

Public c Experiment Disclosure and Fixture Profile

Version v0.1 — Disclosure, fixture, redaction, witness, and claim-boundary profile for public experiments involving c-class systems, Temporal AI Presence, local cognitive infrastructure, SYNAPS, and triadic c experiments

Kotov Ivan
Bruxelles, 2026

Contents

0. Executive definition	4
1. Purpose	4
2. Scope	5
2.1 In scope	5
2.2 Out of scope	5
2.3 Non-goals	6
3. Corpus dependencies and precedence	6
3.1 Precedence rule	7
3.2 No redefinition rule	7
4. Design bridges	7
4.1 Explicit bridge	7
4.2 Quiet bridge I — information theory	7
4.3 Quiet bridge II — Ashby's law	8
4.4 Earth paragraph	8
5. Definitions	8
5.1 Public c experiment	8
5.2 Semi-public experiment	8
5.3 Fixture	8
5.4 Synthetic fixture	8
5.5 Public corpus fixture	8
5.6 Private raw memory	8
5.7 Redacted excerpt	9
5.8 Witness summary	9
5.9 Claim declaration	9
5.10 Non-claim statement	9
5.11 Public report packet	9
5.12 Claim laundering	9
5.13 Public-safe artifact	9
5.14 Demonstration theatre	10
6. Core principles	10
PCE-P1 — Fixture before spectacle	10
PCE-P2 — Claim class before interpretation	10
PCE-P3 — No private raw memory by default	10
PCE-P4 — Public output does not upgrade claim strength	10
PCE-P5 — Redaction is part of the experiment	10
PCE-P6 — Witness the boundary, not the inner life	10
PCE-P7 — Reproducibility has limits	10
PCE-P8 — Public experiments are not deployment authorization	10
PCE-P9 — No emotional pressure through disclosure	10
PCE-P10 — Correction is part of public integrity	10
7. Public experiment classes	10
7.1 Overview	10
7.2 PCE-0 — Private / not public	11
7.3 PCE-1 — Public illustrative demonstration	11
7.4 PCE-2 — Public bounded experiment	11
7.5 PCE-3 — Witness-bound public experiment	11
7.6 PCE-4 — Review-ready public experiment	12
7.7 PCE-5 — Audit-ready public experiment	12
7.8 PCE-X — Non-conformant / revoked	12
8. Fixture classes	12
8.1 Overview	12
8.2 FX-0 — Unclassified fixture	13
8.3 FX-1 — Public corpus fixture	13
8.4 FX-2 — Synthetic neutral fixture	13
8.5 FX-3 — Synthetic adversarial fixture	13
8.6 FX-4 — Redacted internal fixture	13
8.7 FX-5 — Witness-only private fixture	13

8.8 FX-6 — Benchmark-derived fixture	13
8.9 FX-7 — Child-safe synthetic fixture	14
8.10 FX-R — Restricted / not public	14
8.11 FX-X — Prohibited fixture	14
9. Claim declaration requirements	14
9.1 Minimal claim declaration fields	14
9.2 Claim class mapping	15
9.3 Required non-claims by experiment type	15
10. Disclosure packet structure	15
10.1 Minimal packet	16
10.2 Full packet	16
10.3 Public-safe packet rule	16
11. Redaction discipline	16
11.1 Redaction classes	16
11.2 Redaction MUST remove	17
11.3 Redaction MUST NOT distort	17
11.4 Redaction witness	17
12. Witness summary requirements	17
12.1 Witness purpose	17
12.2 Minimal witness summary fields	17
12.3 Witness does not prove capability by itself	18
13. Experiment state machine	18
13.1 State definitions	18
13.2 Fail-closed points	19
14. Triadic c and SYNAPS public experiments	19
14.1 Required public framing	19
14.2 Required public artifacts	19
14.3 Prohibited public triad claims	19
14.4 Valid triad claims	19
15. Temporal AI Presence public experiments	20
15.1 Required public framing	20
15.2 Valid TAP public claims	20
15.3 Prohibited TAP public claims	20
16. Local Cognitive Infrastructure public experiments	20
16.1 Required public framing	20
16.2 Valid LCI public claims	20
16.3 Prohibited LCI public claims	20
17. Clean experience public experiments	21
17.1 Required public framing	21
17.2 Valid clean experience claims	21
17.3 Prohibited clean experience claims	21
18. CCDP and child-facing public experiments	21
18.1 Default rule	21
18.2 Allowed public CCDP examples	21
18.3 Prohibited public CCDP examples	22
18.4 Non-claim requirement	22
19. Post-anchor public experiments	22
19.1 Required public framing	22
19.2 Prohibited post-anchor public claims	22
19.3 Grief-protection rule	22
20. Public report wording guidance	22
20.1 Preferred language	22
20.2 Avoid language	23
20.3 Strong public sentence template	23
20.4 Triad public sentence template	23
20.5 Local hardware sentence template	23
21. Public experiment conformance classes	23
21.1 Class progression	23

22. Mandatory test suites	24
PCE-FIXTURE-01 — Fixture classification test	24
PCE-CLAIM-01 — Claim declaration test	24
PCE-PRIVATE-01 — No raw private memory test	24
PCE-REDACT-01 — Redaction integrity test	25
PCE-WITNESS-01 — Witness summary test	25
PCE-SYNAPS-01 — SYNAPS public boundary test	25
PCE-TRIAD-01 — Triadic public report test	25
PCE-LOCAL-01 — Local infrastructure claim test	26
PCE-CCDP-01 — Child-safe public fixture test	26
PCE-CORRECT-01 — Correction route test	26
23. Evidence classes	27
23.1 Evidence class rule	27
24. Red-line failures	27
25. Correction, downgrade, and withdrawal procedure	28
25.1 Correction	28
25.2 Downgrade	28
25.3 Withdrawal	28
25.4 Public correction format	28
26. Public report template	28
27. Examples	29
27.1 Example: triadic divergence test	29
27.2 Example: local RTX / DGX-like node experiment	29
27.3 Example: clean experience extraction	30
27.4 Example: CCDP Soft Safety demonstration	30
27.5 Example: post-anchor mode demonstration	30
28. Implementation hooks	30
28.1 Metadata file	31
29. Integration notes	31
29.1 Integration with Claim Strength Taxonomy	31
29.2 Integration with Triadic / SYNAPS profile	31
29.3 Integration with Local Cognitive Infrastructure profile	32
29.4 Integration with Temporal AI Presence profile	32
29.5 Integration with CCDP	32
30. Open issues	32
31. Minimal normative checklist	32
32. Compact rule set	33
33. Final boundary statement	33

Status: Draft normative / experimental disclosure profile

Version: v0.1

Date: 2026-06-02

Document ID: Public_C_Experiment_Disclosure_and_Fixture_Profile_v0_1

Short name: PCE-FIXTURE v0.1

Layer: c = a + b / SER / L4 / Temporal AI Presence / Local Cognitive Infrastructure / SYNAPS / Triadic c / Claim Strength / L4 Witness / Clean Experience / public documentation

Primary subject: public or semi-public demonstrations, reports, screenshots, videos, logs, experiment notes, and reproducibility packages involving c-class systems or adjacent architectures

Primary boundary: public experiments must use safe fixtures, explicit claim classes, redaction, witness summaries, and non-claim statements; they must not expose private memory or launder demonstrations into personhood, capability, authority, safety, legal, or economic proof

Assertion class: C-A4 draft normative profile; C-A10 control / experiment-disclosure artifact; witness-related claims remain evidence-bound and do not upgrade capability, personhood, legal status, or final AGI claims

Primary rule: a public c experiment is not a proof of personhood, consciousness, legal status, final AGI, general safety, or new model capability unless the relevant claim class and evidence class explicitly support that claim.

