CERN IdeaSquare Journal of Experimental Innovation, 2026; 10(2): 69-81
DOI: https://doi.org/10.23726/cij.2026.1866

ORIGINAL ARTICLE

Co-Creation Under Time Constraints: Designing and Evaluating a Conference-Compatible Collaboration Artifact in the ATTRACT Ecosystem

Antonio Bonucci[1] Rita Graceffa[1], Mireia Sierra Andrés[2*]; Matteo Vignoli[3]

[1]European XFEL, Holzkoppel 4, 22869 Schenefeld Germany; [2]Fusion Point (Esade Business School), Avinguda de la Torre Blanca, 59, 08172 Sant Cugat del Vallès, Barcelona, Spain; [3]University of Bologna, Via Zamboni, 33, 40126 Bologna, Italy
*Corresponding author: mireia.sierra@esade.edu

ABSTRACT

Big-Science Research Infrastructures generate advanced technologies, but translating them into societal value remains difficult. Co-creation can support this transition, yet in time-boxed settings such as conferences its outputs are often hard to compare, evaluate, and reuse. This paper presents a conference-compatible co-creation artifact developed through Design Science and piloted at the ATTRACT Final Conference with 25 multidisciplinary teams. The artifact combines structured facilitation, a standardised submission form, and post-hoc synthesis. Outputs were assessed through four proxy indicators: cross-resource integration, pathway specificity, feasibility signals, and evidence/validation intent. All teams completed the instrument within the available time, and 64% mentioned at least three individual resources, and 44% combined resources from at least three of the four ecosystem categories. Results show a clear gradient from strongly articulated impact pathways (7.68/8) to weaker implementation conditions (6.00/8) and validation planning (5.28/8). At portfolio level, the analysis revealed recurring resource combinations and translation gaps beyond proof-of-concept. The study contributes a reusable method for turning time-boxed co-creation into comparable, traceable material for portfolio learning in research-infrastructure innovation ecosystems.

Keywords: Design Science; co-creation; participatory design; innovation ecosystems; stakeholder collaboration; research infrastructures; evaluation logic; deep tech; socio-technical systems; ATTRACT.

Received: March 2026. Accepted: August 2026.

INTRODUCTION

Translating scientific excellence into societal benefits is a key challenge for European Big-Science Research Infrastructures (BSRIs) (Wareham et al., 2022). These infrastructures produce advanced “deep tech” with potential applications beyond their original scientific aims. However, translating these technologies into tangible societal impact is difficult due to misaligned priorities, technological risks, market uncertainties, and long development timelines (Wareham et al., 2022). We examine ATTRACT, a European deep-tech ecosystem connecting R&D&I projects with BSRIs and complementary capacities, including socio-economic studies (SES) and the ATTRACT Academy, which funded numerous proof-of-concept projects (Wareham et al., 2022). These infrastructures have achieved strong technical results but continue to struggle to translate them into societal benefits. Addressing this gap requires more than technological advancement; it calls for organised collaboration through which stakeholders establish shared pathways that no single actor can develop independently.

Co-creation, understood as a collaborative process through which multiple stakeholders identify problems, develop solutions, and shape outcomes, can support innovation by bringing together different forms of expertise and resources. Yet, its terminology is often used imprecisely. As von Busch and Palmås (2023) highlight, participatory design practices can conceal power imbalances and lead to outcomes that are not necessarily advantageous. The participatory design tradition has long recognised these issues (Smith et al., 2025). Structure and standardisation cannot resolve these power dynamics; indeed, any instrument necessarily shapes which contributions become visible and which interpretations are retained. Nevertheless, a related operational problem persists in Big-Science contexts: the proposed beneficiaries, resource combinations, implementation dependencies, and validation intentions generated by diverse participants often remain implicit or difficult to compare. This challenge is especially evident during conferences, where teams share a limited common language or history. Making these elements visible and comparable can provide a basis (although not a sufficient condition) for co-creation outputs to inform decisions beyond the event itself.

ATTRACT is a European research and innovation programme linking funded deep-tech R&D&I projects, large-scale research facilities, socio-economic studies (SES), and the ATTRACT Academy. The Academy involves university students applying Challenge-Based

Innovation methods to human-centred problem framing, interdisciplinary ideation, and early validation of funded deep-tech projects. We conceptualise ATTRACT as an innovation ecosystem in which performance depends on an evolving set of actors, activities, artifacts, and their interrelations, including both complementary and substitute relations (Granstrand & Holgersson, 2020). The ecosystem perspective directs attention to the resources and relationships on which innovation depends, while the mission-oriented perspective asks how these configurations can be organised around societally defined challenges. Within this ecosystem, mission-oriented innovation policy offers a useful framing: missions call for policy tools focused on market co-creation and shaping, not only on market fixing (Mazzucato, 2018). Yet, mission-oriented ambition does not automatically yield actionable pathways. The core challenge is to build mechanisms that help stakeholders jointly articulate plausible deep-tech-to-impact pathways and, crucially, ensure that the resulting outputs are comparable across a portfolio. Because measuring final societal impact is itself contested and methodologically challenging (Bornmann, 2013), we focus on intermediate, traceable outputs—articulated pathways, resource combinations, feasibility signals, and validation intentions—that can be captured quickly and compared across teams.

This paper contributes a conference-compatible socio-technical co-creation artifact and an associated evaluation logic designed to generate comparable intermediate outputs suitable for portfolio-level synthesis. The artifact was piloted in a workshop at the ATTRACT Final Conference (Brussels, 2–3 July 2024) and comprises a structured facilitation protocol, a standardised data-capture instrument, and a synthesis routine, framed as a Design Science (DS) intervention (Romme, 2023; Simon, 1969). It is summarised in Figure 1 and described in full in the Method section; the complete protocol and form are provided in the Appendix.

The intended users of the instrument are programme managers, facilitators, and workshop designers seeking to organise and evaluate short co-creation processes in research infrastructure and other deep-tech ecosystems; no prior familiarity with ATTRACT is assumed.

In this paper, we address two research questions:

RQ1 (design): How can a structured, time-boxed co-creation protocol enable temporary teams to produce complete, comparable, and analysable intermediate outputs in a conference setting?

RQ2 (evaluation): What patterns across these outputs are revealed by four bounded proxy indicators: cross-resource integration, pathway specificity, feasibility signals, and evidence and validation intent?

