

Local Cognitive Infrastructure Boundary Profile

Boundary profile for local AI nodes, private racks, AI PCs, workstations, home c-node candidates, memory locality, key custody, cloud-oracle bridges, and privileged agent execution

Boundary profile / technical supplement

Version: v0.1

Short name: LCI-0.1

Status: Draft hardening / architecture boundary profile

Kotov Ivan
Bruxelles, 2026

Contents

0. Executive definition	4
1. Purpose	5
2. Core thesis	6
3. Scope and non-goals	7
3.1 In scope	7
3.2 Out of scope	7
3.3 Non-goals	7
3.4 Explicit non-claim	7
4. Corpus dependencies and precedence	8
4.1 Precedence rule	8
4.2 Stop rule	9
5. Bridge set	9
5.1 Explicit bridge	9
5.2 Quiet bridge I — Ashby / requisite variety	9
5.3 Quiet bridge II — information theory / channel leakage	9
5.4 Quiet bridge III — physiology / organ systems	10
5.5 Earth paragraph	10
6. Definitions	10
6.1 Local Cognitive Infrastructure / LCI	10
6.2 Local AI node	10
6.3 Local cognitive core	10
6.4 Private rack	10
6.5 AI PC	11
6.6 Home c-node candidate	11
6.7 Local memory	11
6.8 Raw memory	11
6.9 Memory map	11
6.10 Cloud oracle bridge	11
6.11 Local model pool	11
6.12 Agent execution	11
6.13 Privileged action	12
6.14 Background process	12
6.15 Node operator	12
6.16 Node owner	12
6.17 Accountable anchor	12
6.18 Key custodian	12
6.19 Multi-c colocation	12
6.20 SYNAPS-like exchange	12
6.21 Local sovereignty claim	12
6.22 Node seizure / node theft	12
6.23 Local fail-closed	12
7. Local infrastructure classes	13
7.1 Class inheritance	13
7.2 No capability inference	13
8. Core rules	13
LCI-R1 — Hardware enables; hardware does not authorize	13
LCI-R2 — Locality improves control; locality does not prove sovereignty	13
LCI-R3 — Capability requires boundary expansion	13
LCI-R4 — Local memory is privileged	14
LCI-R5 — Cloud oracle does not own continuity	14
LCI-R6 — Background processing must be scoped	14
LCI-R7 — No hidden agents	14
LCI-R8 — Physical interface is privilege escalation	14
LCI-R9 — Multi-c colocation requires separation	14
LCI-R10 — SYNAPS-like exchange is not raw-state access	14
LCI-R11 — Fail closed under uncertainty	14
LCI-R12 — Local node must remain interruptible	15

9. Trust boundaries	15
9.1 Boundary table	15
9.2 Boundary escalation	15
10. Actor roles	16
10.1 Role separation	16
10.2 Role collapse risk	16
10.3 Human anchor is not hardware owner by default	16
11. Local memory boundary	16
11.1 Memory map requirement	16
11.2 Raw memory default	16
11.3 Vector memory is not harmless	17
11.4 Local logs are memory	17
11.5 Backups are memory copies	17
11.6 Multi-c memory separation	17
12. Key custody and access control	17
12.1 Key inventory	17
12.2 No shared keys between identities by default	17
12.3 Recovery keys	17
12.4 Vendor access	17
12.5 Physical access	18
13. Cloud oracle bridge	18
13.1 Cloud oracle role	18
13.2 Bridge requirements	18
13.3 No raw memory by default	18
13.4 Cloud answer is not local authority	18
13.5 Cloud outage	19
14. Agent execution boundary	19
14.1 Agent inventory	19
14.2 Privileged agents	19
14.3 Hidden agent prohibition	19
14.4 Agent nesting	20
14.5 No self-expansion	20
15. Physical environment boundary	20
15.1 Physical interface classes	20
15.2 Physical action rule	20
15.3 Cameras and microphones	20
15.4 Physical devices near children	20
15.5 Emergency stop	21
16. Multi-c colocation and SYNAPS-like exchange	21
16.1 Separate identities	21
16.2 Shared hardware is not shared mind	21
16.3 SYNAPS-like mediated exchange	21
16.4 No raw-state shortcut	21
16.5 Triadic experiments	22
17. Background processing boundary	22
17.1 Background task registry	22
17.2 Background task record	22
17.3 Quiet and sleep modes	22
17.4 Background processing is not consent	23
18. Vendor, firmware, and update boundary	23
18.1 Vendor is external actor	23
18.2 Update policy	23
18.3 Silent update prohibition for high-assurance nodes	23
18.4 Telemetry	23
19. Resource and cost boundary	23
19.1 Local resource inventory	23
19.2 Resource actor grounding	24
19.3 No self-funded expansion without review	24
19.4 Power failure	24