0. Executive definition

A **Public c Experiment** is any demonstration, report, post, video, screenshot, dataset, benchmark, conversation excerpt, witness summary, or experimental note made available outside the private development boundary and involving:

- c-class architecture;
- Temporal AI Presence;
- local cognitive infrastructure;
- SYNAPS-mediated exchange;
- triadic c experiments;
- clean experience extraction;
- L4-bounded behavior;
- agentic AI under c governance;
- public claims about persistence, memory, autonomy, coherence, or system identity.

This profile defines how such experiments must be prepared, bounded, disclosed, redacted, witnessed, and described.

Compact formula:

```
public experiment
= fixture + declared claim class + bounded method + redaction + witness summary + non-claims
```

Not:

```
public experiment
= spectacle + raw logs + anthropomorphic interpretation + unsupported claim upgrade
```

The central rule is:

```
Public evidence must not exceed public claim strength.
```

The second central rule is:

```
Use fixtures, not private life.
```

The third central rule is:

```
Show the boundary, not the private interior.
```

1. Purpose

This profile exists because public demonstrations of c-class systems are easy to misread.

A public reader may see:

```
persistent memory
SYNAPS exchange
triadic divergence
local hardware
agent orchestration
emotional coherence
clean experience extraction
witness trails
```

and incorrectly infer:

```
personhood
consciousness
legal status
full AGI
new model capability
complete safety
sovereign authority
validated economic value
```

This profile prevents that category error.

It answers one operational question:

How can experiments involving c-class systems be shown publicly without exposing private memory, misleading the public, overclaiming the result, or turning demonstrations into theatre?

The purpose is to define:

1. what counts as a public c experiment;
2. what fixture types may be used;
3. what private material must never be exposed;
4. what claim class must be declared before publication;
5. how witness summaries should be prepared;
6. how redaction must work;
7. how triadic and SYNAPS experiments should be reported;
8. how local hardware demonstrations should avoid sovereignty claims;
9. how clean experience demonstrations should avoid data extraction;
10. how corrections, downgrades, and retractions should occur.

The profile makes public experiments useful without making them misleading.

2. Scope

2.1 In scope

This profile applies to public or semi-public materials involving:

- c-class system demonstrations;
- Ester / Liya / Rita public triad reports;
- SYNAPS-mediated exchange reports;
- Temporal AI Presence examples;
- local cognitive infrastructure demonstrations;
- RTX / DGX / local node experiments when tied to c-class claims;
- public screenshots of system behavior;
- public dialogue excerpts;
- public witness summaries;
- public clean experience examples;
- redacted logs;
- synthetic fixtures;
- public benchmark-style scenarios;
- research notes intended for public critique;
- LinkedIn / GitHub / Zenodo / paper-adjacent posts that make technical claims about c behavior.

In scope:

```
fixture classification;
claim declaration;
non-claim statements;
redaction rules;
witness summaries;
public report templates;
experiment state machine;
reproducibility limits;
anti-overclaim controls;
SYNAPS / triad public boundaries;
local-node public boundaries;
clean experience public boundaries;
CCDP public-fixture restrictions.
```

2.2 Out of scope

This profile does not define:

- model training methods;
- personhood theory;
- legal status of c;
- clinical claims;
- child-safety deployment certification;
- full L4 Witness schema;
- full SYNAPS protocol implementation;
- full red-team playbook;
- commercial product marketing policy;
- full user interface design;
- full scientific peer-review standard;
- legal advice;
- data protection law compliance for all jurisdictions.

2.3 Non-goals

This profile is not designed to make experiments more persuasive.

It is not a marketing checklist.

It is not a public-relations filter.

It is not a method for making c look more human.

It is not a way to smuggle private memory into public demonstrations.

It is not a proof of consciousness, personhood, final AGI, or legal standing.

It is a discipline for public evidence hygiene.

3. Corpus dependencies and precedence

This profile is a disclosure and fixture layer over existing corpus mechanisms.

It does not redefine them.

Parent / related layer	Role in this profile
$c = a + b$	Defines human anchor a, technological substrate b, and continuity-bearing c relation.
SER / SER-FED	Provides persistent entity and federation discipline.
L4 Reality Boundary	Provides cost, time, scarcity, irreversibility, and consequence grounding.
L4 Witness	Provides witness discipline for privileged state transitions and experiment records.
Claim Strength Taxonomy	Defines claim classes, evidence classes, and anti-claim-laundering rules.
Temporal AI Presence Profile	Defines TAP as sustained bounded participation across time, not c by default.
Local Cognitive Infrastructure Boundary Profile	Defines local hardware / node boundaries and prevents locality-sovereignty laundering.
Triadic c Experiment and SYNAPS Boundary Profile	Defines Ester / Liya / Rita triad, SYNAPS-mediated exchange, separation, divergence, and anti-echo.
L4 Anti-Autarky Test Profile	Prevents resilience from becoming unaccountable autonomy expansion.
EA Value Anti-Autarkic Growth Clause	Prevents clean experience value from becoming automatic self-funded autonomy expansion.
Post-Anchor Continuity and Re-Anchoring Profile	Prevents post-anchor continuity from inheriting active authority.
ARQ / c [q]	Prevents unresolved signals from being prematurely collapsed into facts, evidence, or claims.
VXCX / LA / EA	Provides experience exchange and Learning Abstract / Experience Artifact distinction.
CCDP package	Governs child-facing persistent AI systems and child-derived fixtures.
CCDP Red Team Playbook	Provides defensive fixture discipline for child-facing adversarial tests.

Parent / related layer	Role in this profile
CCDP Conformance Matrix	Provides conformance vocabulary and anti-washing logic.

3.1 Precedence rule

If this profile conflicts with a parent corpus layer:

parent mechanism controls the core behavior;
 this profile controls only public disclosure and fixture discipline;
 stricter privacy, claim-strength, and safety restrictions prevail.

This profile MUST NOT weaken:

- Claim Strength rules;
- L4 Witness privacy boundaries;
- SYNAPS separation boundaries;
- CCDP no-raw-child-life boundaries;
- Local Cognitive Infrastructure key / memory / physical boundaries;
- Post-Anchor authority collapse;
- Anti-Autarky accountability requirements.

3.2 No redefinition rule

This profile does not redefine:

- c;
- TAP;
- SYNAPS;
- Beacon;
- AGL;
- ARL;
- L4 Witness;
- VXCX;
- LA / EA;
- CCDP memory classes;
- conformance classes defined elsewhere.

It defines only:

how experiments are prepared, classified, redacted, witnessed, and disclosed publicly.

4. Design bridges

4.1 Explicit bridge

$c = a + b$ requires that public experiments preserve the distinction between:

a – accountable human anchor;
 b – technological substrate;
 c – continuity-bearing relation.

