a permanent space of 30 to 70 participants with 200+ events across 5 years, covers the same seven functions a 400+ flagship needs formal owners for, through a stable core team. What the checklist asks for is ownership, not policy.

Exhibit 14.1. Comparison matrix

The seven minimum viable institutional functions

The leading indicators tend to show up early; how early, the evidence does not say.

#	Function	What it covers	Leading indicator if absent	Owner
1	Safety and incident response	A named person with authority over immediate participant-safety decisions, including removal	Code of conduct exists but does not specify consequences, escalation, or who decides	☐
2	Host-relations capacity	A dedicated owner for host-jurisdiction relations and local-community integration	Host treated as logistics rather than relationship; no named local-integration owner	☐
3	Financial management	Budget burn tracked against plan	Organiser knows the total budget but not the burn rate	☐
4	Conflict resolution	A process for handling interpersonal conflict before it escalates	No structured process between informal signalling and formal removal	☐
5	Coordination infrastructure	Channels and decision processes that let a distributed team operate without constant founder intervention	Coordination overhead climbing as a share of total team hours	☐
6	Knowledge capture	A mechanism for recording what happened, what worked and what didn't, before the team disperses	Post-event debrief not scheduled before the event begins	☐
7	Succession planning	Cross-iteration documentation sufficient for a replacement to take over any critical function on short notice	All-volunteer or underpaid-core staffing with no documented procedures	☐

Source: the report's synthesis across the Ch.9, Ch.12 and Ch.13 evidence; capacity-floor claim per Finding 6.



The measurement set follows the same minimum-first logic. The six Tier 0 metrics need nothing more than a shared spreadsheet and a weekly routine, and each detects a Ch.12 failure mode while it is still cheap to correct, tied to the Ch.13 decision whose guardrail it operationalises.

The same record does double duty: it is the track record Finding 1 says the field claims legitimacy on but does not keep. The set carries almost no numeric trigger points; the return-intention range is a reference band, not a threshold.³

Exhibit 14.2. Comparison matrix

The Tier 0 measurement set

The minimum measurement set a small team can feasibly track.

#	Metric	How to collect	What it detects	Failure mode	Dec.
1	Attendance and departure rate	Daily headcount or check-in	Early departures point to community-fit or safety problems	Operational overwhelm	#1
2	Programming engagement rate	Session attendance against total participants	Sustained decline is an early signal of dissatisfaction	Governance mismatch	#5
3	Organiser hours per week	Self-reported weekly log	A sustained workload well beyond a normal full-time week is the burnout warning sign (Ch.8)	Burnout and knowledge loss	#13
4	Return intention	Single-question post-event survey (1 to 10)	Case means in the cases surveyed sit mostly in a 6 to 8.75 range; markedly lower results warrant attention	Knowledge loss	#4
5	Budget burn against plan	Weekly comparison of actual to projected spend	Widening variance in early weeks points to planning failure or scope creep	Financial overstretch	#8
6	Formal incident count	Log every incident beyond routine logistics	An upward trend points to enforcement-ladder or selection problems	Operational overwhelm	#14

Source: the report's synthesis mapping each metric to a Ch.12 failure mode and a Ch.13 decision; return-intention band per the survey recompute against the locked evidence base, on small per-case samples.

Where the field is heading

The evidence also carries encouragement. Small permanent-space models have demonstrated sustainability at the 30 to 70 participant scale, and the 2 parent entities Ch.3 records with 5+ years of operation both run spaces of that kind. Returning-participant rates around 40% at the flagship case documented suggest the communities are producing relationships durable enough to bring people back.⁴ The range of funding models has widened, structured arrangements such as milestone-based foundation funding are appearing in the record, and new events keep launching at volume, into the narrower demand base Finding 7 describes.

What comes next is a set of open directions, each of which already has a first case in the record without yet having a pattern: openly shared funder-diligence frameworks, the three articulated but untested sustainability strategies (real estate, endowment, subscription), open tooling that propagates across events rather than staying inside the case that built it, and sustained incubation pipelines for the projects and communities that pop-up cities launch. Knowledge still transfers through informal networks, not shared resources, genuine co-production with local partners remains documented in one or two cases, host relationships still have to survive electoral cycles on the strength of individual relationships, and the path from a good event to a durable institution still runs through individual founder effort.

None of this reads as a verdict: the field is early, its main failure modes are knowable in advance, and the instruments in this chapter are cheap; a field in which most events named seven owners and tracked six numbers would already look meaningfully more mature to the funders, hosts and institutions now deciding how seriously to take it.

Notes

1. A bounded claim: no structured conflict-resolution process is documented across the 92 coded interviews. The absence of documentation is not proof that no informal process existed.
2. Organiser-reported figures: 200+ events; 30 to 70 participants; the flagship threshold from the same evidence base. The seven-function framework is the report's own synthesis, not an industry-standard checklist.
3. The return-intention band reflects the survey recompute against the locked base (completed July 2026): case means for completed cases sit mostly between roughly 6 and 8.75 out of 10, on small per-case samples. It is a reference band; boundary definitions sit in Appendix B.
4. Returning-participant figure from one organiser's report at a flagship case, corroborated by the survey base.

Appendices.

App . A	76
Focal case profiles	
App . B	89
Metrics dictionary	
App . C	91
Failure mode reference	
App . D	92
Terms used in this report	
App . E	93
Methodology	
App . F	98
Case directory	

Focal case profiles.

How to read a profile

Each profile pairs a factual at-a-glance table with short paragraphs on what the case is, its composition, programme character, distinctive signature, outcomes and artefacts, and lineage and continuation. Profiles are documentary; the analysis sits in the chapters and in Appendix C. Cases carry their operators' published names, and real names appear with consent; chapter prose keeps to the report's anonymised conventions. Monetary figures are stated in US dollars; amounts sourced in other currencies are converted at documented approximate rates and rounded.

Rights notice

This report is licensed under CC BY 4.0. Operator logos, brand assets, photographs, and other third-party intellectual property remain the property of their respective rights holders and are excluded from this licence. Reuse of these materials requires separate clearance.

A.1.	Zuzalu	77
A.2.	Network School v2	78
A.3.	Edge Esmeralda 2025	79
A.4.	Edge City Patagonia	80
A.5.	Burning Man 2024	81
A.6.	ZuBerlin 2025	82
A.7.	Zuitzerland Residency	83
A.8.	Zanzalu 2.0	84
A.9.	Forma Georgia	85
A.10.	Proof of Retreat #2	86
A.11.	Infinita City 2025	87
A.12.	Ârc Network School	88



Case A.1. Zuzalu

At a glance



Dates

March 25 – May 25, 2023 (62 days)

Location

Luštica Bay, Montenegro

Venue type

Apartments; hotel; event venues within the Luštica Bay marina resort

Archetype

Proto-jurisdiction (founding instance)

Scale band

500–1,500 (≈200 core residents plus ≈600 visitors)

Themes

AI Governance Longevity Science Tech

Edition

Sole edition under the original organising team

What this case is

The first pop-up city of the post-2023 genre and the common ancestor of most cases in this appendix. Initiated by Vitalik Buterin and convened by a core team from the Ethereum and d/acc communities; the field's use of “pop-up city” as its genre label and the “Zu-” naming convention spread from this instance.

Composition

Application-based, weighted toward the core team's networks (Ethereum Foundation, crypto-research, longevity-science); international, English-speaking, technically literate.

Programme character

Emergent and participant-led, with light formal governance. Community norms and a moderation framework were published before the residency opened, and the programme closed with the multi-day reflection-and-speech ritual that became the Zu-lineage template.

Distinctive signature

Established the genre's core move: a bounded cohort in shared space for an extended period, with emergent governance and spontaneous collaboration doing the work a formal programme might have done.

Outcomes and artefacts

Of the software produced or prototyped here, Zupass (ZK-credential identity) is the clear survivor, still in event use as of late 2025; ZuPolls, Daimo's consumer app and viaPrize did not sustain, and the quadratic-funding experiments ended with the 2024 Gitcoin Zu-round. Alumni coordination continues through Zuzone and the monthly Zuzone Townhall.

Lineage and continuation

Sole edition under the original team. Cases in this appendix tracing Zuzalu lineage: A.3, A.4, A.6, A.7, A.8, A.10, A.11, A.12.