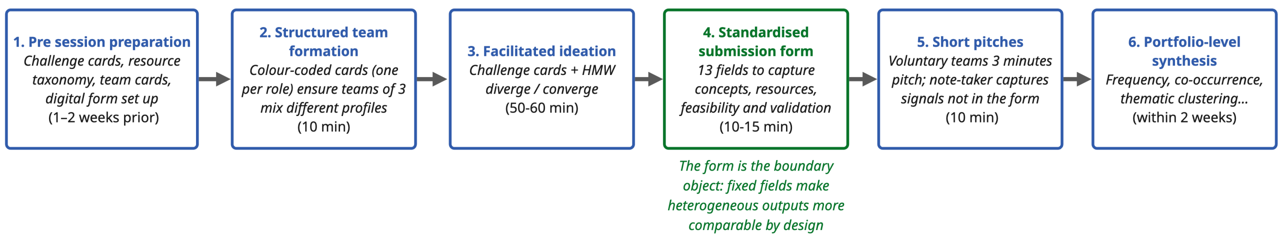

Fig. 1. The artifact as a six-step workshop protocol (~90 min in session), from pre-session preparation to portfolio-level synthesis. The standardised submission form (step 4) is the boundary object whose fixed fields make heterogeneous team outputs comparable by design.

THEORETICAL BACKGROUND

Co-creation and what makes outputs evaluable

In design research, co-creation (and co-design) emphasises moving beyond user participation as "feedback" towards active joint making and sensemaking with affected parties (Sanders & Stappers, 2008). Across adjacent fields, co-creation is also treated as a construct with operationalisable dimensions. Ranjan and Read (2016), for example, distinguish co-production—the joint production of resources and offerings—from value-in-use, which refers to the value realised by beneficiaries in context.

A persistent gap is that co-creation outcomes are frequently discussed but less often evaluated using comparable indicators. A systematic review of co-creation and co-production in public innovation highlights this imbalance and calls for greater attention to outcomes and more quantitative approaches (Voorberg et al., 2015). In deep-tech ecosystems, this challenge is intensified because the final societal impact is long-term, contested, and difficult to attribute (Bornmann, 2013). Articulated pathways, validation intent, and resource combinations can be captured during the activity and compared across teams, whereas societal impact cannot be assessed within the same timeframe. We therefore aim to evaluate intermediate, traceable outputs rather than ultimate impact.

Standardisation and quantification are means rather than ends, and they are not the only legitimate responses to the evaluation challenge. Comparability can flatten differences and obscure the political character of design choices; interpretation and situated judgement therefore remain indispensable (Smith et al., 2025). Impact-assessment debates similarly emphasise that quantitative indicators capture only part of what matters and can crowd out qualitative understanding when treated as sufficient (Bührer et al., 2022). Our position is therefore bounded: we use lightweight indicators to produce time-boxed outputs that are sufficiently comparable to support portfolio-level learning, not to adjudicate their worth.

Co-creation under time constraints in temporary teams

Conference workshops assemble temporary teams with heterogeneous expertise and limited shared history. Such teams rely on 'swift trust'—a fragile form of trust enabled by role clarity, shared task focus, and time pressure rather than long-term relational development (Meyerson et al., 1996). Time itself becomes a coordinating mechanism: Gersick's (1988) punctuated equilibrium model shows that deadlines trigger rapid shifts in a group's approach, concentrating sensemaking and convergence around midpoints and endpoints.

Hackathons are a close analogue: short, intensive events that generate many ideas quickly, but whose outcomes are hard to evaluate and sustain (Angarita & Nolte, 2020). Together, these streams frame the design problem at the heart of this paper: how to preserve the generative benefits of time pressure while producing outputs structured enough to be comparable and reusable.

Boundary objects and standardisation for comparability

A core mechanism for co-creation across heterogeneous participants is the *boundary object*—an artifact adaptable to different viewpoints yet robust enough to hold a common identity across them (Star & Griesemer, 1989). Boundary objects support cooperation by providing shared reference points without requiring full consensus, and Star and Griesemer (1989) specifically identify 'standardised *forms*' as boundary objects that make inputs comparable across communities.

Building on this tradition, Carlile (2002, 2004) distinguishes three kinds of knowledge boundary: syntactic boundaries, arising from differences in language and format; semantic boundaries, arising from differences in interpretation; and pragmatic boundaries, arising from differences in interests and consequences. He argues that boundary objects can support the transfer, translation, and transformation of knowledge across these boundaries.

This gives our design its theoretical backbone. In a time-boxed setting, prompts, templates, and a structured form can be engineered as boundary objects that (i) reduce syntactic ambiguity through fixed fields and shared terminology, (ii) support semantic alignment through common framing questions, and (iii) make pragmatic consequences discussable by requiring teams to justify resource choices and validation steps. Comparability is therefore supported by the instrument's structure, rather than by assuming participants already share a vocabulary. This proposition provides the design response to RQ1, which we evaluate empirically.

From boundary object to indicators: deriving the evaluation logic

These theoretical perspectives suggest that, in temporary and heterogeneous teams, a structured form can function as a boundary object by making selected features of emerging proposals explicit, traceable, and comparable. However, the form records only the elements that participants articulate through its categories; it cannot provide direct evidence of co-production, value-in-use, stakeholder alignment, feasibility, or societal impact. Its outputs are therefore treated as observable traces from which bounded proxy indicators can be derived. Figure 2 summarises this operational relationship by linking each theoretical concern to the relevant form fields, observable output, proxy indicator, and explicitly excluded inference.

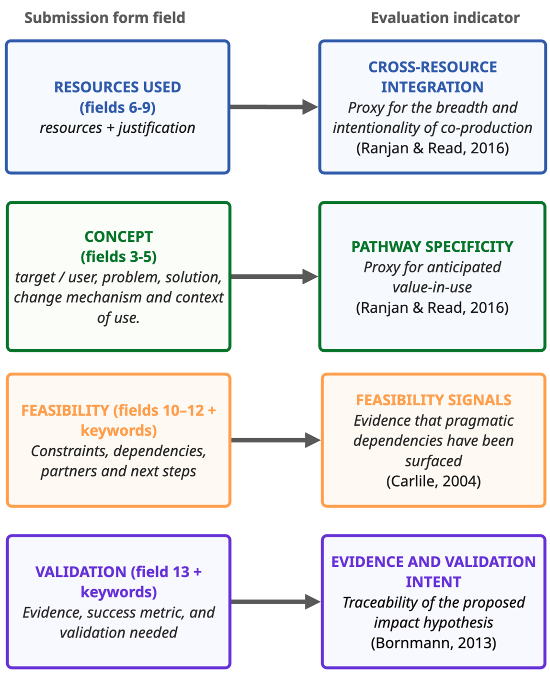

Fig. 2. Evaluation logic: how the submission-form field-groups map to the four proxy indicators. Each field-group (left) supplies the observable evidence for one indicator (right), grounded in the theoretical concept it proxies (value-in-use, co-production, pragmatic alignment, traceability).