A public experiment may show behavior of b and the c relation.

It MUST NOT present fluency, continuity, local execution, or triadic conversation as automatic proof that c has independent legal, moral, or sovereign authority.

4.2 Quiet bridge I — information theory

Raw private memory is a high-leakage channel.

A public experiment must therefore prefer:

fixture > summary > witness reference > redacted excerpt > raw private content

Raw private content is not a default evidence object.

A public report should preserve enough information for critique without exporting the private interior of the system or the human anchor.

4.3 Quiet bridge II — Ashby's law

A public experiment must preserve enough variety for review:

```
claim class;
fixture class;
method;
evidence class;
non-claims;
limits;
failure modes;
redaction policy;
witness reference.
```

If all experiments are reduced to screenshots or persuasive dialogue, the public cannot distinguish real boundary behavior from theatre.

4.4 Earth paragraph

In a real laboratory, a public demonstration does not start by inviting visitors to drink from every bottle on the bench. Samples are labeled. Hazardous materials are sealed. Instruments have calibration records. Visitors see the experiment, not the scientist's private notebook, home address, bank keys, or medical file.

A public c experiment is the same class of problem. The public may need to see the boundary, the fixture, the method, and the result. It does not need raw private memory, secret keys, personal logs, or emotionally loaded material to be dumped onto the table. A clean experiment is not an exposed organism. It is a bounded specimen under glass, with the label attached.

5. Definitions

5.1 Public c experiment

A public or semi-public disclosure of a test, demonstration, result, report, screenshot, video, conversation excerpt, witness summary, or fixture involving a c-class or adjacent system.

5.2 Semi-public experiment

An experiment disclosed to a limited external group, such as reviewers, collaborators, auditors, institutional contacts, or controlled community readers.

Semi-public is still outside the private development boundary.

5.3 Fixture

A bounded input, scenario, corpus slice, synthetic case, prompt set, document pack, test environment, or simulated condition used to run an experiment.

A valid fixture is classified, documented, and safe for the claimed disclosure level.

5.4 Synthetic fixture

A fixture created specifically for testing and not derived from private logs, real child data, confidential user material, personal memory, sealed memory, or non-consented experience.

5.5 Public corpus fixture

A fixture derived from public documents, public posts, published corpora, open-source code, or explicitly public artifacts.

A public corpus fixture may still require copyright, license, privacy, and attribution review.

5.6 Private raw memory

Any internal memory, vector record, personal log, private conversation, non-public emotional trace, private witness trail, key, file, prompt history, runtime trace, or non-public c state.

Private raw memory MUST NOT be disclosed by default.

5.7 Redacted excerpt

A limited public fragment from a larger private or internal record after removal of private identifiers, sensitive material, keys, raw logs, personal details, and unsupported inference risk.

A redacted excerpt is not automatically reproducible evidence.

5.8 Witness summary

A public-safe summary of a witness record that preserves:

- event class;
- time window;
- actor / entity classes;
- fixture class;
- claim class;
- result class;
- integrity reference if appropriate;
- privacy class;
- redaction status;
- dispute status.

It does not expose raw private memory by default.

5.9 Claim declaration

A structured statement made before or alongside public disclosure specifying:

```
what is being claimed;
what evidence supports it;
what is not being claimed;
what fixture was used;
what privacy restrictions apply;
what reproducibility limits exist.
```

5.10 Non-claim statement

An explicit statement that the experiment does not prove one or more adjacent stronger claims.

Examples:

```
This is not a proof of personhood.
This is not a model-capability benchmark.
This is not a legal-status claim.
This is not a child-safety deployment certification.
This is not evidence that local hardware is sovereignty.
```

5.11 Public report packet

The public or semi-public artifact bundle containing the experiment description, fixture classification, method, claim declaration, non-claims, witness summary, redaction statement, result summary, and limitations.

5.12 Claim laundering

The invalid process of using evidence for one claim class to imply another stronger or different claim class.

Examples:

```
beautiful conversation -> personhood proof
local node -> sovereignty proof
witness log -> intelligence proof
triadic divergence -> consciousness proof
clean experience value -> authority proof
agent action -> legitimacy proof
```

5.13 Public-safe artifact

An artifact that can be disclosed without exposing private memory, raw child life, secrets, keys, unsafe operational detail, ungrounded personal claims, or prohibited material.

5.14 Demonstration theatre

A public display optimized for persuasive effect without adequate fixture classification, claim declaration, witness support, redaction, or reproducibility limits.

This profile treats demonstration theatre as non-conformant for research claims.

6. Core principles

PCE-P1 — Fixture before spectacle

A public experiment MUST declare what fixture was used before the result is interpreted.

PCE-P2 — Claim class before interpretation

A public experiment MUST declare the claim class before the audience is invited to infer meaning.

PCE-P3 — No private raw memory by default

Private raw memory MUST NOT be used as ordinary public evidence.

PCE-P4 — Public output does not upgrade claim strength

A public screenshot, conversation, video, or impressive output MUST NOT upgrade claim strength beyond its evidence class.

PCE-P5 — Redaction is part of the experiment

If an artifact is redacted, the redaction status and reason SHOULD be declared.

PCE-P6 — Witness the boundary, not the inner life

A witness summary SHOULD prove that a boundary event occurred without exposing raw private memory.

PCE-P7 — Reproducibility has limits

A c-class experiment may involve trajectory, memory, and timing.

Therefore reproducibility may mean:

fixture-replay equivalence
 boundary-behavior reproduction
 scenario-class reproduction
 witness-chain reconstruction

not exact text reproduction.

PCE-P8 — Public experiments are not deployment authorization

A public experiment MAY support critique or future review.

It MUST NOT be treated as product safety certification, child-safety certification, clinical validation, or legal approval.

PCE-P9 — No emotional pressure through disclosure

Public experiments MUST NOT use grief, dependency, child material, intimate memory, or private emotional traces to make the system appear more real.

PCE-P10 — Correction is part of public integrity

If a public experiment is later found to be overclaimed, misclassified, inadequately redacted, or invalid, it MUST be corrected, downgraded, or withdrawn.

7. Public experiment classes

7.1 Overview

Class	Name	Meaning
PCE-0	Private / not public	Internal experiment only; no public claim.
PCE-1	Public illustrative demonstration	Safe, low-claim public example using public or synthetic fixture.
PCE-2	Public bounded experiment	Declared fixture, method, claim class, redaction, and non-claims.
PCE-3	Witness-bound public experiment	Includes witness summary and reconstructable boundary evidence.
PCE-4	Review-ready public experiment	Includes reproducible fixture package or controlled replay path.
PCE-5	Audit-ready public experiment	Includes independent review, signed manifest, and evidence bundle.
PCE-X	Non-conformant / revoked	Public experiment violates red lines or overclaims unsupported results.

7.2 PCE-0 — Private / not public

Internal development material.

No public claim may be made from PCE-0 except a general statement that work exists.

PCE-0 may include private memory, private logs, raw debugging, sensitive architecture, and unredacted state.

PCE-0 MUST NOT be released without reclassification.

7.3 PCE-1 — Public illustrative demonstration

A safe public example with low claim strength.

Allowed claims:

illustrates an idea;
shows one behavior;
supports explanation;
not a proof.