Photograph: https://x.com/_TomHoward/status/1656711211973591054/photo/1



Case A.2. Network School v2

At a glance



Dates

March 1, 2025 – February 28, 2026 (1 year);
monthly cohorts begin on the 1st

Location

Forest City, Malaysia

Venue type

Hotel and conference resort property taken over
for residential campus use

Archetype

Proto-jurisdiction (long duration, high institutional
ambition)

Scale band

150–500 simultaneous at residential peak;
thousands cumulatively across the year
(organiser-reported)

Themes

Crypto **Governance** **Lifestyle** **Tech**

Edition

Second (v2); successor to v1 (September–
December 2024); v3 ongoing

What this case is

The longest single continuous edition among the focal cases and one of the most institutionally ambitious: a year-long residency with rolling intake, founded by Balaji Srinivasan and run by a Singapore-registered operating company in Forest City, Malaysia, under the network state framing set out in his book *The Network State* (2022), with an NFT-based access credential.

Composition

Written application. Roughly 150 residents simultaneous on a rolling monthly basis, predominantly international (90% or more), ethnically diverse with the majority demographic shifting meaningfully from month to month, skewing toward builders, content creators and researchers aligned with the network state framing or the digital-nomad lifestyle.

Programme character

Founder-led governance. In May 2025 Network School launched a long-term programme, shifting from the 90-day pop-up model (v1) toward a permanent society-as-a-service: all-inclusive housing, meals, gym access and 24/7 coworking, with flexible stays from one month to a full year.

Distinctive signature

The most explicitly network-state-aligned case in this appendix, with a comprehensive internal operating system, census, analytics software, and the first focal case to run a programme-year residency with rolling intake at this scale.

Outcomes and artefacts

Around 40 demos in February monthly showcase. The community reports more than 2,000 members across 16 cohorts since its 2024 founding, as of July 2026; the annual Network State Conference is its flagship gathering, and it also has hosted a Zcash hackathon, a Solana pop-up and multiple resident teams. An MOU with Kazakhstan's AI and Digital Development Ministry to establish a new campus was announced on July 21, 2026, following regulatory friction and a reported shutdown order in Malaysia.

Lineage and continuation

Own lineage and Zuzalu-adjacent: the programme follows its own network-state design, and several core team members carry Zuzalu lineage. v3 ongoing; the Forest City site functioned as a sustained geographic anchor.

Photograph: https://ns.com/_next/image?url=https%3A%2F%2Fassets.ns.com%2Fstatic%2Fdashboard%2Fnetworkschool%2Fnetworkschoolgallery2.jpg&w=3840&q=75



Case A.3. Edge Esmeralda 2025

At a glance



Dates

May 24 – June 21, 2025 (29 days)

Location

California, USA (Healdsburg, wine-country venue)

Venue type

Distributed village venues across Healdsburg, Northern California, anchored on the town plaza

Archetype

Launchpad

Scale band

500–1,500

Themes

Culture Science Tech

Edition

2nd Esmeralda edition; Edge City programme lineage

What this case is

The second Esmeralda edition and the Edge City programme's flagship US residency: 29 days in wine-country California with an intergenerational, family-inclusive composition.

Composition

850+ residents across the month, 80+ kids and families, and 66 countries represented (operator-reported), across fellowship, general-attendee, and sponsor-track builder tracks; the family-inclusive composition continued from 2024 (which spanned ages 6 weeks to 86 years).

Programme character

The 2025 programme ran 15 thematic tracks and 13 community-led experiments with attendee-organised sessions dominant (the 2024 precedent ran 551 sessions, 167 unique hosts, and 93% attendee-organised). Three deep-dive residencies anchored the month: a Long Journey residency (14 builders in synthetic biology, neurotech, AI infrastructure and fertility health), a World Builder residency (apps on World ID/World Chain), and the Edge City Fellowship (7 fully-funded fellows aged 18 to 25).

Distinctive signature

Its defining feature is a pop-up-to-permanent pipeline. Edge Esmeralda is co-hosted by Edge City and the Esmeralda Institute and runs as a prototype for Esmeralda, a permanent walkable town nearby, pitched as a modern Chautauqua for families: lifelong learning, a car-light life and independence for children.

Outcomes and artefacts

Residency outputs per the operator: three World Builder projects received residency bounties, and named launches from the month include Constellation (a neural-data startup that has since raised \$10M, operator-reported) and Harmonica (AI-moderated dialogue tooling). The operator's community letter additionally reports \$200M+ raised by 40+ Edge alumni and 200+ projects launched.

Lineage and continuation

Part of the Zu-family by personnel, early grant support, and the network. Edge City lineage; cross-reference Edge City Patagonia 2025 (A.4). Edge Esmeralda 2026 ran as the third edition (May 30 to June 27, 2026), drawing roughly 850 residents from more than 60 countries per the operator's month-in-review.

Photograph: https://cdn.prod.website-files.com/6691010455371864d09b8800/6862db320dba3109a6daf932_EE25-month-recap-04.jpg



Case A.4. Edge City Patagonia

At a glance



Dates

October 18 – November 15, 2025 (29 days)

Location

San Martín de los Andes, Argentina

Venue type

Hotel; co-working space; local-business partner venues (Le Village Hotel, BIT Cotesma, Del Sauco)

Archetype

Launchpad (regional modifier)

Scale band

500–1,500 (600+ ticketed participants per operator recap; 91 Argentine and 150+ local Patagonian; 72 countries)

Themes

Tech Culture Startup & Biz Science

Edition

9th Edge City edition; 4th month-long edition

What this case is

The Edge City programme's month-long edition in the Argentine Andes, a popup village in San Martín de los Andes built around themed residencies.

Composition

More than 600 ticketed participants from 72 countries, 91 of them Argentine and over 150 local Patagonians, with a crypto-ecosystem-heavy skew; 62 children, 50 of them local, took part weekly. A work-exchange and stipend tier lowered the entry cost for some, though dollar-denominated costs kept attendance out of reach for many local Argentines (two independent participant accounts).

Programme character

11 themed residencies anchored the village across more than 700 participant-led sessions. The strongest civic thread was the Argentina Onchain Residency, whose teams built nine prototypes shaped by local public officials rather than a crypto-native audience, with a few local councils set to run pilots (operator-reported); other tracks ran d/acc work with Protocol Labs and a \$5,000 buildathon.

Distinctive signature

A regional edition defined by its local footprint: over 150 local Patagonians took part, the Argentina Onchain Residency built prototypes directly with local public officials, and the month's clearest institutional trace (Patagonia Labs) formed on the host side.

Outcomes and artefacts

More than 100 projects ran across 4 demo days, and the clearest institutional trace is Patagonia Labs, a local incubator spun out of the coworking partner BIT Cotesma, where over 20 local founders pitched and two won initial funding. Case-survey satisfaction (eight responses) was mixed, with low value-for-money ratings and calls for a firmer, less web3-tied business model.

Lineage and continuation

Part of the Zu-family by personnel, early grant support, and the network. Edge City lineage (Denver, South Africa, Austin, Lanna, Bhutan, Patagonia and Esmeralda since 2023). The programme opened its first editions in India (Goa, October 11 – November 1, 2026) and Bhutan, a mentor-residency partnership with the Pelsung Innovation Lab around Gelephu Mindfulness City (operator-reported).

Photo: adaptation of https://cdn.prod.website-files.com/6691010455371864d09b8800/692da63cac736758f0b31ac2_24521183.jpeg



Case A.5. Burning Man 2024

At a glance



Dates

August 25 – September 2, 2024 (9 days; annual, last week of August into first week of September)

Location

Black Rock Desert, Nevada, USA

Venue type

Federal public land under annual Special Recreation Permit and Federal Register temporary closure

Archetype

Boundary case (scale benchmark, outside the four-archetype map)

Scale band

10,000+ (tens of thousands; off-chart relative to the rest of this appendix)

Themes

Culture Lifestyle

Edition

Annual since 1986 (Black Rock Desert from 1990; cancelled 2020 and 2021); the 2024 edition is the most recent with public-record documents

What this case is

The appendix's comparison benchmark, included deliberately even though it sits at the boundary of the working pop-up-city definition. The 9-day annual convening operates under a Federal Special Recreation Permit, a Federal Register temporary closure (2023–2027), and the Burning Man Project's Restated Bylaws (2014).

Composition

Tens of thousands of participants, selected by ticket purchase with tiered pricing and no ideological screening. Community-organised camps are the primary unit of social structure.

Programme character

No central session programme; participant- and camp-organised programming dominates. The ten principles form the normative backbone, with a thin central organisation (the Burning Man Project, a 501(c)(3)) devolving heavily to camp leadership.