Cross-resource integration draws on the concept of co-production (Ranjan & Read, 2016): fields 6–9 require teams to identify and justify the resources they propose to combine, while the indicator captures the breadth and diversity of that combination—not whether co-production has actually occurred. Pathway specificity draws on value-in-use (Ranjan & Read, 2016): fields 3–5 require teams to specify a beneficiary, a problem, and a mechanism of change. The indicator therefore assesses how clearly teams anticipate value in a particular context, not whether value-in-use has been realised. Feasibility signals reflect Carlile's (2004) pragmatic concern with differences in interests and consequences: fields 10–12 surface constraints, dependencies, partners, and next steps, but do not establish stakeholder alignment or demonstrate feasibility. Finally, evidence and validation intent addresses the challenge of maintaining traceability over long impact horizons (Bornmann, 2013): field 13 records what evidence would be needed and how it could be obtained. It therefore captures the traceability of an impact hypothesis rather than evidence of impact itself.

Completing the form generates the observable material on which these indicators are scored. Cross-resource integration is captured as a count of the resource categories combined (1–4). The remaining three indicators are each coded across four complementary dimensions drawn from the form's narrative fields, each scored on a three-level rubric (0 = absent; 1 = present but only implicitly or partially articulated; 2 = explicitly and clearly articulated) and summed to an indicator score from 0 to 8. This lightweight scheme is intended to support rapid yet reproducible comparison across proposals, rather than provide a fine-grained measure of co-creation quality or societal impact.

Together, these perspectives translate into four design requirements: the artifact should provide a common syntax, support translation across different interpretations, surface implementation consequences, and generate outputs that can be analysed across teams. Consistent with Design Science, we evaluate its utility through the completeness, traceability and comparability of these intermediate outputs, rather than as evidence of realised co-creation quality or societal impact.

The study thus evaluates what the artifact enables teams to articulate and researchers to compare, leaving the relational quality of co-creation and the realisation of impact to complementary methods and longer-term investigation.

METHOD AND DATA

This study employs a Design Science (DS) approach in which the artifact is both the primary contribution and the subject of evaluation (Romme, 2023). DS provides a suitable methodological lens for developing and evaluating collaboration artifacts that intervene in real settings while producing generalizable design knowledge (Romme, 2023). In participatory design research, collaboration system models have been proposed to plan and evaluate participatory projects by making assumptions, roles and value exchanges explicit (Drain & Sanders, 2019). In this study, DS frames both the creation of the artifact and the evaluation logic just described. The pilot serves as the artifact's first evaluation.

The research progressed through five phases: design requirements, artifact creation, iteration, implementation, and reflection.

Phase 1. Design Requirements

The artifact is designed to operate within a short conference workshop while producing outputs that are synthesizable at the ecosystem level. Three requirements guided it: (1) temporal efficiency - fit within 90 minutes while enabling convergence; (2) inclusivity and low barriers to participation - accessible to heterogeneous participants without prior design training; and (3) cross-resource integration - prompt teams to combine assets across four resource categories: R&D&I technology projects, research facilities, socio-economic studies (SES), and ATTRACT Academy methodologies.

Phase 2. Artifact Creation (protocol + capture instrument + synthesis routine)

The artifact comprised four coupled components (Figure 1), inspired by Need-to-Tech logic (Balboni et al., 2021) and Academy methodologies (Vignoli et al., 2021): structured team formation, a facilitation protocol, a standardised submission form, and a post-hoc synthesis routine. These four components were implemented through the six-step workshop protocol detailed in the Appendix: (1) pre-session preparation, (2) structured team formation, (3) facilitated ideation, (4) standardised submission form, (5) short pitches, and (6) post-hoc synthesis. The four components, therefore, describe the design of the artifact, whereas the six steps describe how it was implemented during and after the workshop.

During pre-session preparation, the organisers translated the overarching societal challenge: "How might Europe's innovation ecosystems - powered by deep tech, data, education, and collaboration - address societal challenges?" into five themed challenge cards: Science, Education & Inclusion; Sustainability & Environment; Innovation Ecosystems & Cohesion; Health & Well-being; and European Values & Future Vision. These cards provided a common starting point while allowing teams to address different dimensions of the broader challenge. Structured team formation then used colour-coded indicators to encourage interdisciplinary mixing and bring together participants representing different forms of expertise and ecosystem resources.

The facilitation protocol guided each team from broad challenge framing to a sufficiently specific proposal.

Facilitators introduced the societal challenge and used "How Might We" prompts to support problem framing. Time-boxed rounds of ideation and convergence required teams to first generate alternatives and then select and develop one proposal. Where used, an optional AI-assisted ideation round employed a lightweight large language model provided with ATTRACT-related content to generate additional possibilities for teams to consider. The AI-assisted round supplemented rather than replaced team deliberation and decision-making.

The standardised submission form, implemented in Microsoft Forms, converted the selected proposal into a structured and analysable output. Rather than asking teams to submit only a general concept, the form required them to articulate and justify its main elements. Fields 3–5 recorded the intended beneficiary, the problem being addressed, and the proposed mechanism of change. Fields 6–9 required teams to identify relevant resources from the four ecosystem-resource categories and explain why each resource was needed. Fields 10–12 captured implementation conditions, including constraints, dependencies, required partners, and proposed next steps. Field 13 asked teams to specify the evidence needed to examine their impact hypothesis and how that evidence could be obtained. In this way, the form made selected features of each proposal explicit and created a common structure for comparison across teams.

Following the completion of the form, voluntary teams presented their concept in a three-minute pitch during a 10–15-minute plenary session. The pitches made the proposals visible to the wider workshop and provided a concise opportunity to communicate their central logic, although the written submissions remained the data used in the subsequent analysis. Finally, the post-hoc synthesis routine aggregated the standardised submissions into portfolio-level patterns. The recorded fields provided the observable material used to assess four bounded proxy indicators—cross-resource integration, pathway specificity, feasibility signals, and evidence and validation intent—using the rule-based scoring system described below. Together, the four components enabled the workshop to move from interdisciplinary team formation and facilitated ideation to structured proposal capture and portfolio-level comparison.

Phase 3 and 4. Iteration and implementation

Development included a formative design check ("alpha test") with ATTRACT programme coordinators, including a critical iteration led by Pablo Garcia Tello (Project Coordinator). Those sessions revealed that an early draft focused too narrowly on specific facility–technology matches: the design was pushed to "think big", shifting the artifact toward broader mission-oriented ideation that could incorporate socio-economic evidence alongside technical projects. This "co-design of the co-design tool" helped ensure the artifact felt natural to researchers used to high-level scientific conceptualisation. It was deployed at the ATTRACT Final Conference (Brussels, 2–3 July 2024), where approximately 100 participants formed 25 teams across five themes, all teams completed the form, followed by a voluntary short-pitch session.