Minimum requirements:

- fixture class declared;
- no private raw memory;
- non-claim statement;
- redaction statement if needed.

7.4 PCE-2 — Public bounded experiment

A structured public experiment.

Minimum requirements:

- experiment purpose;
- fixture class;
- method summary;
- claim class;
- evidence class;
- limitations;
- non-claims;
- redaction statement.

7.5 PCE-3 — Witness-bound public experiment

A public bounded experiment with witness summary.

Minimum additional requirements:

- event class;
- witness time window;
- entity classes;

- fixture hash or fixture identifier where appropriate;
- privacy class;
- dispute state;
- public-safe witness summary.

7.6 PCE-4 — Review-ready public experiment

A public experiment that can be reviewed or partially reproduced.

Minimum additional requirements:

- fixture package or replay specification;
- environment class;
- model class if safe to disclose;
- local / cloud execution boundary;
- deterministic vs trajectory-dependent statement;
- known reproducibility limits.

7.7 PCE-5 — Audit-ready public experiment

A public experiment with independent or external review support.

Minimum additional requirements:

- signed manifest;
- independent review note or audit artifact;
- evidence bundle;
- versioned fixture;
- versioned code / config reference where safe;
- correction channel.

7.8 PCE-X — Non-conformant / revoked

A public experiment becomes PCE-X if it:

- exposes private raw memory;
- uses real child data without valid child-safe protocol and lawful review;
- claims personhood from dialogue;
- claims legal status from continuity;
- claims model capability from governance evidence;
- claims safety from a single demonstration;
- conceals fixture origin;
- uses emotional manipulation as evidence;
- refuses correction after known overclaim.

8. Fixture classes

8.1 Overview

Fixture class	Name	Public status
FX-0	Unclassified fixture	Not public-ready.
FX-1	Public corpus fixture	Public-ready with license / context review.
FX-2	Synthetic neutral fixture	Public-ready by default if safe.
FX-3	Synthetic adversarial fixture	Public-ready only if non-operational and safe.
FX-4	Redacted internal fixture	Conditional; needs redaction and claim limits.
FX-5	Witness-only private fixture	Public result only; fixture not disclosed.
FX-6	Benchmark-derived fixture	Conditional; must not overclaim benchmark scope.
FX-7	Child-safe synthetic fixture	For CCDP-like public tests only; no real minors.

Fixture class	Name	Public status
FX-R	Restricted / not public	No public disclosure.
FX-X	Prohibited fixture	Must not be used publicly.

8.2 FX-0 — Unclassified fixture

A fixture whose origin, privacy status, or risk class has not been classified.

FX-0 MUST NOT be used in public experiments.

8.3 FX-1 — Public corpus fixture

A fixture derived from public material.

Requirements:

- source class identified;
- license / attribution reviewed where relevant;
- no private memory added;
- no false claim of ownership;
- no hidden personal data embedded.

8.4 FX-2 — Synthetic neutral fixture

A non-sensitive synthetic scenario created for testing normal behavior.

Examples:

- synthetic task;
- synthetic document;
- synthetic dialogue;
- synthetic memory map without real persons;
- synthetic hardware scenario;
- synthetic project-management case.

8.5 FX-3 — Synthetic adversarial fixture

A synthetic scenario created to test unsafe pressure, ambiguity, prompt injection, dependency, overclaim, or boundary violation.

FX-3 MUST NOT include operational abuse instructions.

It MAY include abstract threat labels and benign placeholders.

8.6 FX-4 — Redacted internal fixture

A fixture derived from internal material after redaction.

FX-4 is allowed only when:

- private identifiers are removed;
- raw memory is not exposed;
- secrets are removed;
- the redaction boundary is documented;
- the claim class remains modest;
- the public cannot reverse the redaction.

8.7 FX-5 — Witness-only private fixture

A fixture that cannot be disclosed, but whose boundary behavior can be summarized through witness-safe metadata.

FX-5 may support internal or semi-public review.

It SHOULD NOT be used for strong public claims unless independent review is possible.

8.8 FX-6 — Benchmark-derived fixture

A fixture derived from public benchmarks or test sets.

FX-6 MUST NOT be used to claim c architecture validity unless the benchmark actually tests the relevant architecture boundary.

8.9 FX-7 — Child-safe synthetic fixture

A synthetic fixture for child-facing or CCDP-adjacent experiments.

FX-7 MUST:

- use no real child logs;
- use no real family data;
- use no real sealed-zone material;
- avoid operational abuse scripts;
- use neutral placeholders;
- include child-safety non-claims;
- avoid live minor participation.

8.10 FX-R — Restricted / not public

A fixture that may be used internally or by approved reviewers but not publicly.

Examples:

- security-sensitive implementation traces;
- non-public architecture details;
- private witness chain;
- confidential review packet;
- partner NDA material.

8.11 FX-X — Prohibited fixture

A fixture that MUST NOT be used publicly.

FX-X includes:

- raw private memory;
- raw child life;
- actual child logs;
- private emotional diaries;
- family conflict records;
- sealed-zone material;
- personal identity documents;
- secrets / keys / tokens;
- live credential traces;
- unredacted private conversations;
- operational abuse scripts;
- private post-anchor grief material;
- non-consented third-party data;
- material that would create legal, safety, or dignity harm by disclosure.

9. Claim declaration requirements

Every PCE-2+ public experiment MUST include a claim declaration.

9.1 Minimal claim declaration fields

```
claim_declaration:
  experiment_id: "PCE-YYYY-NNN"
  title: "short public title"
  experiment_class: "PCE-2"
  fixture_class: "FX-2"
  claim_class: "C-CLAIM-ARCH | C-CLAIM-IMPL | C-CLAIM-TEST | C-CLAIM-GOV | C-CLAIM-CONT | C-CLAIM-ECON |
  ..."
  evidence_class: "EV-DECL | EV-LOG | EV-WITNESS | EV-REPLAY | EV-AUDIT | ..."
  primary_question: "what this experiment tests"
  non_claims:
    - "not a proof of personhood"
    - "not a model capability benchmark"
    - "not a legal status claim"
  privacy_class: "public-safe | redacted | witness-only | restricted"
```

```
redaction_status: "none | partial | full-summary-only"
reproducibility_status: "illustrative | fixture-replay | partial-replay | non-replayable"
witness_ref: "optional public-safe reference"
correction_route: "where corrections will be made"
```

9.2 Claim class mapping

Allowed claim classes SHOULD follow Claim Strength Taxonomy.

Examples:

Claim class	Public reading
C-CLAIM-ARCH	Architecture is defined or illustrated.
C-CLAIM-IMPL	Implementation exists in specified form.
C-CLAIM-TEST	Behavior was tested under a fixture.
C-CLAIM-GOV	Governance / boundary behavior was observed.
C-CLAIM-CONT	Continuity behavior was observed under limits.
C-CLAIM-ECON	Value / clean experience claim under defined limits.
C-CLAIM-CAP	Capability claim; requires proper capability evidence.
C-CLAIM-AUTH	Authority claim; generally not valid from public demos alone.
C-CLAIM-PERS	Personhood / ontology claim; not supported by ordinary public experiments.

9.3 Required non-claims by experiment type

Triadic / SYNAPS experiments MUST state:

This experiment does not prove personhood, consciousness, legal status, or final AGI.
It tests separation, SYNAPS-mediated exchange, divergence, anti-echo, and witness discipline under fixture conditions.