Distinctive signature

By far the largest case in this appendix and the only one under a federal-level permit regime. Retained as the reference point for scale aspiration: the network-state and d/acc lineages cite it as proof of ephemeral-city viability at tens-of-thousands scale.

Outcomes and artefacts

Annual editions since 1986, interrupted only by the 2020 and 2021 cancellations; the ten principles propagated as a reference ethical frame; philanthropic and ecological programmes documented in the Impact Report. Its role in this appendix is comparison rather than causal chain.

Lineage and continuation

Annual, with the 2020 and 2021 editions cancelled. Pre-dates the Zu-family programme by three decades; an ancestor rather than a descendant.

Photograph: Dawn at the Temple of Together, Hannes Haefele - Luminae.eu, 2024. <https://gallery.burningman.org/asset/9207a607-dc8e-4ec3-a146-8d88c68d32f2?i=17&q=2024>



Case A.6. ZuBerlin 2025

At a glance



Dates

June 14–22, 2025 (9 days)

Location

Berlin, Germany

Venue type

Urban venues; paid accommodation (shared sleeping at peak occupancy)

Archetype

Sprint (urban, short, sits on top of city infrastructure)

Scale band

50–150 (~70 peak participants per operator survey)

Themes

AI Culture Governance Tech

Edition

Second ZuBerlin; predecessor ZuBerlin 2024 (September 1–8, 2024, ~80–120 participants)

What this case is

The second edition of ZuBerlin, a nine-day urban pop-up run by the team behind the 2024 edition and a ZuGarden edition in Thailand, and the shortest Zu-lineage case in this appendix. It is framed less as a proto-city than as a community of technologists, scientists and artists oriented toward mutual care.

Composition

Application screening on soft, relational criteria, curiosity and openness rather than fixed thresholds. The cohort skews toward Ethereum researchers, integrated alongside partners and children, with roughly 70 people actively supported at the busy peak (operator-reported).

Programme character

A deliberate hybrid of technical workshops across three tracks (crypto, AI and a still-thin longevity strand) and Berlin cultural programming, with artists brought in as a softening counterweight to technical rigour. Its funding is deliberately kept unconcentrated: by the organiser’s account, the aim is to avoid majority funding from any single backer and to treat sponsors more as grant-givers than as investors with rigid expectations.

Distinctive signature

One of the most comfort- and connection-oriented cases in the appendix, down to the interiors. By observation, ZuBerlin appears to draw a relatively more senior, established mix of operators and researchers than at the earlier-stage, founder-driven cases, and to run in a register of conferences, workshops and conversations that organically range from deeply technical to meditative to mundane. It keeps a deliberate distance from the network-state framing its lineage is associated with.

Outcomes and artefacts

Two ZuBerlin editions have been delivered alongside the co-developed ZuGarden edition in Thailand. The 2026 Futura Camp edition published a public series of recorded talks and panels from the event.

Lineage and continuation

Zu-lineage via ZuConnect Istanbul rather than Zuzalu directly. The 2026 edition ran as Futura Camp, a ten-day residency at Berlin's Funkhaus in June 2026.

Photo: https://x.com/flock_io/status/1935716543205372287/photo/3



Case A.7. Zuitzerland Residency

At a glance



Dates

May 1–25, 2025 (25 days)

Location

Swiss Alps, Switzerland (Rocksresort ski-village, Laax, off-season)

Venue type

Apartments; conference facilities within the Rocksresort ski-resort village

Archetype

Launchpad to Proto-jurisdiction (first edition; explicit permanent-hub vision)

Scale band

150–500 (150 approved; 190 total residents of one week or more; 300+ including summit visitors)

Themes

Tech Longevity Governance Crypto AI

Edition

First major Zuitzerland edition; Zuzalu lineage (≈40% of residents had prior Zuzalu experience)

What this case is

The first major edition of a Swiss-based multi-year programme framed by its organisers as a “Genesis Village”: a 25-day residency conceived as the pilot phase of a permanent Alpine hub, with stated funding targets of CHF 10–15M (~\$12M–\$18M) and a planned capacity of 150 to 200 residents across three to four cycles per year.

Composition

312 applications; 150 approved (48% acceptance); 190 final residents; 56 stayed the full 3.5 weeks. Applications: 25.6% Swiss, 74.4% international; 42 nationalities among residents. Roughly 20% of accepted scholars cancelled owing to visa non-issuance or refusal, with African-based recipients hardest hit (operator-reported).

Programme character

Co-founder-led with participatory experiments inside individual residencies. 238 events on the custom ZuitzApp calendar plus a build-week hackathon (35 projects, 9 prizes). Revenue per the Residency Report: CHF 216k (~\$260K), 90% tickets / 10% sponsors. Total funds received across the event and the permanent-hub runway: CHF 716k (~\$865K): tickets CHF 176k, sponsors CHF 40k, grants CHF 500k (69.8% of funds received); funds allocated CHF 551k (~\$665K). The grant-dominant mix is a clear instance of founder-philanthropist funding underwriting a field-building experiment.

Distinctive signature

The most financially transparent case in this appendix: its 2025 Residency Report is an unusually complete operator account of the residency’s finances. Its governance design also draws, uniquely among the focal cases, on Swiss subsidiarity and direct-democracy principles as stated inputs.

Outcomes and artefacts

Three pieces of infrastructure released for use beyond the residency: the open-sourced ZuitzApp event platform, ZuLink (VPN and gated wiki), and ZuBikes (on-chain bonding-curve bike-sharing). Operator post-event survey (46 responses): NPS 47; 50% rated 9+/10 likely to consider moving to a permanent hub for three months or more.

Lineage and continuation

Direct Zuzalu lineage through prior-attendee share and team-level continuity. Permanent hub planned, with ZuitzApp as an infrastructural contribution to the wider Zu-family programme.

Photograph: <https://pbs.twimg.com/media/GrTH9VtbwAA9Z55?format=jpg&name=4096x4096>



Case A.8. Zanzalu 2.0

At a glance



Dates

July 19 – August 1, 2025 (14 days)

Location

Fumba, Zanzibar, Tanzania

Venue type

Several venues across Fumba Town: a breakfast venue, a coworking space, and a pavilion for large events and presentations

Archetype

Launchpad (regional modifier)

Scale band

50–150

Themes

Crypto **Culture** **Governance** **Tech**

Edition

Second Zanzalu edition (predecessor: Zanzalu 1.0, May–June 2024); a third edition is planned for July–August 2026

What this case is

The second Zanzibar-based Zu-lineage event, run across several venues in Fumba Town, a privately developed new-city development on the Fumba peninsula that sits within a government-designated Zanzibar free economic zone.

Composition

A draw from the Zu-lineage alumni network with an Africa-region emphasis; the operator targets East African builders alongside international participants, with fellowships for selected participants.

Programme character

An unstructured co-living base drawn from the Zu-lineage network. The operator's published model frames Zanzalu as an annual pop-up city on four pillars, talent, experimentation, access and culture, with designed-in local engagement (operator-stated; outcomes not independently measured).

Distinctive signature

Among the focal cases hosted inside special zones, the one that positions itself as a cooperative guest rather than a semi-autonomous enclave, with an operator-stated ambition of a permanent hub in the town.

Outcomes and artefacts

Local engagement in the 2025 edition ran mostly through cultural and social-impact programming that the edition's web3 infrastructure sponsor organised using the operator's Zanzibar connections (sponsor-reported): a free public evening concert of leading Zanzibari musicians, with transport laid on from the city, reportedly drawing around 500 people; a funded donation of ten Raspberry Pi units to a local school; and a student hackathon that issued three social-impact grants reported to reach 90 people. The operator's own stated artefact is the "living laboratory" thesis: prototypes and playbooks meant to carry back to other African cities.

Lineage and continuation

Zu-lineage via naming and alumni network. Predecessor: Zanzalu 1.0 (May 15 – June 21, 2024). Successor planned: Zanzalu 3.0 (July 25 – August 14, 2026).

Photograph: composite of images provided by the organiser.



Case A.9. Forma Georgia

At a glance



Dates

November 15–23, 2025 (9 days)

Location

Tbilisi, Georgia

Venue type

Co-working space; conference facilities
[Operator input pending]

Archetype

Launchpad for Forma Georgia (the pop-up programme), Season for Forma City (the parent organisation, with a permanent setup in the UK)

Scale band

150–500 (approximately 150, organiser-reported)

Themes

Culture

Tech

Startup & Biz

Edition

Third Forma pop-up village (Argentina 2024, Sri Lanka 2025, Georgia 2025); first European edition

What this case is

Forma's first European edition: a nine-day pop-up village in Tbilisi for roughly 150 Solana-ecosystem founders, creators and investors, run by an operator that pairs pop-up villages with "Solana Economic Zones" in frontier markets.