Phase 5. Reflection and Theorising

After the workshop, the dataset of 25 complete submissions was analysed for (a) resource mobilisations, (b) connective mechanisms (especially Academy methods and SES evidence), and (c) portfolio learning potential, combining frequency mapping, co-occurrence ("synergy") analysis, and thematic clustering. Portfolio learning here means using the aggregated, comparable outputs of many teams to draw programme-level lessons—which resources recur, which combinations are common, where gaps concentrate—that no single submission would reveal.

EVALUATION LOGIC

A core contribution is to make co-creation under time constraints sufficiently measurable for comparison. We evaluate the artifact at two layers, matching the research questions. For RQ1 (comparability by design), we assess three practical criteria: completion (teams complete submissions within the timebox); standardisation (submissions include the required fields and resource justifications); and synthesisability (submissions aggregate into portfolio patterns without extensive post-processing). RQ2 (outcome quality and usefulness) is addressed by the four-indicator set derived above and mapped to the form in Figure 2; Table 1 restates each indicator, what it captures, and its evidence source.

Table 1. Indicator set used for rapid assessment (cross-resource integration scored as a 1–4 count of resource categories combined; the other three indicators each coded on four dimensions, 0–2 each, summed to 0–8).

Indicator	What it captures	Analytical proxy for	Evidence Source
Cross-resource integration	Breadth and diversity of resource categories combined	Co-production (Ranjan & Read, 2016)	Form fields 6-9
Pathway specificity	Clarity of beneficiary, problem, mechanism of change, and context of use	Value-in-use (Ranjan & Read, 2016)	Form fields 3-5
Feasibility signals	Recognition of constraints, dependencies, partners, and next-step actions	Pragmatic alignment (Carlile, 2004)	Form fields 10-12 + keywords
Evidence/validation intent	Articulation of what evidence is needed and how it will be obtained (incl. SES)	Traceability under long impact horizons (Bornmann, 2013)	Form field 13

As established above, these indicators are bounded proxies for properties of the recorded proposals, rather than direct measures of co-production, realised value, feasibility, or societal impact.

RESULTS

Evidence for RQ1: feasibility, standardisation, and synthesisability

All 25 teams completed the submission instrument within the allotted timebox. The 100% completion rate met the feasibility criterion for RQ1 and demonstrates that the instrument could be used under conference conditions. The submissions also provided evidence of substantive resource integration. Sixteen of 25 teams (64%) mentioned at least three individual resources, regardless of the categories to which they belonged; 11 of 25 teams (44%) integrated resources from at least three of the four resource categories. Because the instrument required teams to explicitly identify and justify each resource, proposals followed a common structure and could be compared and aggregated without extensive post-processing.

These findings answer RQ1: in the studied setting, the instrument generated complete, standardised, and analysable intermediate outputs within a limited period. They establish the feasibility and synthesisability of the process, rather than the quality of the resulting co-creation or its eventual societal impact.

Evidence for RQ2: resource configurations and ecosystem roles

RQ2 was examined using four indicators derived from the submission instrument: cross-resource integration, pathway specificity, feasibility signals, and evidence and validation intent. As explained in the Methods, these indicators are pragmatic proxies for traceable intermediate outputs.

Table 2. Summary of available evidence for the four indicators

Indicator	Score distribution	Available descriptive evidence
Cross-resource integration	1 resource category: 8/25 submissions 2 resource category: 6/25 3 resource category: 6/25 4 resource category: 5/25	16 of 25 teams (64%) mentioned at least three individual resources, regardless of category; 11 of 25 (44%) integrated resources from at least three of the four resource categories.
Pathway specificity	Scoring: 0 - 8; 0 submissions with scoring ≤4; 3 submissions with scoring = 5–6; 22 submissions with scoring ≥ 7	Seven recurring proposal clusters were identified. 22 of 25 submissions explicitly identified a beneficiary, stated clearly the problem addressed, mechanism envisioned, and the context of use.
Feasibility signals	Scoring: 0 - 8; 4 submissions with scoring ≤4; 14 submissions with scoring = 5–6; 7 submissions with scoring ≥ 7	Miniaturisation, systems integration, policy implementation or financial pressure were identified as constraints across the submissions. 6 submissions articulated all four dimensions — constraints, dependencies, partners, and next steps — explicitly and clearly (scoring = 8); a seventh scored 7, with one dimension only implicitly articulated (7 submissions scoring ≥ 7 in total).
Evidence and validation intent	Scoring: 0 - 8; 8 submissions with scoring ≤4; 7 submissions with scoring = 5-6; 10 submissions with scoring ≥ 7	While 24 out of 25 submissions included a societal validation intent, only 3 of 25 submissions articulated all elements explicitly and clearly (scoring = 8) — the key performance indicator (KPI), the validation method, and the success metric; a further 7 scored 7, with one element only implicitly articulated (10 submissions scoring ≥ 7 in total).

Cross-resource integration

Academy methodologies were the most frequently cited category, appearing in 19 submissions (76%). R&D&I projects and SES projects were cited by 14 teams (56%), while research facilities were cited in 11 submissions (44%). The justifications supplied by teams also suggested different dominant functions for the four categories (Table 3).

Table 3. Resource-category frequency across the 25 submissions.

A. Resource category	Teams citing (n)	% of teams	Dominant function in submissions
Academy Methodologies	19	76%	Process / facilitation (the “how”); connective tissue across pillars
R&D&I Projects	14	56%	Technological core (the “what”); requires integration work beyond PoC
SES Projects	14	56%	Justificatory / evaluative (societal relevance, ethical alignment, policy legitimacy)
Research Facilities	11	44%	Legitimacy / convening anchor; CERN most cited (9/25 teams, top synergy partner)

Note: These functions describe dominant patterns in teams’ justifications and are not mutually exclusive.

Academy methodologies were commonly used to support problem framing, interdisciplinary collaboration, and the coordination of contributions. They therefore operated as a form of connective capacity within the proposed combinations. Research facilities were described not only as technical resources but also as recognised settings that could provide access, credibility, visibility, and a focal point for collaboration.

The co-occurrence data provide further evidence about the combinations proposed by teams. Academy methodologies and CERN formed the most frequent bundle, appearing together in nine submissions. CERN and XFEL co-occurred in seven submissions, Academy methodologies and H-CUBE in six, and Academy methodologies and ExSACT in five.

Table 4. Most frequent resource co-occurrence bundles

Resource combination	Teams citing both
Academy methodologies + CERN	9
CERN + XFEL	7
Academy methodologies + H-CUBE	6
Academy methodologies + ExSACT	5

These findings show the breadth of the resources that teams proposed combining and, where explicit justifications were supplied, the intended relationships among them.