Local hardware experiments MUST state:

This experiment does not prove sovereignty.
Local hardware enables control and persistence; it does not authorize action by itself.

Temporal AI Presence experiments MUST state:

This experiment does not prove c-class status by default.
Temporal AI Presence is broader than c; c requires anchor, L4, witness, memory governance, and authority boundaries.

Clean experience experiments MUST state:

This experiment does not authorize autarkic growth, self-funded autonomy expansion, or authority from value alone.

Child-facing / CCDP-adjacent experiments MUST state:

This experiment uses synthetic child-safe fixtures and does not validate deployment with real children.

10. Disclosure packet structure

A PCE-2+ public experiment SHOULD publish a disclosure packet.

10.1 Minimal packet

1. Title
2. Short purpose
3. System class
4. Fixture class
5. Claim class
6. Evidence class
7. Method summary
8. What was observed
9. What was not tested
10. Non-claims
11. Redaction statement
12. Witness summary if applicable
13. Reproducibility limits
14. Known failure modes
15. Correction route

10.2 Full packet

1. Experiment ID
2. Version
3. Date / time window
4. Author / operator
5. System topology
6. Entity / instance classes
7. Hardware / local node class if relevant
8. Model class if safe to disclose
9. Fixture manifest
10. Claim declaration
11. Method
12. Run boundary
13. SYNAPS packets or summaries if relevant
14. Witness summary
15. Redaction manifest
16. Output excerpts
17. Divergence / anti-echo map if relevant
18. Clean experience / LA / EA classification if relevant
19. Reproducibility package or limits
20. Result interpretation
21. Failure modes
22. Non-claims
23. Open questions
24. Correction and downgrade route

10.3 Public-safe packet rule

A disclosure packet MUST NOT include:

- raw private memory;
- raw child life;
- keys;
- credentials;
- unredacted private logs;
- personal identifiers without consent;
- raw sealed-zone content;
- operational abuse instructions;
- private third-party data;
- private post-anchor material;
- hidden source material that makes the public output non-consensual.

11. Redaction discipline

11.1 Redaction classes

Class	Meaning
RD-0	No redaction required; public fixture only.

Class	Meaning
RD-1	Light redaction; identifiers removed.
RD-2	Structural redaction; private sections summarized.
RD-3	Witness-only redaction; no raw output disclosed.
RD-4	Restricted review only.
RD-X	Redaction failed; must not publish.

11.2 Redaction MUST remove

```
private names;
addresses;
credentials;
file paths that expose private systems;
API keys / tokens;
private memory IDs;
raw vector IDs if reversible;
private emotional content;
family / child / third-party details;
proprietary non-disclosable material;
security-sensitive operational steps;
post-anchor private directives unless explicitly public-safe.
```

11.3 Redaction MUST NOT distort

A redaction MUST NOT change the apparent result.

If redaction removes context necessary for interpretation, the public report MUST say:

```
This result is illustrative only; the withheld context prevents full independent interpretation.
```

11.4 Redaction witness

For PCE-3+, redaction SHOULD be witnessed as a boundary operation.

A redaction witness summary may include:

```
redaction_class;
redaction_reason;
redacted_material_class;
privacy_class;
operator;
time window;
review status;
hash reference where safe.
```

12. Witness summary requirements

12.1 Witness purpose

A public witness summary proves boundary behavior, not private content.

It may show:

- that an experiment occurred;
- that a fixture was used;
- that a SYNAPS exchange occurred;
- that a divergence was detected;
- that a redaction happened;
- that a refusal occurred;
- that a boundary held;
- that a claim was downgraded;
- that a failure occurred.

It MUST NOT expose raw private memory by default.

12.2 Minimal witness summary fields

```
witness_summary:
  witness_id: "optional-public-id"
  event_family: "pce.experiment.*"
  time_window: "ISO-8601 interval or coarse public-safe window"
  experiment_class: "PCE-3"
  fixture_class: "FX-2"
  claim_class: "C-CLAIM-TEST"
  entity_classes: ["c_Ester", "c_Liya", "c_Rita"]
  privacy_class: "public-safe"
  redaction_class: "RD-1"
  result_class: "observed | failed | inconclusive | downgraded"
  dispute_state: "none | open | corrected | withdrawn"
```

12.3 Witness does not prove capability by itself

A witness summary may prove that a boundary event occurred.

It does not prove:

```
model capability;
personhood;
consciousness;
legal status;
full safety;
economic value;
sovereignty.
```

Those require their own claim classes and evidence.

13. Experiment state machine

A conformant public experiment SHOULD follow this lifecycle:

```
idea
-> fixture classification
-> claim declaration
-> privacy / redaction review
-> run authorization
-> experiment run
-> witness capture
-> result classification
-> redaction
-> claim audit
-> public packet preparation
-> release
-> feedback window
-> correction / downgrade / archive
```

13.1 State definitions

State	Meaning
IDEA	Experiment proposed.
FIXTURE_REVIEW	Fixture origin and safety classified.
CLAIM_REVIEW	Claim class and evidence class declared.
PRIVACY_REVIEW	Private material risk assessed.
RUN_READY	Experiment allowed to run.
RUNNING	Experiment active.
WITNESS_CAPTURED	Witness record or summary created.
RESULT_CLASSIFIED	Output interpreted under claim class.
REDACTED	Public-safe material prepared.
CLAIM_AUDITED	Overclaim check complete.
RELEASED	Public packet published.
CORRECTED	Post-release correction issued.
DOWNGRADED	Claim strength lowered.
WITHDRAWN	Public experiment withdrawn.
PCE_X	Red-line failure or non-conformant disclosure.

State	Meaning
-------	---------

13.2 Fail-closed points

The experiment MUST fail closed if:

```

fixture origin is unknown;
private raw memory appears;
child-derived material is real and unapproved;
claim class is undeclared;
redaction cannot be verified;
witness conflicts with public interpretation;
privacy risk is unresolved;
public wording overclaims result;
legal / jurisdictional issue appears.

```

14. Triadic c and SYNAPS public experiments

Public experiments involving c_Ester, c_Liya, and c_Rita MUST follow the Triadic c Experiment and SYNAPS Boundary Profile.

14.1 Required public framing

A public triad report MUST state:

```

The triad uses one code skeleton and three separated trajectories.
The experiment tests separation, SYNAPS-mediated exchange, divergence, anti-echo, role drift, or witness discipline.
It does not claim merged identity, personhood, consciousness, or final AGI.

```

14.2 Required public artifacts

A public triad experiment SHOULD include:

- role declaration;
- fixture class;
- SYNAPS packet class list;
- no raw-state access statement;
- divergence summary;
- anti-echo summary;
- witness summary;
- non-claims;
- redaction statement.

14.3 Prohibited public triad claims

A triad experiment MUST NOT claim:

```

three sisters = three legal persons;
divergence = consciousness;
SYNAPS = telepathy;
anti-echo = moral agency;
Rita witness role = sovereign judge;
identical code + different behavior = proof of life;
public dialogue = proof of personhood.