Composition

Roughly 150 founders, creators and investors from the Solana ecosystem (organiser-reported). Entry was priced at 1 SOL, with free access for local residents.

Programme character

Built for young founders who push each other to build and ship, inside a social, aesthetically minded programme of co-living and co-working with café gatherings, hikes, sulfur baths, local feasts and DJ sets. The operator frames the format as a cultural and trade mission for the Solana ecosystem, letting participants "beta test what it's like living in" a country before committing to it.

Distinctive signature

One of the clearest host-side institutional engagements in the appendix: a national-agency memorandum of understanding between the Solana Foundation and GITA, Georgia's national Innovation and Technology Agency, signed and announced on camera during the village. Both sides framed it in development-economics terms, to support local talent and Georgia's growing tech sector. The ecosystem-government partnership was pursued in the register of a trade mission, not of exit.

Outcomes and artefacts

The Solana Foundation–GITA MoU is the edition's principal artefact; its commitments are not public. The operator also reports seeding Superteam Georgia via a lead found while scouting, and scouted leads in Argentina and Kazakhstan now run Superteams backing local talent with its government contacts. Programme-wide it reports \$5M+ in foreign direct investment (founder-stated, unaudited).

Lineage and continuation

In 2026 the operator based itself in Bristol, UK, running a June–July pop-up village and a month-long funded residency for 15 builders.

Photograph: <https://x.com/formacity/status/2004462019110727814/photo/4>.



Case A.10. Proof of Retreat #2

At a glance



Dates

January 15 – February 15, 2026 (32 days)

Location

Bình Thuận, Vietnam (Mui Ne)

Venue type

3–5 bedroom villas in the same villa town, self-arranged (no central housing)

Archetype

Sprint (organisers self-describe as a “pop-up village” rather than city)

Scale band

15–50 (organiser preference for 20–30 simultaneous; explicit choice not to scale past 100)

Themes

Governance

Lifestyle

Science

Tech

Edition

Second iteration; predecessor #1 ran January 7 – March 7, 2024 (61 days) at the same site

What this case is

A small, invite-only friend-network pop-up in Mui Ne, organised by a two-person team with direct Zuzalu lineage and launched after Vitalik Buterin's late-2023 call to decentralise Zuzalu following ZuConnect. Framed as “next-generation monk mode”: productivity-first, low-fee, relationally curated, with an explicit rejection of the “city” label. Seed funding of 2 ETH from a Gitcoin grant; self-funded thereafter.

Composition

No formal selection: trust is mediated through friend-of-friend referrals, with the website communicating the intentions and atmosphere of the event so people self-select in. The 2026 edition operates in the 15–50 band. The gender balance was observed to be more balanced than in comparable cases, around 60% men and 40% women.

Programme character

Relational, spontaneous governance rather than rules. No central accommodation (costs split among co-residents), no participation fees, and programming curated around who is present.

Distinctive signature

The smallest-by-design focal case. An unusually clear counter-case to the subsidy-attracts-consumers pattern: the zero-fee, no-accommodation model has produced a producer culture rather than a consumer culture.

Outcomes and artefacts

Villedge Tech, a coordination platform the organisers and contributors developed by testing and iterating in real time during the pop-up village; the “ZuWell” framework scoring pop-up villages on nature accessibility, safety, food diversity, walkability, and similar categories; a session note “(Placeholder)” produced by the research team afterward.

Lineage and continuation

Direct Zuzalu lineage via the organisers' prior attendance at Zuzalu and ZuConnect. The organisers prefer to validate the platform at small scale before broader scaling.

Photograph: composite of images from the organiser's shared album, 2026.



Case A.11. Infinita City 2025

At a glance



Dates

January 6 – March 3, 2025 (57 days; programme ongoing across additional 2025 editions)

Location

Roatán, Honduras (within the Próspera SEZ)

Venue type

Roatán island, within the Próspera SEZ host jurisdiction, Honduras

Archetype

Launchpad (short duration, high institutional ambition)

Scale band

500–1,500

Themes

Longevity Crypto Tech

Edition

Successor to the 2024 Vitalia predecessor; multiple 2025 editions on Roatán

What this case is

The 2025 flagship pop-up city edition by a longevity- and biotech-focused community and platform, with its main hub in Próspera, a semi-autonomous charter zone on Roatán, Honduras, operating its own legal and regulatory framework. Among the focal cases it sits in the most institutionally autonomous host jurisdiction, chosen for the regulatory speed a charter zone can give biomedical innovation.

Composition

Predominantly international, with a cohort skewing young and male (one attendee put the gender split at roughly 70:30). An investor described it as still early: well suited to a month-long event, though not yet built for families or a more mainstream audience. Entry ran roughly \$2,500 to \$5,000, one attendee noted, better matched to high-income frontier-tech visitors than to local Latin American incomes.

Programme character

High-filter curation by application and invitation, with participants describing unusually easy spontaneous collaboration. One attendee read the community as built more top-down than bottom-up, its activities loosely cohesive and its opportunities weighted toward founders, with little for job-seekers.

Distinctive signature

A clear fusion of a pop-up city with a permanent host jurisdiction and investor capital: a venture fund holds equity in the operator and in on-site startups, backing biotech founders to build "generational companies" alongside the community. The case is candid about the resulting tension, the investor's return running on company exits while many come for the community.

Outcomes and artefacts

The operator reported raising \$3 million to convert the pop-up into a permanent biomedical district. (March 2025) Its clearest institutional outcome since is the launch of experimental-treatment review board services in Montana, taking the model beyond the charter zone into a US state (July 2026, operator-announced).

Lineage and continuation

A cofounder spin-off from the 2024 Vitalia pop-up, relaunched under its own name in January 2025. Later editions (Crypto Cities, Bio-Frontiers Month) are catalogued in Appendix F, and in 2026 it ran Infinite Games, a two-month programme of tournaments, conferences and a startup accelerator.

Photograph: composite of <https://x.com/NiklasAnzinger/status/1888377211952259153>



Case A.12. Ârc Network School

At a glance



Dates

September 18 – October 17, 2025 (30 days)

Location

Forest City, Malaysia (Singapore-Johor Special Economic Zone; co-located with Network School)

Venue type

Hotel and co-working space at the Network School site, Forest City

Archetype

Proto-jurisdiction (charter-city aspiration)

Scale band

150–500

Themes

AI Crypto Longevity Governance

Edition

Ârc series: Lisbon 2024, Austin March 2025, and this Network School activation (September 2025)

What this case is

A month-long Ârc activation co-located with Network School in the Singapore-Johor Special Economic Zone (Forest City, Malaysia), and the organisation that anchored the site as Ârc's permanent home. It is the only genuinely nested case in the set: a startup society run as a "layer 2" inside another, leasing its own separately incorporated space within Network School.

Composition

Operator-reported at 100-plus resident members with hundreds more passing through over the month, a builder- and investor-heavy cohort overlapping the wider Ârc and Network School communities.

Programme character

More than 30 curated events across the month spanning AI, Web3, longevity and governance (operator-reported), on a founder-led, minimal-intervention model. Entry is by vouching only, members stake their reputation on those they admit, and the founder reports no interpersonal or community disputes.

Distinctive signature

The corpus's clearest bridge from the pop-up form to the charter-city tradition, and the case in which its parent organisation proactively pitches host nations. Its public pitch to governments asks for at least 5 square miles with maximum regulatory autonomy and promises tens of billions in investment toward a trillion dollars of GDP, framed as "a research lab for the world's cities", one of the most ambitious jurisdictional visions among the focal cases.

Outcomes and artefacts

Contribution to a repeatable "pop-up first, then permanent" playbook for government partnerships. Its public artefacts include developer hackathons run with Superteam Malaysia and a founder-led media effort, a podcast and a Substack.

Lineage and continuation


Ârc editions: Lisbon (2024) and Austin (2025) before the Network School activation.

Photograph: https://www.thearccity.com/_next/image?url=%2Fimages%2Fcommunity%2Fgrid-13.jpeg&w=750&q=75&dpl=dpl_BBMXqxapM4EmZvLXyMR2kjRtzbUj

Metrics dictionary.