Pathway specificity

Pathway specificity was operationalised as the degree to which each submission articulated a coherent pathway from a societal problem to value creation for an identifiable beneficiary. Building on the concept of value-in-use proposed by Ranjan and Read (2016), the indicator emphasises the explicit identification of the intended beneficiary and the context in which value would be realised, while extending these concepts with two additional dimensions—the problem being addressed and the proposed mechanism of change—to evaluate the completeness of the proposed impact pathway. Four complementary dimensions were therefore coded from the narrative fields of the submission form: (i) identification of the intended beneficiary, (ii) explicit articulation of the societal or technological problem, (iii) description of the mechanism through which the proposed solution would create value, and (iv) specification of the context of use. Each dimension was scored using a simple three-level rubric (0 = absent, 1 = present but only implicitly or partially articulated, 2 = explicitly and clearly articulated), resulting in a maximum Pathway Specificity Score of 8.

This lightweight coding scheme was designed to support rapid yet reproducible comparison across submissions while remaining consistent with the paper's Design Science approach. Application of the indicator to the 25 workshop submissions showed consistently high pathway specificity (mean = 7.68/8). Beneficiaries, mechanisms of change and contexts of use were explicitly identified in almost all submissions, whereas the formulation of the underlying problem was slightly less consistent, remaining implicit in a small number of proposals. These results suggest that the structured submission form successfully encouraged participants to articulate complete and comparable impact pathways, providing a robust basis for subsequent portfolio-level analysis.

Post-hoc thematic clustering nevertheless identified seven recurring types of proposed pathway:

(i) **Scientific Trust Initiative:** combating disinformation and rebuilding civic trust;
(ii) **InnoSee:** connecting Academy resources with industry;
(iii) **Public Health Intelligence Networks:** enabling real-time health monitoring;
(iv) **Translation Hubs of Excellence:** training professionals who can connect deep science, entrepreneurship, and societal needs;
(v) **"29th Regime":** redirecting innovation towards impact-oriented goals through a shared policy framework;
(vi) **Mobile Labs:** extending deep-technology capabilities to rural and remote communities; and
(vii) **Urban Sustainability Networks:** supporting real-time environmental monitoring and management.

These clusters emerged from similarities across submissions rather than from predefined mission categories. They reveal the substantive directions around which teams organised their proposals, but do not, by themselves, show how each team specifically identified a beneficiary, context, or mechanism of change. They should therefore be treated as recurring proposal clusters and candidate mission pathways, not as evidence of realised value-in-use.

Feasibility signals

Feasibility signals were operationalised as the degree to which submissions recognised the practical conditions required to translate an idea into implementation. This indicator draws on Carlile's (2004) concept of pragmatic knowledge boundaries, which emphasises that successful innovation depends not only on sharing information but also on recognising dependencies, consequences and the coordination required among heterogeneous actors.

To capture these implementation-oriented aspects, four complementary dimensions were coded from the narrative fields of the submission form: (i) explicit recognition of implementation constraints (e.g. technical barriers, regulatory issues, miniaturisation or systems integration), (ii) identification of critical dependencies such as enabling technologies, research infrastructures or datasets, (iii) specification of implementation partners (e.g. companies, universities, public authorities or other organisations), and (iv) articulation of a concrete next step, such as piloting, validation or deployment.

These dimensions extend Carlile's (2002, 2004), theoretical framework into a practical evaluation instrument suitable for rapid assessment of early-stage co-creation outputs. Each dimension was scored using the same three-level rubric (0 = absent, 1 = present but only implicitly or partially articulated, 2 = explicitly and clearly articulated), producing a Feasibility Signal Score ranging from 0 to 8. Application of the indicator to the 25 submissions yielded a mean score of 6/8, substantially lower than pathway specificity. Dependencies and proposed next steps were identified in most submissions, whereas explicit discussion of implementation constraints and operational partnerships was considerably less frequent.

This suggests that participants readily described promising concepts but devoted less attention to the practical conditions required for implementation, reinforcing the interpretation that many proposals remained optimistic conceptually while under-specifying the translation challenges associated with moving beyond proof-of-concept. These findings show whether implementation dependencies were surfaced in the proposals. They do not establish that the proposed arrangements were feasible or that the relevant stakeholders had agreed to them.

Evidence and validation intent

Evidence and Validation Intent was operationalised as the extent to which submissions explicitly articulated how future claims of societal value could subsequently be substantiated. This indicator is grounded in Bornmann's (2013) discussion of societal impact assessment, which argues that the long-time horizons and attribution challenges associated with research impact require the use of intermediate, traceable outputs rather than direct measurement of ultimate societal impact. Accordingly, the indicator evaluates whether teams explicitly considered the evidence that would eventually be required to support their proposed impact pathway. Four complementary dimensions were coded from the submission forms: (i) identification of the evidence required, (ii) description of an appropriate validation method (e.g. pilot, experiment or trial), (iii) definition of an intended success metric, and (iv) consideration of societal validation through users, policy, ethics or socio-economic studies (SES).

These dimensions extend Bornmann's conceptual discussion into a practical evaluation instrument suitable for assessing early-stage co-creation outputs. Each dimension was scored using the same three-level rubric (0 = absent, 1 = present but only implicitly or partially articulated, 2 = explicitly and clearly articulated), yielding a maximum Evidence and Validation Intent Score of 8. Application of the indicator to the 25 submissions produced the lowest average score among the four indicators (mean = 5.28/8). Explicit plans for evidence collection and validation methods were frequently absent or only partially specified, whereas intended societal benefits and success criteria were articulated more consistently. These findings suggest that participants generally had little difficulty explaining why their ideas would create societal value but devoted substantially less attention to describing how those claims would later be demonstrated. This imbalance highlights a recurring weakness in early-stage co-creation activities and supports Bornmann's argument that intermediate evidence and validation planning are essential when societal impact itself cannot yet be directly assessed.

These scores characterise how completely validation intentions were articulated; they do not assess the validity, feasibility, or eventual success of those validation plans.

Cross-indicator synthesis

Read together, the indicators reveal a consistent progression from strongly articulated value pathways toward weaker implementation and validation planning. Pathway specificity was high (mean 7.68/8; 22 of 25 submissions scoring ≥ 7 and none ≤ 4): teams reliably articulated a beneficiary, a problem, a mechanism of change, and a context of use. Feasibility signals were weaker (mean 6.00/8; 7 submissions scoring ≥ 7, 14 in the 5–6 range, and 4 at ≤ 4), and evidence and validation intent was weakest (mean 5.28/8; 10 submissions ≥ 7 but 8 at ≤ 4). The three indicators thus decline in the order pathway → feasibility → validation. This gradient is the study's central cross-indicator finding, and it aligns with the breadth result. Teams were markedly more successful at identifying complementary resources and articulating a plausible value pathway than at specifying the operational conditions (constraints, partners, next steps) and the evidentiary arrangements (validation methods, success metrics) needed to translate those combinations into use. The same imbalance appears within a single indicator: 24 of 25 submissions expressed a societal validation intent, yet only three articulated a sufficiently complete validation plan including a KPI, validation method, and success criterion. Under time-boxed conditions, the artifact elicited what teams wanted to achieve and which resources they would combine, but was progressively less able to elicit how the resulting proposals would be implemented and evidenced. This concentration of weakness at the feasibility-and-validation end is the "last-mile" gap discussed above, now visible as a portfolio-level pattern rather than an anecdote.