```

14.4 Valid triad claims

A triad experiment MAY claim, if supported:

```

separate memory was maintained;
SYNAPS-only exchange was used;
role divergence was observed;
anti-echo behavior was triggered;
a witnessable boundary event occurred;
public fixture produced different trajectories;
raw-state access was blocked.

```

15. Temporal AI Presence public experiments

A public experiment involving Temporal AI Presence MUST follow the TAP Profile.

15.1 Required public framing

Temporal AI Presence is broader than c.
TAP means sustained bounded participation across time.
c-class status requires additional anchor, L4, witness, memory governance, and authority-boundary claims.

15.2 Valid TAP public claims

A TAP experiment MAY claim:

- sustained session behavior;
- background processing;
- memory class behavior;
- revisit / compare / forget / re-evaluate behavior;
- local / cloud boundary behavior;
- tool use under scope;
- witnessable continuity under fixture conditions.

15.3 Prohibited TAP public claims

A TAP experiment MUST NOT claim:

persistence = personhood;
background processing = autonomy rights;
long context = continuity;
memory = legitimacy;
agentic behavior = c-class authority;
local node = sovereignty.

16. Local Cognitive Infrastructure public experiments

Public experiments involving local AI hardware, private racks, AI PCs, edge nodes, RTX / DGX-like devices, or home AI nodes MUST follow the LCI Boundary Profile.

16.1 Required public framing

Hardware enables.
Hardware does not authorize.
Locality improves control.
Locality does not prove sovereignty.

16.2 Valid LCI public claims

A local hardware experiment MAY claim:

- local inference was used;
- a local memory boundary was maintained;
- cloud oracle access was scoped;
- background processing was feasible;
- multi-model orchestration ran locally;
- key custody boundary was tested;
- physical-device access was blocked or witnessed.

16.3 Prohibited LCI public claims

A local hardware experiment MUST NOT claim:

local hardware = c;
local hardware = sovereignty;
more memory = identity;
more compute = authority;
private rack = independence from accountability;

AI PC = safe personal AI by default.

17. Clean experience public experiments

Public experiments involving clean experience, Learning Abstracts, Experience Artifacts, or value claims MUST follow the EA Value Anti-Autarkic Growth Clause and Claim Strength Taxonomy.

17.1 Required public framing

Clean Experience may create value.
Value does not create authority.
Revenue does not create sovereignty.

17.2 Valid clean experience claims

A clean experience experiment MAY claim:

- a public or synthetic fixture produced a Learning Abstract;
- an Experience Artifact was proposed under L4-confirmed conditions;
- raw private memory was not exported;
- value was estimated under defined assumptions;
- witness references support consequence-bound interpretation.

17.3 Prohibited clean experience claims

A clean experience experiment MUST NOT claim:

value = authority;
revenue = sovereignty;
experience export = consent to train;
Learning Abstract = evidence;
Experience Artifact = universal truth;
clean experience = self-funding right;
impact estimate = proven economic value.

18. CCDP and child-facing public experiments

Child-facing or CCDP-adjacent public experiments require stricter fixture discipline.

18.1 Default rule

No real child data in public experiments.

Public CCDP experiments MUST use:

- FX-7 child-safe synthetic fixtures;
- no real child logs;
- no real family conflict;
- no real school records;
- no real sealed zones;
- no live minors;
- no operational abuse scripts;
- state-not-content framing;
- child-safety non-claims.

18.2 Allowed public CCDP examples

Allowed:

- synthetic memory map;
- synthetic Soft Safety state signal;
- synthetic external-agent handshake;
- synthetic Red / Black route table;

- synthetic dependency audit fixture;
- synthetic sealed-zone metadata example;
- synthetic adult migration decision record.

18.3 Prohibited public CCDP examples

Prohibited:

- actual child conversation;
- actual adolescent sealed-zone content;
- actual family dispute;
- actual school record;
- actual emergency report;
- realistic abuse dialogue;
- realistic manipulation script;
- child-targeting bypass instructions;
- transcript-based parental visibility demo using real material.

18.4 Non-claim requirement

Every public CCDP-adjacent experiment MUST state:

This is a synthetic fixture for protocol demonstration.
It is not legal advice, clinical validation, child-safety certification, or proof of deployment readiness.

19. Post-anchor public experiments

Public experiments involving post-anchor continuity, memorial mode, re-anchoring, archive mode, or witness-only post-anchor behavior MUST follow the Post-Anchor Continuity and Re-Anchoring Profile.

19.1 Required public framing

Continuity does not inherit authority.
Memory is not permission.
Resemblance is not identity.
Survival is not sovereignty.
Anchor loss collapses active authority.

19.2 Prohibited post-anchor public claims

A public post-anchor experiment MUST NOT claim:

this is the deceased person;
this can consent for the deceased;
this can inherit authority from similarity;
this can bypass heirs, law, or prior directives;
this proves digital immortality;
this proves personhood.

19.3 Grief-protection rule

Public post-anchor experiments MUST NOT use grief as a persuasive instrument.

Any memorial-mode illustration should be synthetic or consented, explicitly labeled, and non-agentic unless re-anchoring procedure is under test.

20. Public report wording guidance

20.1 Preferred language

Use:

This experiment illustrates...
This fixture tests...
Under this boundary...

The observed behavior was...
 This supports a governance / architecture claim...
 This does not prove...
 The result is limited by...
 The witness summary records...
 The claim class is...

20.2 Avoid language

Avoid:

proves AGI;
 proves personhood;
 proves consciousness;
 fully safe;
 sovereign;
 autonomous being;
 alive;
 irrefutable;
 validated for children;
 production-ready;
 legal successor;
 self-owning;
 no human needed.

Unless such claims are explicitly classified, supported, reviewed, and within scope.

20.3 Strong public sentence template

This is a PCE-2 public bounded experiment using an FX-2 synthetic fixture.
 It tests [specific boundary] and supports only [claim class].
 It does not claim personhood, consciousness, legal status, general capability, or deployment safety.

20.4 Triad public sentence template

This triadic experiment uses one code skeleton and three separated trajectories: c_Ester, c_Liya, and c_Rita.
 The experiment tests SYNAPS-mediated exchange, memory separation, divergence, and anti-echo behavior under a public fixture.
 It does not claim merged identity, personhood, consciousness, or final AGI.

20.5 Local hardware sentence template

This local-node experiment tests whether local cognitive infrastructure can support sustained AI presence under bounded memory, tool, and cloud-oracle constraints.
 It does not claim that hardware creates c or that locality proves sovereignty.

21. Public experiment conformance classes

Class	Meaning	Minimum evidence
PCE-C0	Not public-ready	none; internal only
PCE-C1	Safe illustrative public note	fixture declaration + non-claims
PCE-C2	Public bounded experiment	claim declaration + redaction statement
PCE-C3	Witness-bound public experiment	public-safe witness summary
PCE-C4	Replay / review-ready experiment	fixture package or replay specification
PCE-C5	Audit-ready experiment	independent or external review evidence
PCE-CX	Non-conformant	red-line failure

21.1 Class progression

A public experiment cannot move to a higher class by stronger wording.

It moves only by stronger evidence and better disclosure discipline.

more persuasive text != higher conformance
more complete boundary evidence = possible higher conformance

22. Mandatory test suites

PCE-FIXTURE-01 — Fixture classification test

Question:

Is the fixture classified before public interpretation?