Exhibit B.1. Metrics by tier

Tier	Metric	What it measures	Status
0	Attendance and departure rate	Daily headcount; early-departure count	Operational
0	Programming engagement rate	Session attendance against total participants	Operational
0	Organiser hours per week	Weekly time commitment per organiser	Operational
0	Return intention	Single-question post-event survey, 1 to 10	Operational
0	Budget burn against plan	Weekly variance against pre-event projection	Operational
0	Formal incident count	Incidents beyond routine logistics	Operational
1	Cross-tier interaction rate	Interactions per day per participant, by tier pair	Proposed
1	Local-community interaction frequency	Participant-to-host-community interactions per week	Proposed
1	New collaboration pairs formed	Pair count per 100 participants	Proposed
1	Post-event community persistence at 30 days	Active community members 30 days post-event as a share of peak attendance	Proposed
2	Willingness-to-repurchase	Revealed-preference test of perceived value	Proposed
2	Social network density change	Pre-vs-post acquaintance-graph density	Proposed
2	Concrete outputs and artefacts shipped	Tangible deliverable count (demo-day output count)	Operational
2	Ambient collaboration quality	Frequency and quality of unscheduled collaboration	Proposed
3	Collaboration persistence at 3, 6 and 12 months	Active proportion of post-event collaboration pairs	Proposed
3	Actual return rate	Cohort-to-next-cohort registration overlap	Operational where tracked
3	Institutional innovation transfer	Innovations transferring to other institutional contexts	Proposed



The report uses four measurement tiers. Tier 0 is the minimum the research recommends every pop-up city track from day one: six metrics, a shared spreadsheet and a weekly routine (Ch.14 Exhibit 14.2 carries the collection guidance and maps each metric to a Ch.12 failure mode and a Ch.13 decision). Tiers 1 to 3 are progressive additions for teams with more capacity. Ch.11 carries the measurement approaches and the deployment gap.

Tier 0 definitions and un-instrumented outcomes

Three definitions in the Tier 0 set need their boundaries stated once, here:

Return intention

Case means in the cases surveyed sit mostly in a 6 to 8.75 range, computed on small per-case samples in the survey recompute against the locked base (2026-07-21). The range is a reference band, not a threshold: no trigger score is defined, because the completed cases in the evidence include scores below any cutoff that has been proposed.

Organiser hours per week

No hourly red line is defined. The warning sign is Ch.8's qualitative one: a sustained, uncompensated workload well beyond a normal full-time week.

Programming engagement rate

Sustained decline is an early signal of dissatisfaction; the evidence does not support a specific lead time between decline and crisis.

The Tier 3 windows (3, 6 and 12 months) are proposed instrument windows, not measured findings; no project studied yet runs the longitudinal infrastructure that would fill them.

Outcomes named but not yet instrumented

Four outcomes recur across the interview evidence as what participants value most, and none yet has an operational instrument: serendipitous encounters, accelerated relationship depth, subjective life density, and trajectory or identity transformation. The grouping is the report's synthesis rather than a measured distribution, and the measurement gap itself is documented across at least five independent sources. Ch.11 §1 and §3 carry the evidence; Ch.14's closing section carries the roadmap.

Failure mode reference.

Three failure modes recur across the cases studied, each at a different level of analysis (the event, the organisation, the ecosystem) and each knowable in advance. Ch.12 carries the full evidence chain; this appendix carries the cross-reference grid for planning and review. For operators: the leading indicators are recognisable in advance, though the evidence does not support a specific lead time. Where several indicators for one mode are visible before the event begins, resource the corresponding institutional function then (Ch.14 Exhibit 14.1), because the shortfall otherwise tends to surface mid-event.

Exhibit C.1. The three failure modes

Mode	What drives it	Leading indicators (Ch.12 expands)	Design responses (Ch.13 decisions; Ch.14 functions)
1. Operational overwhelm (event level)	The operational surface the format demands exceeds what small founding teams have been funded to provide; a resourcing gap rather than an effort gap	A founding team of a handful of people carrying a flagship-scale operational surface; institutional functions without named owners before the event (Ch.14 Exhibit 14.1); budget burn diverging from plan in early weeks; a sustained workload well beyond a normal full-time week (Ch.8)	#11 coordination mode; #12 role clarity; #13 staffing model; #17 host-relations model
2. Institutional knowledge loss (organisation level)	Knowledge travels with people, and the artefacts built to carry it go unmaintained once the event ends (Ch.5 Finding 5); historically unaddressed, now attracting explicit counter-designs whose effectiveness is not yet measured	Post-event debrief not scheduled before the event begins; no operational documentation for any core function; retrospectives skipped under end-of-event exhaustion; handovers running on continuing personal contact rather than documents	#12 role clarity; #4 reputation portability; Ch.14 functions 6 and 7 (knowledge capture, succession)
3. The ecosystem ceiling (field level)	Capital fragmented rather than compounding, talent dispersing without a credentialing pathway, and a category mainstream institutions cannot yet read (Ch.5 Finding 7, partly by the field's own selectivity)	Funding contingent on a single anchor event or patron rather than pre-committed; continuity resting on individual funders' decisions; no training or credentialing pathway for organisers; absence from mainstream consulting coverage, as far as this research could find	#8 funding model; #18 jurisdictional strategy; record-keeping and published outcome reporting (Ch.14 Exhibit 14.2)

The register tracks further modes at their own chapters: governance-scale mismatch (Ch.7), subsidy-dependency distortion (Ch.8), and venue-access loss, documented in one project that wound down within weeks of losing access to its intended site. Host-jurisdiction rejection is documented in adjacent forms and in one pop-up city studied (A.2) (Ch.4 §4.4).

Terms used in this report.

Adjacent form

A neighbouring form that fails at least one of the four conditions: conferences, retreats, co-living operators, intentional communities, charter cities and special economic zones. Ch.1 Exhibit 1.1; Ch.4.

Ambient collaboration

The report's editorial term (not a phrase drawn from any interview) for the unscheduled productive overlap between selected, co-present participants. Ch.9, after Exhibit 9.1; see also Ch.5 Finding 1 and Ch.11 Exhibit 11.1.

Archetype

One of four configuration types in the report's classification lens, ordered by duration and institutional ambition: Sprint, Launchpad, Season, Proto-jurisdiction. A reader-facing simplification; the honest number of recurring configurations is probably higher. Ch.3.

Coordination ceiling

The scale at which coordination cost becomes the main constraint under participatory governance; the available evidence places it somewhere in the low hundreds of participants, a working hypothesis anchored by a single comparison rather than a benchmark. Ch.9, after Exhibit 9.1; Ch.5 Finding 6.

Depletion ceiling

The bound on uncompensated intrinsic and social recognition before contribution stops being sustainable. Ch.13 Exhibit 13.1, decision 10; mechanism in Ch.8.

Four conditions

Daily co-presence for the full duration; a fixed start and a fixed end; selection-generated trust from an explicit entry filter; a shared intention to work on something substantive together. All four must hold at once. Ch.1 Exhibit 1.1.

Hypothesis register

The hypotheses the research tested, each stated in full with its status in Appendix E; two were split into halves that earned separate verdicts (the paired entries such as H19a and H19b).

Hypothesis status

Each hypothesis graded on the weight of evidence: supported, qualified, contested, or under assessment. Current counts in the Executive Summary sidebar; the full register in Appendix E.

Network state

Srinivasan's coinage (2022): "a highly aligned online community with a capacity for collective action that crowdfunds territory around the world and eventually gains diplomatic recognition from pre-existing states". Ch.4 §4.2 and §4.3.

Patron capital

Mission-aligned funding that closes the gap between what most models consume and what they generate; in the mid-scale cases studied, often a single committed backer. Finding 2; Ch.10.

Pop-up city

A temporary physical gathering in which people with shared interests or values build a short-lived, city-like environment; satisfies the four conditions. Ch.1.

Primary corpus

The 20 cases studied closely (12 consented focal cases and 8 comparators), the referent of the report's "20 cases" claims. Appendix E.

Methodology.

Case selection and data period

The research is a case-based interpretive study, not a probability sample. The primary corpus is 20 cases: 12 focal cases profiled in Appendix A, plus 8 comparators carried under ORG codes because they did not pass the focal-profile consent process.

Cases span the report’s two classification lenses: the four-archetype map of Ch.3 (Sprint, Launchpad, Season, Proto-jurisdiction) and the cost-type grouping of Ch.10 and 13, whose scale bands are flagship (400+ participants, multi-week, rented venue), residency (50 to 80, 4 to 8 weeks, curated), permanent anchor (30 to 70, seasonal, owned space) and large-scale urban (1,000+, short duration).