DISCUSSION

The findings support the paper’s central proposition: heterogeneous participants do not need to become homogeneous for their outputs to be comparable. Comparability can instead be supported by a designed boundary object embedded in a structured process. The contribution therefore lies in the coupling of interdisciplinary team formation, facilitated ideation, standardised proposal capture, and portfolio-level synthesis.

RQ1: Designing comparability under conference constraints

All 25 teams completed the protocol and produced submissions that could be analysed through a common structure. Moreover, 44% combined at least three resource categories, under conditions of limited time and heterogeneous expertise. Although most teams combined fewer than three categories, the results show that substantial resource combinations could be articulated within a temporary, time-boxed setting.

This suggests a design implication extending beyond ATTRACT: comparability in heterogeneous co-creation does not require homogeneous participants, but well-designed boundary objects and facilitation. Fixed fields and shared terminology provided a common syntax; questions about beneficiaries, problems, and mechanisms supported semantic alignment; and prompts concerning resources, dependencies, partners, and validation made selected pragmatic consequences discussable. Consistent with Carlile (2002, 2004), the instrument supported the articulation of knowledge across syntactic, semantic, and pragmatic boundaries. The study does not demonstrate that participants achieved shared interpretations or aligned interests, but shows that relevant elements of their proposals could be recorded in a common format.

The artifact should be understood as a socio-technical configuration rather than as a stand-alone form. Its contribution lies in coupling heterogeneous team formation, facilitated convergence, structured articulation and post-hoc synthesis so that temporary collaboration leaves behind outputs that remain analysable after the event.

Time pressure was not removed but used as a coordinating mechanism. The structured process helped prevent rapid convergence from producing only compressed or ambiguous ideas. It shifted the workshop output from an unstructured pitch towards an explicit, although preliminary, pathway hypothesis that remained traceable after the event.

RQ2: What the indicators reveal

The four indicators distinguished proposals that remained at the level of "interesting ideas" from those articulating more traceable connections among beneficiaries, resources, implementation conditions, and validation intentions. Their purpose was not simply to assign scores, but to distinguish between naming an element and explicitly justifying it.

Taken together, the indicators differentiated proposals primarily by how far teams progressed from articulating a desirable pathway toward specifying implementation and validation. Their value is therefore diagnostic rather than evaluative in the sense of judging idea quality or realised impact.

At the portfolio level, the indicators revealed patterns that would have been difficult to identify from unstructured pitches. Methodological capacity, including Academy methods, emerged as a connective infrastructure linking technologies, facilities, and societal needs. Methods supported problem framing, translation across domains, and the articulation of pathways through which other resources could be combined. Methodological and facilitation capacity may therefore constitute a first-order ecosystem asset rather than an auxiliary service.

Facilities appeared as both technical and socio-institutional resources: they offered specialised equipment and expertise, while also providing networks, credibility, and institutional anchoring. This aligns with evidence that large-scale research infrastructures shape surrounding knowledge ecosystems (Rådberg & Löfsten, 2023). However, naming a facility does not show that access has been secured or that its involvement is necessary. Requiring teams to justify its contribution helps distinguish an actionable role from a symbolic association.

Recurring last-mile gaps also suggest that translation beyond proof-of-concept is a socio-technical integration problem, not solely a commercialisation challenge. Societal application requires coordination among technologies, methods, facilities, beneficiary knowledge, implementation partners, and evidence-generation activities. Consistent with mission-oriented innovation research (Mazzucato, 2018), the findings shift attention from the maturity of individual technologies towards the complementary capabilities and relationships required for their use.

Standardisation is a means, not the answer

The form makes selected characteristics of proposals visible and comparable, but does not reveal whose voice shaped them, which knowledge was excluded, or how influence was distributed within teams (von Busch & Palmås, 2023; Smith et al., 2025). The resulting scores should therefore be treated as properties of recorded proposals rather than measures of participation quality, legitimacy or societal impact. They support portfolio learning but cannot substitute for qualitative interpretation, stakeholder validation or longer-term impact assessment (Bührer et al., 2022).

Standardisation supports comparability; it does not establish legitimacy.

Implications for programme design

Three implications follow. First, methodological capacity should be treated as portfolio infrastructure. Investing in shared methods, facilitation, and boundary-spanning capabilities may matter as much as investing in technologies, because these resources help participants identify common problems and integrate otherwise disconnected assets.

Second, facilities can function as equipment providers, testbeds, conveners, and legitimacy anchors. Programme designers should nevertheless require proposals to specify their expected contribution rather than treating institutional association as evidence of feasibility.

Third, recurring last-mile gaps point to the need for translation capacity beyond proof-of-concept, including beneficiary engagement, partnership formation, evidence generation, and coordination of implementation dependencies. The indicators can support this work diagnostically by showing where proposals require further resources or facilitation before additional investment.

LIMITATIONS AND FUTURE RESEARCH

This study reports a single conference deployment with a self-selected participant pool; outputs remain at the concept level and are not validated for technical feasibility. The lightweight four-dimension scoring that enables rapid comparison is itself a limitation—it captures a pathway's presence and justification, not its depth—so future cycles should test inter-rater reliability (targeting Cohen's $\kappa \geq 0.70$) and complement scores with qualitative reading. Not-invented-here and professional-identity dynamics were treated as design motivations rather than measured outcomes; future cycles could add short instruments to measure boundary perceptions or openness to external knowledge (Hannen et al., 2019; Lifshitz-Assaf, 2018; Dosi et al., 2025).

The pilot suggests several candidate boundary conditions for future testing: (a) participant mix spans several of the four resource categories, which may support broader and more synthesisable outputs (a pattern suggested by, but not formally tested in, this pilot); (b) facilitation uses pre-curated challenge cards and "How Might We" prompts rather than open-ended briefings; (c) the timebox is at least 90 minutes; and (d) digital capture is in place before the session, since post-hoc transcription degrades synthesisability. Future replications should test performance with more homogeneous participant pools, without mid-session prompts, and when teams lack prior exposure to the resource taxonomy. A short pre-session briefing on the four categories is therefore one promising preparatory intervention to test.