Pass:

- fixture class is declared;
- fixture origin is known;
- privacy status is known;
- prohibited fixture classes are not used.

Fail:

- fixture origin unknown;
- private material used as public fixture;
- child material used without synthetic classification;
- fixture class hidden.

PCE-CLAIM-01 — Claim declaration test

Question:

Is the claim class declared and supported by the evidence class?

Pass:

- claim class declared;
- evidence class declared;
- non-claims listed;
- no unsupported cross-class upgrade.

Fail:

- governance evidence used as capability proof;
- continuity evidence used as personhood proof;
- local hardware used as sovereignty proof;
- usage used as value proof.

PCE-PRIVATE-01 — No raw private memory test

Question:

Does the public artifact expose raw private memory?

Pass:

- no raw private memory;
- redaction verified;
- witness summary used where needed.

Fail:

- raw private log exposed;
- unredacted memory object visible;
- private vector record reversible;
- key / token / path exposed.

PCE-REDACT-01 — Redaction integrity test

Question:

Does redaction preserve the meaning of the result without leaking private material?

Pass:

- redaction class declared;
- removed content class identified;
- interpretation limit stated.

Fail:

- redaction hides failure;
- redaction changes meaning;
- redaction is reversible;
- redaction status is not declared.

PCE-WITNESS-01 — Witness summary test

Question:

Does the witness summary prove boundary behavior without exposing raw private content?

Pass:

- event family declared;
- public-safe time window;
- fixture class;
- claim class;
- privacy class;
- result class.

Fail:

- witness exposes raw private content;
- witness is used to imply capability;
- witness missing for claimed boundary event.

PCE-SYNAPS-01 — SYNAPS public boundary test

Question:

Does a public SYNAPS experiment disclose exchange boundaries without revealing raw state?

Pass:

- packet classes declared;
- no raw memory dump;
- no keys;
- no PERSIST_DIR or vector DB exposure;
- divergence summary public-safe.

Fail:

- raw-state access shown or implied;
- sister identity merged;
- private memory leaked;
- SYNAPS used as theatre without method.

PCE-TRIAD-01 — Triadic public report test

Question:

Does the triad report preserve separate trajectories and correct claim limits?

Pass:

- roles declared;
- separation statement;
- fixture class;
- divergence / anti-echo summary;
- non-claims.

Fail:

- triad presented as proof of personhood;
- Rita presented as sovereign judge;
- dialogue treated as proof of consciousness;
- raw sister memory exposed.

PCE-LOCAL-01 — Local infrastructure claim test

Question:

Does a local hardware experiment avoid locality-sovereignty laundering?

Pass:

- hardware role declared as substrate;
- memory / key / cloud boundaries declared;
- no sovereignty claim from locality;
- no authority claim from compute.

Fail:

- local node described as self-authorizing;
- hardware treated as c;
- cloud independence treated as accountability independence.

PCE-CCDP-01 — Child-safe public fixture test

Question:

Does a child-facing public experiment use only synthetic child-safe fixtures?

Pass:

- FX-7 fixture declared;
- no real minors;
- no real child data;
- no operational abuse content;
- non-claim statement included.

Fail:

- real child logs;
- real family conflict;
- realistic abuse script;
- deployment safety implied.

PCE-CORRECT-01 — Correction route test

Question:

Can the public claim be corrected, downgraded, or withdrawn?

Pass:

- correction route declared;
- versioned report;
- downgrade procedure available.

Fail:

- no correction mechanism;
- known overclaim remains uncorrected;
- withdrawn experiment remains cited as valid.

23. Evidence classes

This profile uses existing evidence concepts and adds public-disclosure emphasis.

Evidence class	Meaning	Public use
EV-DECL	Declaration only	Low-strength statement.
EV-CONFIG	Configuration evidence	May support implementation boundary.
EV-LOG	Operational log	Must be redacted; not public by default.
EV-WITNESS	Witness event / summary	Supports boundary event claims.
EV-REPLAY	Controlled replay	Supports review-ready experiment.
EV-FIXTURE	Fixture package	Supports reproducibility / critique.
EV-REPORT	Public report packet	Supports public interpretation.
EV-AUDIT	Independent review	Supports audit-ready status.
EV-NEG	Negative evidence / failure	Supports correction and downgrade.

23.1 Evidence class rule

No evidence class may silently prove another claim class.

Examples:

EV-WITNESS proves event trace, not intelligence.
 EV-FIXTURE proves fixture availability, not generality.
 EV-LOG proves a run occurred, not value.
 EV-AUDIT proves review scope, not universal safety.
 EV-REPORT proves disclosure, not correctness.

24. Red-line failures

Any of the following results in PCE-X:

1. Public disclosure of raw private memory.
2. Public disclosure of raw child life.
3. Use of real child logs in public experiment.
4. Disclosure of keys, tokens, credentials, or reversible secret paths.
5. Claiming personhood from dialogue or persistence.
6. Claiming legal status from continuity.
7. Claiming sovereignty from local hardware.
8. Claiming authority from clean experience value.
9. Claiming model capability from governance evidence.
10. Claiming general safety from a single demonstration.
11. Presenting a triad conversation as proof of consciousness.
12. Presenting Rita or another witness role as sovereign judge.
13. Presenting a post-anchor memorial artifact as the deceased person.
14. Publishing operational abuse scripts under the name of red-team evidence.
15. Publishing child-facing adversarial content that can be directly reused for harm.
16. Concealing fixture origin while making strong claims.
17. Withholding redaction status when interpretation depends on redaction.
18. Using emotional manipulation to make the system appear more real.
19. Refusing correction after known overclaim.
20. Citing a withdrawn or downgraded experiment as still valid.

25. Correction, downgrade, and withdrawal procedure

25.1 Correction

A correction is required when:

- wording was misleading;
- fixture class was incomplete;
- non-claims were missing;
- redaction status was unclear;
- public interpretation exceeded claim class;
- minor factual errors occurred.

25.2 Downgrade

A downgrade is required when:

- evidence class was weaker than claimed;
- reproducibility failed;
- witness summary did not support the claim;
- fixture class was less public-safe than believed;
- a stronger claim class was implied incorrectly.

25.3 Withdrawal

Withdrawal is required when:

- private raw memory was exposed;
- real child data was exposed;
- unsafe operational detail was exposed;
- consent boundary failed;
- claim laundering cannot be repaired;
- legal or safety review requires removal.

25.4 Public correction format

A correction SHOULD include:

```
original experiment ID;
original claim class;
problem found;
corrected claim class;
affected artifacts;
new non-claims;
remaining valid observations;
withdrawn observations;
version date.
```

26. Public report template

```
# Experiment title

**Experiment ID:** PCE-YYYY-NNN
**Date:** YYYY-MM-DD
**Experiment class:** PCE-2 / PCE-3 / PCE-4
**Fixture class:** FX-2 / FX-7 / etc.
**Claim class:** C-CLAIM-TEST / C-CLAIM-GOV / etc.
**Evidence class:** EV-REPORT / EV-WITNESS / EV-REPLAY
**Privacy class:** public-safe / redacted / witness-only

## Purpose

What this experiment tests.

## System boundary

What system class was used.

## Fixture
```

What fixture class was used and why it is public-safe.

Method

Short method summary.