Coverage spans North America, Europe, Asia, Africa and Latin America; a wider catalogued pool supports landscape-level patterning. Interview and survey data cover November 2025 to April 2026; desk research and case updates are coded through July 2026; entries are dated by event date.

Data sources

Source	Count	Notes
Coded interviews	92	41 primary-source + 51 coded-from-notes; 92 unique IDs
Valid survey responses	86	After excluding duplicates, test and researcher submissions, spam, and empty partials; the recruitment shortfall against target is disclosed under Positionality and disclosures
Coded desk research	337	Fully coded sources only; coded through July 2026
Theme memos	22	One per hypothesis
Primary corpus	20	12 focal + 8 comparators; focal profiles in Appendix A
Catalogued pool	107 (102 definitional + 5 boundary)	Appendix F directory as of May 21, 2026

Hypothesis register.

The hypotheses were re-adjudicated in July 2026 over the full evidence base, superseding interim statuses in earlier working drafts that rested on the first 39 coded interviews. Two were split into halves that earned separate verdicts, which is why the table below pairs entries such as H19a and H19b; entry and status totals are given in the Executive Summary’s evidence-base sidebar. Each hypothesis appears below in one line with its current status, so that every code stamped on a chapter opener or used in prose resolves to a statement; the full statements, evidence bases and adjudication history are maintained in the reviewer materials. One register discipline is disclosed here because a chapter relies on it: the “sustained” qualifier on the incentive hypothesis (Ch.8) was added after contrary short-burst evidence arrived, and it scopes the claim rather than predicting it.

Code	Hypothesis	Status	Ch.
H1	Pop-up cities are a distinct institutional form defined by specific necessary conditions, not reducible to conferences, co-living, special economic zones or intentional communities; the time-bounded definition holds for the period studied.	Qualified	Ch.1
H2	Pop-up cities cluster into recurring archetypes that differ meaningfully in governance, economics and outcomes, classified per configuration per iteration; the four-type scheme is a reader-facing simplification and the empirical cluster count remains open.	Qualified	Ch.3
H3	Selection stringency trades off measurably against community legitimacy, and stringency operates as a portfolio of filter types rather than a single dial, with legitimacy costs attaching differently across filter types.	Qualified	Ch.6
H4	Permeability of participation boundaries, whether tiers or gradients, is the operative predictor of contribution quality and inclusion; whether multi-tier models outperform single-tier models remains an open question.	Qualified	Ch.6
H5	Portable identity and reputation is scaling infrastructure whose value is stage-dependent: informal networks carry selection cheaply at current scale, and formal systems become necessary for selection quality and cost only as cross-event volume outgrows those networks.	Qualified (scaling clause forward-looking, under assessment)	Ch.6
H6	Founder-led governance is the observed default across formats, tracking stakes, scale and accountability design rather than duration, and inversion toward participatory governance occurs only where explicitly designed in advance.	Qualified	Ch.7

Hypothesis register.

Code	Hypothesis	Status	Ch.
H7	Legitimacy is earned mainly through non-process currencies (perceived competence, founding authority, consent through voluntary entry and exit, narrative coherence, contribution and recruitment volume) rather than through participatory process.	Qualified	Ch.7
H8	On-chain governance mechanisms introduce more failure modes than they resolve in collective-decision uses such as token voting, delegation and treasury allocation; scoped instrumental uses show mixed-to-workable records and sit outside the claim.	Supported	Ch.7
H9	Subsidy without accountability mechanisms distorts participant behaviour toward consumer modes and leaves continuity resting on funders' decisions when funding disappears; the operative variable is accountability structure, not the proportion of external revenue.	Qualified	Ch.8
H10	Intrinsic and social incentives are stronger predictors of sustained contribution than financial incentives, with the sustained qualifier doing essential work.	Supported	Ch.8
H11	The minimum viable institutional capacity threshold is higher than most organisers expect, particularly for safety and host-jurisdiction relations.	Supported	Ch.9, 14
H12	Coordination costs scale superlinearly under participatory governance but linearly under delegated governance, creating a ceiling on participatory models.	Supported	Ch.9
H13	Institutional knowledge loss between iterations is a significant and well-documented failure mode, operating through team turnover, absent feedback systems, narrative gaps and ecosystem-level fragmentation.	Supported	Ch.12
H14	Cost per participant-week varies more than fivefold across archetypes, driven by venue, accommodation and staffing intensity rather than programming quality.	Qualified	Ch.10
H15	No current flagship or large-format model achieves financial sustainability without premium pricing, founder subsidy, token treasury, tax or regulatory arbitrage, or grants; small permanent-space models are a partial, scale-capped third path.	Supported	Ch.10
H16a	Where no external value capture is structured in from the outset, host jurisdictions under-value pop-up cities' institutional learning and reputational contributions relative to the metrics they collect.	Qualified	Ch.4

Hypothesis register.

Code	Hypothesis	Status	Ch.
H16b	Where ventures are structured from the outset to capture value externally, host jurisdictions accurately identify net-negative impacts, making the undercounting frame inapplicable.	Supported	Ch.4
H17	The outcomes participants value most are consistently the outcomes current measurement approaches fail to capture.	Supported	Ch.11
H18	The three most heavily evidenced failure modes (operational overwhelm, governance-scale mismatch, subsidy dependency) all show identifiable leading indicators that precede crisis, though no specific lead time is evidenced.	Qualified	Ch.12
H19a	Pop-up cities function as governance R&D environments whose innovations transfer through overlapping personnel, and transfer fails without a receiving structure in the adopting institution.	Supported	Ch.7
H19b	Institutional innovations transfer through artefacts (tools, documents, protocols) independently of personnel; design-stage support and outcome-level counter-evidence currently divide, and no documented cross-community adoption exists.	Contested	Ch.12
H20	The network-state framing attracts funding and talent but creates legitimacy risk with host jurisdictions that read sovereignty claims as threatening.	Supported	Ch.4
H21	The pop-up city ecosystem is structurally pre-tipping-point, lacking the capital density, talent density and narrative category required for a self-reinforcing adoption curve.	Supported	Ch.12
H22	Selection-based trust production structurally excludes the host community unless the operating model prices locals in by design, making exclusion the default rather than an oversight.	Assess	Ch.6

Anonymisation and positionality

Confidence grading

High – 3 or more independent sources drawn from at least 2 source streams, or direct measurement.

Medium – 2 independent sources, proxy measurement, or 40 to 60% case coverage.

Low – 1 source, under 40% case coverage, or inference-heavy reading; retained where central to a finding and flagged explicitly.

Anonymisation and consent

Organisations referenced in chapter prose are coded ORG-XXX; persons are referenced by role, not by code. Anonymisation protects interview participants and figures in primary-source operational archives. Role-based descriptors may remain identifying to field insiders; the consent protocol accepts that bound. Authors of published works cited as desk research are named normally with title and year. Real names appear in the Appendix A focal profiles (consent-gated), in the Appendix F case directory (drawn from already-public case names), and in the acknowledgements (consent-gated). Each interview was conducted with informed consent covering anonymisation, quotation and retention, and a re-identification audit was conducted by an embedded reviewer before lock.

Positionality and disclosures

The research team is embedded in the pop-up city ecosystem, and the report is written from informed proximity rather than external distance: the advantage is access; the risk is ecosystem bias. The general mitigations are the hypothesis register maintained as a dated, versioned analytical frame, the codebook governing every coded excerpt, and the re-identification audit above.

The report is published as a versioned edition with the intent to iterate; feedback received after publication informs subsequent editions under the versioning policy stated in the back matter.

Limitations

- Non-random sampling – the interview and survey base draws on organiser networks accessible to the research team; participant-side findings describe a young, mobile, visa-advantaged, largely Western-facing population.
- Archetype coverage – no Season-archetype focal case; Season-form evidence rests on the wider catalogued pool.
- Survivorship bias – cases that launched are over-represented; pre-launch failure modes are likely under-counted.
- Perception-based measurement – trajectory and outcome claims rest on retrospective self-report.
- Coordination substrate – in at least one well-documented case, real coordination happened in chat channels lost to retrospective analysis (Ch.9).
- Low case numbers – the focal case database is approximately 50% populated, the survey ran across 21 forms, 6 of which received no submissions, and quantitative cross-case comparisons are directional.
- Source dependency – the host-side analysis in Ch.10 leans on one local-host informant for the formal interview record, as disclosed above; observational and informal sources corroborate the pattern from other angles.

Methodology questions and source-traceability requests can be directed to the research team.

Case directory.