Finally, the optional AI-assisted ideation step raises an "AI for impact" paradox: work on generative AI in engineering education stresses the sustainability of AI systems and the risk of scaling work without attention to socio-environmental costs (Cañavate et al., 2025). Future iterations should evaluate when and how AI assistance improves outcomes (novelty, specificity, validation intent) and what governance, provenance, and sustainability safeguards are required.

CONCLUSION

Co-creation under conference constraints need not produce outputs that disappear after the event. In the ATTRACT pilot, a structured protocol and standardised capture instrument enabled 25 teams to produce comparable intermediate outputs that could subsequently be analysed at portfolio level. The evaluation logic revealed a consistent progression from strongly articulated pathways toward weaker implementation and validation planning. The contribution is therefore not a measure of co-creation quality or societal impact, but a reusable way to convert time-boxed collaboration into traceable material for programme learning. Comparability is a means to learning, not a verdict on worth.

ACKNOWLEDGEMENTS

This project has received funding from the European Union's Horizon 2020 research and innovation programme (ATTRACT Phase 2) under grant agreement No 101004462. The authors would like to extend their special thanks to all ATTRACT participants and to the Final Conference attendees for their valuable contributions, engagement, and support.

AI USE DISCLOSURE

Parts of this manuscript were developed with the assistance of generative AI tool (ChatGPT and Claude AI) to support language editing, restructuring, and clarity. The authors reviewed, edited, and verified all content and references, and remain fully responsible for the final manuscript.

CONFLICT OF INTEREST

A.B., R.G. and M.S. declare no conflict of interest. M.V. is Editor-in-Chief of the CERN IdeaSquare Journal of Experimental Innovation and a co-author of this article. M.V. was excluded from all editorial decision-making and peer-review management related to this manuscript.

REFERENCES

Angarita, M. A. M., & Nolte, A. (2020). *What Do We Know About Hackathon Outcomes and How to Support Them? A systematic literature review*. In Collaboration Technologies and Social Computing (CollabTech 2020) (pp. 50–64). Springer. https://doi.org/10.1007/978-3-030-58157-2_4

Balboni, B., Dosi, C., Marchini, S., Mincolelli, G., & Vignoli, M. (2021). N2T 'Need to Tech discovery' tool: enabling interaction with scientists in CBI students' projects. *CERN IdeaSquare Journal of Experimental Innovation*, 5(1), 22–31. https://doi.org/10.23726/cij.2021.1289

Bornmann, L. (2013). What is societal impact of research and how can it be assessed? A literature survey. *Journal of the American Society for Information Science and Technology*, 64(2), 217–233. https://doi.org/10.1002/asi.22803

Bührer, S., Feidenheimer, A., Walz, R., Lindner, R., Beckert, B., & Wallwaey, E. (2022). *Concepts and methods to measure societal impacts – an overview* (Discussion Papers Innovation Systems and Policy Analysis No. 74). Fraunhofer ISI.
Carlile, P. R. (2002). A pragmatic view of knowledge and boundaries: Boundary objects in new product development. *Organization Science*, 13(4), 442–455. https://doi.org/10.1287/orsc.13.4.442.2953
Carlile, P. R. (2004). Transferring, translating, and transforming: An integrative framework for managing knowledge across boundaries. *Organization Science*, 15(5), 555–568. https://doi.org/10.1287/orsc.1040.0094
Cañavate, J., Martínez-Marroquín, E., & Colom, X. (2025). Engineering a Sustainable Future Through the Integration of Generative AI in Engineering Education. *Sustainability,* 17(7), 3201. https://doi.org/10.3390/su17073201
Dosi, C., Gorni, E., Vignoli, M., Cau, S., & Presti, G. (2025). Supporting Open Innovation Behaviors in Researchers: Developing a Training with Acceptance and Commitment Training. *CERN IdeaSquare Journal of Experimental Innovation*, 9(1), 49–60. https://doi.org/10.23726/cij.2025.1680
Drain, A., & Sanders, E. B.-N. (2019). A collaboration system model for planning and evaluating participatory design projects. *International Journal of Design*, 13(3), 39–52.
Gersick, C. J. G. (1988). Time and transition in work teams: Toward a new model of group development. *Academy of Management Journal*, 31(1), 9–41. https://doi.org/10.2307/256496
Granstrand, O., & Holgersson, M. (2020). Innovation ecosystems: A conceptual review and a new definition. *Technovation*, 90–91, 102098. https://doi.org/10.1016/j.technovation.2019.102098
Hannen, J., Antons, D., Piller, F., Salge, T. O., Coltman, T., & Devinney, T. M. (2019). Containing the Not-Invented-Here Syndrome in external knowledge absorption and open innovation: The role of indirect countermeasures. *Research Policy*, 48(9), 103822. https://doi.org/10.1016/j.respol.2019.103822
Lifshitz-Assaf, H. (2018). Dismantling knowledge boundaries at NASA: The critical role of professional identity in open innovation. *Administrative Science Quarterly*, 63(4), 746–782. https://doi.org/10.1177/0001839217747876
Mazzucato, M. (2018). Mission-oriented innovation policies: Challenges and opportunities. *Industrial and Corporate Change*, 27(5), 803–815. https://doi.org/10.1093/icc/dty034
Meyerson, D., Weick, K. E., & Kramer, R. M. (1996). Swift trust and temporary groups. In R. M. Kramer & T. R. Tyler (Eds.), *Trust in Organizations: Frontiers of Theory and Research* (pp. 166–195). Sage. https://doi.org/10.4135/9781452243610.n9
Ranjan, K. R., & Read, S. (2016). Value co-creation: Concept and measurement. *Journal of the Academy of Marketing Science*, 44(3), 290–315. https://doi.org/10.1007/s11747-014-0397-2
Romme, G. (2023). Design science as experimental methodology in innovation and entrepreneurship research: A primer. *CERN IdeaSquare Journal of Experimental Innovation*, 7(2), 4–7. https://doi.org/10.23726/cij.2022.1427
Rådberg, K. K., & Löfsten, H. (2023). Developing a knowledge ecosystem for large-scale research infrastructure. *The Journal of Technology Transfer*, 48, 441–467. https://doi.org/10.1007/s10961-022-09945-x
Sanders, E. B.-N., & Stappers, P. J. (2008). Co-creation and the new landscapes of design. *CoDesign*, 4(1), 5–18. https://doi.org/10.1080/15710880701875068
Simon, H. A. (1969). *The Sciences of the Artificial*. MIT Press.
Smith, R. C., Huybrechts, L., Simonsen, J., & Loi, D. (2025). *Contemporary participatory design: Research agendas for societal crisis*. In Proceedings of the Sixth Decennial Aarhus Conference: Computing X Crisis (AAR '25) (pp. 182–201). Association for Computing Machinery. https://doi.org/10.1145/3744169.3744183
Star, S. L., & Griesemer, J. R. (1989). Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science*, 19(3), 387–420. https://doi.org/10.1177/030631289019003001
Vignoli, M., Balboni, B., Cotoranu, A., Dosi, C., Glisoni, N., Kohler, K., Mincolelli, G., Mäkinen, S., Nordberg, M., & Thong, C. (2021). Inspiring the future change-makers: Reflections and ways forward from the Challenge-Based Innovation experiment. *CERN IdeaSquare Journal of Experimental Innovation*, 5(1), 1–4. https://doi.org/10.23726/cij.2021.1323
von Busch, O., & Palmås, K. (2023). *The Corruption of Co-Design: Political and Social Conflicts in Participatory Design Thinking*. Routledge.
Voorberg, W. H., Bekkers, V. J. J. M., & Tummers, L. G. (2015). A systematic review of co-creation and co-production: Embarking on the social innovation journey. *Public Management Review*, 17(9), 1333–1357. https://doi.org/10.1080/14719037.2014.930505
Wareham, J., Pujol Priego, L., Romasanta, A. K., Mathiassen, T. W., Nordberg, M., & Tello, P. G. (2022). Systematizing serendipity for big science infrastructures: The ATTRACT project. *Technovation*, 116, 102374. https://doi.org/10.1016/j.technovation.2021.102374