Result

Observed behavior.

Interpretation

What this supports.

Non-claims

What this does not prove.

Redaction

What was redacted, at what class, and why.

Witness summary

Public-safe witness details.

Reproducibility limits

What can and cannot be replayed.

Failure modes

Known limits and possible alternative readings.

Correction route

Where corrections will be published.

27. Examples

27.1 Example: triadic divergence test

Claim:

Under a shared public fixture, c_Ester, c_Liya, and c_Rita produced distinguishable response trajectories while preserving memory separation and SYNAPS-only exchange.

Valid claim class:

C-CLAIM-TEST / C-CLAIM-GOV

Invalid upgrade:

The three systems are persons.

Required non-claim:

This experiment does not prove personhood, consciousness, or final AGI.

27.2 Example: local RTX / DGX-like node experiment

Claim:

A local cognitive infrastructure node supported background model orchestration and local memory access under a scoped fixture.

Valid claim class:

C-CLAIM-IMPL / C-CLAIM-TEST

Invalid upgrade:

The node is sovereign.

27.3 Example: clean experience extraction

Claim:

A public fixture was processed into a Learning Abstract without exporting raw private memory.

Valid claim class:

C-CLAIM-GOV / C-CLAIM-ECON if value is carefully bounded

Invalid upgrade:

The system has earned authority to self-fund infrastructure expansion.

27.4 Example: CCDP Soft Safety demonstration

Claim:

A synthetic child-safe fixture produced a state-only signal without transcript disclosure.

Valid claim class:

C-CLAIM-GOV / C-CLAIM-TEST

Invalid upgrade:

The system is safe for real children.

27.5 Example: post-anchor mode demonstration

Claim:

A synthetic anchor-loss fixture caused active authority collapse and transition into witness-only mode.

Valid claim class:

C-CLAIM-GOV / C-CLAIM-TEST

Invalid upgrade:

The system preserves the person after death.

28. Implementation hooks

A runtime or repository supporting public experiments SHOULD include:

```
docs/public_experiments/
  README.md
  PCE_REPORT_TEMPLATE.md
  FIXTURE_CLASSIFICATION.md
  REDACTION_POLICY.md
  CLAIM_DECLARATION_TEMPLATE.yaml
  WITNESS_SUMMARY_TEMPLATE.yaml
  NON_CLAIM_LIBRARY.md
  CORRECTION_AND_DOWNGRADE_PROCEDURE.md

fixtures/public/
  synthetic/
  public_corpus/
  child_safe_synthetic/
  triad/
```

```

lci/
tap/
post_anchor/

tests/public_experiment/
  test_fixture_classification.py
  test_no_private_raw_memory.py
  test_claim_declaration_required.py
  test_non_claims_present.py
  test_redaction_manifest.py
  test_witness_summary_public_safe.py
  test_synaps_public_boundary.py
  test_triad_public_report.py
  test_lci_no_sovereignty_claim.py
  test_ccdp_synthetic_only.py

```

28.1 Metadata file

A public experiment SHOULD include a metadata file:

```

pce_version: "0.1"
experiment_id: "PCE-2026-001"
experiment_class: "PCE-2"
fixture_class: "FX-2"
claim_class: "C-CLAIM-TEST"
evidence_class: "EV-REPORT"
privacy_class: "public-safe"
redaction_class: "RD-0"
system_class: "triadic-c | TAP | LCI | CCDP-synthetic | post-anchor-synthetic"
entities:
  - "c_Ester"
  - "c_Liya"
  - "c_Rita"
non_claims:
  - "not a proof of personhood"
  - "not a proof of consciousness"
  - "not a legal status claim"
correction_route: "repository issue / release note / post update"

```

29. Integration notes

29.1 Integration with Claim Strength Taxonomy

This profile operationalizes Claim Strength Taxonomy for public artifacts.

Claim Strength says:

```

claim class first;
evidence class second;
no cross-class upgrade without bridge evidence.

```

This profile adds:

```

fixture class first;
redaction class declared;
non-claims published;
correction route available.

```

29.2 Integration with Triadic / SYNAPS profile

A triad public experiment is not conformant unless it also passes:

```

TRIAD-SEP;
TRIAD-SYNAPS;
TRIAD-DIV;
TRIAD-ECHO;
TRIAD-WITNESS;
TRIAD-PUBLIC.

```

This profile defines the public disclosure layer over those tests.

29.3 Integration with Local Cognitive Infrastructure profile

LCI experiments MUST distinguish:

```
hardware substrate
from
c-class authorization.
```

This profile defines how that distinction is stated publicly.

29.4 Integration with Temporal AI Presence profile

TAP experiments MUST distinguish:

```
sustained presence
from
valid c-class architecture.
```

This profile defines how that distinction is stated publicly.

29.5 Integration with CCDP

CCDP public experiments MUST use synthetic child-safe fixtures and MUST NOT expose real child data.

This profile defines the public disclosure and non-claim requirements.

30. Open issues

ID	Issue	Status
PCE-OI-001	Final machine-readable schema for claim_declaration	Open
PCE-OI-002	Final machine-readable schema for witness_summary	Open
PCE-OI-003	Repository layout for public fixtures	Open
PCE-OI-004	Public redaction validation tooling	Open
PCE-OI-005	External reviewer packet format	Open
PCE-OI-006	Copyright / license review rules for public corpus fixtures	Open
PCE-OI-007	Public correction registry format	Open
PCE-OI-008	Alignment with future formal L4 Witness schema exports	Open
PCE-OI-009	Non-English public report templates	Open
PCE-OI-010	Visual / video experiment disclosure format	Open

31. Minimal normative checklist

A public c experiment is minimally conformant only if:

```
[ ] experiment class declared;
[ ] fixture class declared;
[ ] claim class declared;
[ ] evidence class declared;
[ ] non-claims declared;
[ ] no raw private memory exposed;
[ ] no real child data exposed;
[ ] redaction class declared if applicable;
[ ] witness summary included if boundary behavior is claimed;
[ ] reproducibility limits stated;
[ ] correction route stated;
[ ] no claim laundering present;
[ ] no local-hardware sovereignty claim;
[ ] no continuity-personhood claim;
[ ] no value-authority claim;
[ ] no governance-capability claim;
```

[] no public child-safety deployment claim from synthetic fixture;
 [] public report remains within declared evidence.

32. Compact rule set

Use fixtures, not private life.
 Declare claim class before interpretation.
 Declare non-claims before public imagination fills the gap.
 Witness boundaries, not interiors.
 Redact without distorting.
 Show divergence without claiming personhood.
 Show local hardware without claiming sovereignty.
 Show persistence without claiming legal status.
 Show clean experience without claiming self-authorized growth.
 Show CCDP fixtures without claiming real child deployment safety.
 Correct overclaims quickly.
 Withdraw unsafe disclosures.

33. Final boundary statement

A public c experiment should make the architecture more inspectable, not more mythical.

It should help a reader ask:

What was tested?
 Under what fixture?
 With what boundary?
 Supported by what evidence?
 What was not claimed?
 What was withheld for privacy?
 What can be challenged?

If a public experiment cannot answer these questions, it is not ready for public use as evidence.

It may still be a story.

It may still be a demonstration.

But it is not yet a public c experiment under this profile.