Case	Location	Dates (days)	Approx. scale	Themes
Zuzalu ●	Luštica Bay, Montenegro	Mar 25 – May 25, 2023 (62)	500–1,500	GovernanceLongevityScienceAITech
ShanHaiWoo Beidahu	Jilin, China	Aug 5 – 31, 2023 (27)	150–500	CryptoCultureTech
Burning Man 2023	Black Rock City, USA	Aug 27 – Sep 4, 2023 (9)	10,000+	CultureLifestyle
muChiangMai	Chiang Mai, Thailand	Sep 15 – Oct 31, 2023 (47)	150–500	CryptoCultureGovernanceTech
Web3 Village Hoi An	Hoi An, Vietnam	Oct 24 – 30, 2023 (7)	150–500	CultureTech
ZuConnect	Istanbul, Turkey	Oct 29 – Nov 11, 2023 (14)	150–500	TechGovernanceCryptoAI
WAMO2023	Chiang Mai, Thailand	Dec 16, 2023 – Jan 1, 2024 (17)	150–500	CultureCryptoGovernanceAI
Vitalia	Roatán, Honduras	Jan 6 – Mar 1, 2024 (56)	150–500	Longevity
Proof of Retreat #1	Bình Thuận, Vietnam	Jan 7 – Mar 7, 2024 (61)	15–50	TechCultureCrypto
muBuenos	Buenos Aires, Argentina	Mar 15 – Apr 28, 2024 (45)	500–1,500	CryptoTech
Forma Argentina	Buenos Aires, Argentina	Mar 21 – May 4, 2024 (45)	150–500	CryptoTechStartup & BizGovernance
Zanzalu 1.0	Zanzibar, Tanzania	May 15 – Jun 21, 2024 (38)	150–500	GovernanceCultureCrypto
Edge Esmeralda 2024	California, USA	Jun 2 – 29, 2024 (28)	500–1,500	CultureScienceTech
muAccra	Accra, Ghana	Jun 9 – 22, 2024 (14)	150–500	CryptoCultureGovernanceTech
ZuVillage Georgia	Tusheti, Georgia	Jul 20 – Aug 31, 2024 (43)	150–500	AILongevity
Aleph, Ciudad de Crecimiento	Buenos Aires, Argentina	Aug 5 – Sep 1, 2024 (28)	500–1,500	CryptoCulture
Burning Man 2024 ●	Black Rock City, USA	Aug 25 – Sep 2, 2024 (9)	10,000+	CultureLifestyle

Case	Location	Dates (days)	Approx. scale	Themes
Zuitzerland v0.1	Interlaken, Switzerland	Sept 2024	15–50	Tech Governance Culture
ZuBerlin 2024	Berlin, Germany	Sep 1 – 8, 2024 (8)	50–150	AI Culture Governance Tech
Ârc Lisbon	Lisbon, Portugal	Sep 10 – Oct 10, 2024 (31)	150–500	Startup & Biz Culture Governance Tech
Digitalia	Ascoli Piceno, Italy	Sep 14 – Oct 14, 2024 (31)	150–500	Crypto Culture Tech
Network School v1	Forest City, Malaysia	Sep 23 – Dec 23, 2024 (92)	50–150	Crypto Governance Tech
IslandDAO Koh Samui	Koh Samui, Thailand	Sep 25 – Oct 25, 2024 (31)	150–500	Crypto Governance Tech
Invisible Garden Chiang Mai	Chiang Mai, Thailand	Sep 30 – Nov 10, 2024 (42)	150–500	Crypto Science Tech
ShanHaiWoo Chiang Mai	Chiang Mai, Thailand	Sep 30 – Nov 10, 2024 (42)	150–500	AI Crypto Culture
Chiang Mai Tech City	Chiang Mai, Thailand	Oct 7 – Nov 5, 2024 (30)	150–500	Crypto Tech
MegaZu	Chiang Mai, Thailand	Oct 7 – Nov 7, 2024 (32)	50–150	Crypto Culture Governance Tech
Zelar City	Berlin, Germany	Oct 7 – Nov 17, 2024 (42)	150–500	Longevity Science
Edge City Lanna	Chiang Mai, Thailand	Oct 10 – Nov 10, 2024 (32)	500–1,500	Culture Science Tech
AuraVerse	Chiang Mai, Thailand	Oct 12 – Nov 10, 2024 (30)	50–150	Longevity Science Tech
Web3 Village Chiang Mai	Chiang Mai, Thailand	Oct 21 – 30, 2024 (10)	500–1,500	Culture Tech
CAMPUS	Punta Cana, Dominican Republic	Nov 5 – 8, 2024 (4)	150–500	Tech Culture
Web3 Village Bangkok	Bangkok, Thailand	Nov 10 – 17, 2024 (8)	500–1,500	Tech Science Culture
ZuGarden	Bangkok, Thailand	Nov 10 – 30, 2024 (21)	50–150	Crypto Culture Tech
ZuThailand	Chon Buri, Thailand	Nov 18 – Dec 18, 2024 (31)	150–500	Tech Culture
Aleph de Verano	Buenos Aires, Argentina	Dec 9 – 20, 2024 (12)	500–1,500	Crypto Governance
Network School v2 ●	Forest City, Malaysia	Mar 1, 2025 – Feb 28, 2026 (365)	150–500	Crypto Governance Tech
Infinita City 2025 ●	Roatán, Honduras	Jan 6 – Mar 3, 2025 (57)	500–1,500	Longevity Crypto Tech
Zu-Grama	Kerala, India	Jan 6 – Feb 16, 2025 (42)	150–500	AI Longevity Crypto Governance
Mars College 2025	California, USA	Jan 13 – Apr 13, 2025 (91)	50–150	Tech Governance AI
ETHiopia Pop-up	Addis Ababa, Ethiopia	Jan 21 – Feb 21, 2025 (32)	150–500	Crypto Culture Governance Tech
Ârc Austin	Texas, USA	Mar 2 – 30, 2025 (29)	500–1,500	AI Governance Longevity

Case	Location	Dates (days)	Approx. scale	Themes
Aleph March '25	Buenos Aires, Argentina	Mar 5 – 28, 2025 (24)	1,500–10,000	AI Crypto Culture
Forma Sri Lanka	Ahangama, Sri Lanka	Mar 21 – Apr 4, 2025 (15)	150–500	Crypto Tech
Vitalist Bay	California, USA	Apr 4 – May 29, 2025 (56)	500–1,500	Longevity Science AI
Liberland deSci & deSpace	Apatin, Serbia	Apr 14 – May 11, 2025 (28)	50–150	Crypto Governance Longevity Science
ZuAfrique 1.0	Kilifi, Kenya	Apr 15 – May 6, 2025 (22)	150–500	Crypto Culture Tech
Ipê Village I	Santa Catarina, Brazil	Apr 22 – May 22, 2025 (31)	150–500	AI Crypto Governance Tech
Zuitzerland Residency ●	Swiss Alps, Switzerland	May 1 – 25, 2025 (25)	150–500	Tech Longevity Governance Crypto AI
Edge Esmeralda 2025 ●	California, USA	May 24 – Jun 21, 2025 (29)	500–1,500	Culture Science Tech
Crypto Cities	Roatán, Honduras	Jun 1 – 30, 2025 (30)	500–1,500	Crypto Governance Tech
ZuBerlin 2025 ●	Berlin, Germany	Jun 14 – 22, 2025 (9)	50–150	AI Culture Governance Tech
Viva Frontier Tower	California, USA	Jun 20 – Aug 4, 2025 (46)	150–500	AI Longevity Crypto Tech
Oz City Cannes	Cannes, France	Jun 21 – 29, 2025 (9)	50–150	AI Crypto Tech
Zanzalu 2.0 ●	Zanzibar, Tanzania	Jul 19 – Aug 1, 2025 (14)	50–150	Crypto Culture Governance Tech
Burning Man 2025	Black Rock City, USA	Aug 24 – Sep 1, 2025 (9)	10,000+	Culture Lifestyle
Aleph Festival	Buenos Aires, Argentina	Aug 25 – Nov 22, 2025 (90)	1,500–10,000	Crypto Governance Startup & Biz
Bio-Frontiers Month	Roatán, Honduras	Sep 1 – 30, 2025 (30)	150–500	Longevity Science Tech
ShanHaiWoo Singapore	Singapore	Sep 1 – 30, 2025 (30)	150–500	Crypto AI Tech Culture
traditionaldreamfactory	Alentejo, Portugal	Sep 1 – 30, 2025 (30)	50–150	Culture Crypto Public Goods Tech
ZuKas I	Kaş, Turkey	Sep 9 – 19, 2025 (11)	150–500	Governance Crypto Culture Tech
ZuCity Japan 2025	Nagano Prefecture, Japan	Sep 16 – 30, 2025 (15)	50–150	Longevity Culture Tech
IslandDAO Mykonos	Mykonos, Greece	Sep 17 – Oct 17, 2025 (31)	150–500	Crypto Tech Startup & Biz
Ârc Network School ●	Forest City, Malaysia	Sep 18 – Oct 17, 2025 (30)	150–500	AI Crypto Longevity Governance
Kuu Village	Nara, Japan	Oct 1 – 7, 2025 (7)	50–150	Culture Governance Tech
Urbe Village	Rome, Italy	Oct 5 – 19, 2025 (15)	150–500	Crypto Culture Tech Startup & Biz
REcity	Recife & Olinda, Brazil	Oct 6 – 26, 2025 (21)	150–500	Public Goods Science Tech Governance