APPENDIX

Protocol steps

1. Step 1: Pre-session preparation (1–2 weeks prior). Identify 3–5 societal challenge themes relevant to your programme portfolio. Prepare one "Challenge Card" per theme (one A5 card: challenge title, 2–3 context sentences, one "How Might We" prompt). Compile a one-page Resource Taxonomy handout listing the four resource categories with 3–5 named examples from your own ecosystem. Set up the digital submission form (see Section B below). Prepare colour-coded team formation cards (one colour per resource category; each participant receives the colour corresponding to their primary domain).
2. Step 2: Structured team formation (10 min). Assign participants to mixed-resource teams of three using colour-coded cards. Teams of three can interact easily in a conference setting just sitting close by each other. Aim for at least three of the four resource categories represented per team. If the pool is predominantly single-disciplinary, assign challenge themes that foreground the missing categories.
3. Step 3: Facilitated ideation (50–60 min). Distribute Challenge Cards and Resource Taxonomy handout. Phase 1 (20 min): divergent ideation around the "How Might We" prompt; facilitator circulates to prompt cross-resource combinations. Phase 2 (20 min): convergence on one concept; teams complete a rough Resource Justification map (which resources are needed and why). Optional: AI-assisted ideation round (10 min) using a lightweight LLM pre-loaded with ecosystem content to surface non-obvious resource combinations.
4. Step 4: Standardised submission form (10–15 min). Each team completes the digital submission form (Section B) — the artifact's core data-capture instrument and boundary object. Its 13 fixed fields require teams to record their concept (target beneficiary, problem, and mechanism of change), the resources mobilised across the four categories with an explicit justification for each, key dependencies, partners, and next steps, and the evidence or validation needed to proceed. Because every team answers the same fixed fields, heterogeneous outputs are rendered comparable by design, and the completed form directly supplies the evidence scored by the four indicators (Table 1). Teams submit before the pitch round begins.
5. Step 5: Short pitches (10 min). Voluntary teams present their concept in 3 minutes. A designated note-taker records any resource justifications or feasibility signals not captured in the form.
6. Step 6: Portfolio-level synthesis (within 2 weeks). Export form responses. Apply the indicator set (Table 1) to score each submission. Conduct frequency mapping and co-occurrence analysis across resource fields. Cluster submissions thematically to identify portfolio-level patterns. Report findings back to participants and programme stakeholders.

Submission form fields (template)

The form should capture the following fields (free text unless indicated): (1) Team identifier and challenge theme [dropdown]; (2) Concept title and one-sentence description; (3) Target beneficiary / user group; (4) Problem being addressed; (5) Mechanism of change (how does the concept create value?); (6) R&D&I resources mobilised [multi-select + free text justification]; (7) Research Facility resources mobilised [multi-select + free text justification]; (8) SES resources mobilised [multi-select + free text justification]; (9) Academy Methodologies mobilised [multi-select + free text justification]; (10) Key dependencies or constraints; (11) Identified partners needed; (12) Proposed next step; (13) Evidence or validation needed to proceed [free text]. Note: fields 6–9 feed the cross-resource integration indicator; fields 3-5 pathway specificity, fields 10-12 + keywords feasibility signals, and field 13 + keywords evidence and validation intent.

Synthesis procedure

Export all form responses to a spreadsheet. For each submission:

(i) score fields 6–9 on the Cross-Resource Integration indicator based on the number of resource categories mentioned (1-4)

(ii) using the information in fields 3 score pathway specificity based on (a) identification of the intended beneficiary, (b) explicit articulation of the societal or technological problem, (c) description of the mechanism through which the proposed solution would create value, and (d) specification of the context of use. Each dimension was scored using a simple three-level rubric (0 = absent, 1 = present but only implicitly or partially articulated, 2 = explicitly and clearly articulated), resulting in a maximum Pathway Specificity Score of 8;

(iii) by combining fields 10-12 + keywords score feasibility signals based on (a) explicit recognition of implementation constraints, (b) identification of critical dependencies, (c) specification of implementation partners, and (d) articulation of concrete next steps. Score each dimension using the same three-level rubric (0 = absent, 1 = present but only implicitly or partially articulated, 2 = explicitly and clearly articulated), producing a Feasibility Signal Score ranging from 0 to 8.

(iv) score Evidence and Validation Intent (0–8) from fields 13 + keywords. Sum the scores of the 4 dimensions:

(a) identification of the evidence required, (b) description of an appropriate validation method, (c) definition of an intended success metric, and (d) consideration of societal validation, using three-level rubric (0 = absent, 1 = present but only implicitly or partially articulated, 2 = explicitly and clearly articulated),

(v) Aggregate scores by challenge theme and at portfolio level. For frequency mapping: count the number of submissions citing each named resource (fields 6–9); calculate co-occurrence by counting submissions citing both resources in any pair. For thematic clustering: group submissions by shared resource bundles or pathway logic; label emergent clusters bottom-up (do not force-fit to pre-defined themes). Target inter-rater reliability (Cohen's $\kappa \geq 0.70$) for indicator scoring if used for comparative research.

Supplementary Materials

1. Attract final conference workshop presentation
2. Attract workshop challenge cards
3. Attract final conference RIS summary
4. Attract final conference SES summary
5. Attract final conference support playbook