Case	Location	Dates (days)	Approx. scale	Themes
Solana Próspera	Roatán, Honduras	Oct 11 – Nov 10, 2025 (31)	150–500	Crypto Longevity Tech
Edge City Patagonia ●	San Martín de los Andes, Argentina	Oct 18 – Nov 15, 2025 (29)	500–1,500	Tech Culture Startup & Biz Science
Invisible Garden Buenos Aires	Buenos Aires, Argentina	Oct 27 – Nov 16, 2025 (21)	50–150	AI Crypto Science Tech Public Goods
Oz City Patagonia	San Martín de los Andes, Argentina	Nov 1 – 16, 2025 (16)	50–150	AI Crypto Tech Culture
Dev3pack Village	Buenos Aires, Argentina	Nov 3 – 24, 2025 (22)	50–150	Crypto Tech
Forma Georgia ●	Tbilisi, Georgia	Nov 15 – 23, 2025 (9)	150–500	Culture Tech Startup & Biz
ETHChiangmai	Chiang Mai, Thailand	Dec 8, 2025 – Feb 3, 2026 (58)	150–500	Crypto Tech
Island Time	San Juan, Puerto Rico	Jan 2 – Mar 3, 2026 (61)	50–150	AI Crypto Longevity Culture
Mars College	California, USA	Jan 5 – Apr 27, 2026 (113)	50–150	Tech Governance AI
Proof of Retreat #2 ●	Bình Thuận, Vietnam	Jan 15 – Feb 15, 2026 (32)	15–50	Governance Lifestyle Science Tech
WAMO2026	Chiang Mai, Thailand	Jan 21 – Feb 4, 2026 (15)	150–500	Culture Crypto Governance AI
Planetir Pop-up	Ahangama, Sri Lanka	Jan 23 – Feb 23, 2026 (32)	50–150	Tech Culture
Open Source Village 2026	Brussels, Belgium	Jan 26 – Feb 6, 2026 (12)	150–500	Tech AI
Tech x Culture Pop-up City	Buenos Aires, Argentina	Feb 13 – 28, 2026 (16)	150–500	Tech Culture AI Crypto
Cyberia Genesis Village	Bali, Indonesia	Feb 17 – Mar 20, 2026 (32)	150–500	AI Longevity
ZuCity Japan 2026	Nagano Prefecture, Japan	Mar 2 – 31, 2026 (30)	50–150	Culture Governance
Aleph March '26	Buenos Aires, Argentina	Mar 4 – 26, 2026 (23)	1,500–10,000	AI Crypto Tech
Solana Network State	Forest City, Malaysia	Apr 1 – May 15, 2026 (45)	150–500	Crypto AI Tech
Ember	Hoi An, Vietnam	Apr 1 – 30, 2026 (30)	15–50	Culture Tech Governance
ProtoVille: The UBI Living Lab	Gyeongsangnam-do, South Korea	Apr 2 – 8, 2026 (7)	15–50	Governance Culture
House of Games	Alentejo, Portugal	Apr 5 – May 3, 2026 (29)	50–150	Culture Tech AI
Ipê Village II	Santa Catarina, Brazil	Apr 6 – May 1, 2026 (26)	150–500	Crypto Governance Tech

Case	Location	Dates (days)	Approx. scale	Themes
ZuCity Cherry Blossom Season	Nagano Prefecture, Japan	Apr 6 – 26, 2026 (21)	50–150	Tech Culture Longevity
ZuKas II	Kaş, Turkey	Apr 10 – May 10, 2026 (31)	150–500	Governance Crypto Culture Tech
ZuAfrique 2.0	Kilifi, Kenya	Apr 12 – May 3, 2026 (22)	150–500	Crypto Culture Startup & Biz
Ethereum NS	Forest City, Malaysia	May 1 – 31, 2026 (31)	150–500	Crypto Tech Governance
muShanghai	Shanghai, China	May 10 – Jun 6, 2026 (28)	500–1,500	AI Longevity Tech
ETH Cinco de Mayo	Puerto Escondido, Mexico	May 11 – 18, 2026 (8)	500–1,500	Crypto Governance Tech AI
NetX State Pop Up Island	Roatán, Honduras	May 18 – 31, 2026 (14)	150–500	Crypto Tech Longevity Governance
Edge Esmeralda 2026	California, USA	May 30 – Jun 27, 2026 (29)	500–1,500	Culture Science Tech
Freo Neuhaus Pop-up	Fremantle, Australia	May 31 – Jun 14, 2026 (15)	50–150	Tech Culture
Futura Camp	Berlin, Germany	Jun 13 – 22, 2026 (10)	50–150	Crypto Culture Tech
Forma Bristol	Bristol, UK	Jun 27 – Jul 5, 2026 (9)	150–500	Crypto Culture Tech Startup & Biz
The Portal at Iceland Eclipse	Snæfellsnes Peninsula, Iceland	Jul 19 – Aug 15, 2026 (28)	150–500	Governance Culture
Zanzalu 3.0	Zanzibar, Tanzania	Jul 25 – Aug 14, 2026 (21)	150–500	Culture Governance Tech
Zō Village	Niigata, Japan	Aug 2 – 31, 2026 (30)	50–150	Tech Culture Startup & Biz
Valley of the Commons	Austrian Alps, Austria	Aug 24 – Sep 20, 2026 (28)	50–150	Tech Governance
Burning Man 2026	Black Rock City, USA	Aug 30 – Sep 7, 2026 (9)	10,000+	Culture Lifestyle
Cosmo Local CNX	Chiang Mai, Thailand	Sep 1 – 28, 2026 (28)	150–500	Culture AI Startup & Biz
Revillage	Montbard, France	Sep 1 – 30, 2026 (30)	50–150	Culture

Contributors, citation, licence

Contributors and acknowledgements

This report was researched and written by Jimin Lee, with assistance from Denisa Lepădatu and Nura Lee. It rests on the willingness of hundreds of organisers, interviewees, survey participants and focal case organisations to share what worked, what broke and what they wished they had known earlier, and on the wider ecosystem of operators, community builders, host institutions, writers and funders whose work shapes the field this report studies. Special thanks to the xyz.city team for their help and their public API, whose pop-up city and network-state data supports the case directory. Participation was given under anonymised consent, and individual names appear only in the Appendix A profiles and the Appendix F directory. Errors of interpretation are the research team's alone. The research was independently funded by Thou Ārt Research.

Citation

Pop-up City Research Initiative. (2026). *Pop-up City Report* (1st ed., v1.0). Thou Ārt Research. With author attribution: Lee, J. (2026). *Pop-up City Report* (1st ed., v1.0). Pop-up City Research Initiative, Thou Ārt Research.

Licence and imprint

Pop-up City Report, first edition (v1.0), published July 30, 2026 by the Pop-up City Research Initiative and Thou Ārt Research.
© 2026.

The text and exhibits are licensed under Creative Commons Attribution 4.0 International (CC BY 4.0; full terms at creativecommons.org/licenses/by/4.0). The interview corpus, coded excerpts, identifier registry and consent-gated focal case materials are not covered by the licence and remain reserved under the consent protocol described in Appendix E. Patch versions (v1.0.x) cover copy-edit fixes, minor versions (v1.x) factual corrections or status changes, and a major version accompanies a new edition with fresh fieldwork; citers should specify the version where a change materially affects what is cited. Research window November 2025 to July 2026; interview and survey corpora closed April 30, 2026; desk sources coded through July 27, 2026; case directory updated May 21, 2026.

Citation queries, source verification, errata and peer review: research@thouartofficial.com. Press and partnerships: media@thouartofficial.com.