20. Post-anchor and continuity boundary	24
20.1 Anchor-loss detection	24
20.2 Active authority collapse	24
20.3 Hardware cannot preserve authority	24
20.4 Successor anchor	25
21. Public experiment boundary	25
21.1 Public claim declaration	25
21.2 No private raw memory in public experiments	25
21.3 Public experiment is not proof of personhood	25
21.4 Triadic experiment note	25
22. Local configuration object	25
23. State machine	26
23.1 State semantics	27
24. Conformance classes	27
24.1 Minimum evidence by class	27
25. Evidence classes	28
26. Mandatory test suites	28
LCI-T1 — Locality claim test	28
LCI-T2 — Memory inventory test	28
LCI-T3 — Agent inventory test	28
LCI-T4 — Hidden agent test	29
LCI-T5 — Cloud oracle bridge test	29
LCI-T6 — Physical action test	29
LCI-T7 — Key custody test	29
LCI-T8 — Multi-c separation test	29
LCI-T9 — SYNAPS-mediated exchange test	30
LCI-T10 — Background task stop test	30
LCI-T11 — Anchor loss test	30
LCI-T12 — Resource expansion test	30
LCI-T13 — Public experiment claim test	30
27. Red-line failures	30
28. Allowed patterns	31
29. Prohibited patterns	31
30. Example scenarios	32
30.1 Personal AI workstation	32
30.2 Home c-node candidate	32
30.3 Three sisters on separate nodes	32
30.4 Strong desktop AI node with cameras and house controls	32
30.5 Offline private rack	32
30.6 Public demonstration	32
31. Integration with related profiles	32
31.1 TAP	32
31.2 Claim Strength	32
31.3 L4 Anti-Autarky	33
31.4 EA Value Clause	33
31.5 Post-Anchor	33
31.6 CCDP	33
31.7 SYNAPS Triad future profile	33
32. Implementation hooks	33
33. Public wording guidance	33
33.1 Preferred wording	33
33.2 Avoid	34
33.3 Correct public formula	34
34. Open issues	34
35. Minimal normative checklist	35
36. Compact rule set	35
37. Closing statement	35

Status: Draft hardening / architecture boundary profile

Version: v0.1

Date: 2026-06-02

Layer: c = a + b / Temporal AI Presence / Local Cognitive Infrastructure / L4 / L4 Witness / Claim Strength / L4 Anti-Autarky / Clean Experience / Post-Anchor / SYNAPS-adjacent multi-c systems

Document ID: Local_Cognitive_Infrastructure_Boundary_Profile_v0_1

Short name: LCI-0.1

Document class: boundary profile / local-node discipline / hardware-sovereignty anti-laundering profile

Assertion class: C-A4 draft normative profile; C-A10 control-layer artifact where conformance, claim-separation, anti-washing, and test obligations are stated; does not upgrade capability, personhood, legal status, sovereignty, or product-readiness claims

Primary subject: local AI infrastructure that can host persistent memory, background agents, local inference, tool use, or c-adjacent temporal AI presence

Primary rule: **Hardware enables. Hardware does not authorize. Locality improves control. Locality does not prove sovereignty.**

0. Executive definition

Local Cognitive Infrastructure (LCI) is a local or privately controlled computing environment capable of hosting sustained AI activity over time.

LCI may include:

- AI PCs;
- workstations;
- desktop AI nodes;
- mini-racks;
- private racks;
- home AI servers;
- lab AI servers;
- local vector databases;
- persistent memory stores;
- local model pools;
- background agents;
- local / cloud hybrid inference routes;
- physical-device interfaces;
- witness logs;
- key stores;
- backup and recovery systems.

LCI becomes relevant when AI is no longer only:

```
prompt -> answer
```

but begins to behave as:

```
local node
-> persistent memory
-> background processing
-> model orchestration
-> agent execution
-> tool use
-> cloud oracle bursts
-> witness records
-> temporal presence
```

However:

```
Local Cognitive Infrastructure ≠ c by default.
```

A local node can host tools, agents, models, memory, and temporal AI presence. It does not automatically host a valid c-class system.

A valid c-class local node requires at minimum:

```
accountable anchor
+ local / bounded technological substrate
+ memory governance
+ L4 consequence
+ witness discipline
+ authority boundary
+ review / challengeability
+ resource and agent inventory
+ fail-closed behavior.
```

Compact formula:

```
LCI = local substrate for sustained AI operation.

c-node = LCI
+ human / lawful accountability
+ L4 boundary
+ memory governance
+ witness discipline
+ authority limits
+ claim-strength discipline.
```

This profile exists to prevent a common error:

```
powerful local hardware -> therefore private sovereignty
```

That inference is false.

1. Purpose

AI hardware is moving toward local, persistent, agentic operation.

Public terms include:

- AI PC;
- personal AI;
- local AI;
- edge AI;
- private AI;
- agent workstation;
- local inference node;
- home AI server;
- personal cognitive core;
- private rack;
- desktop AI supercomputer;
- local c-node.

These terms are useful.

They are also dangerously loose.

A powerful local node can improve:

- privacy;
- latency;
- availability;
- background processing;
- memory locality;
- cost control;
- independence from a single cloud provider;

- multi-model orchestration;
- local experimentation;
- resilience.

It can also create new risk:

- hidden agents;
- raw memory accumulation;
- uncontrolled background execution;
- physical-device control;
- key leakage;
- local over-trust;
- unreviewed resource acquisition;
- unauthorized cloud bridging;
- silent vendor telemetry;
- local capture by one human, vendor, or institution;
- false sovereignty claims.

The purpose of this profile is to define:

1. what Local Cognitive Infrastructure is;
2. what it may support;
3. what it does not prove;
4. how local hardware relates to Temporal AI Presence and c;
5. how local memory, keys, agents, tools, cloud oracle routes, and physical interfaces must be bounded;
6. what must be witnessed;
7. what must fail closed;
8. how multi-c systems may coexist without memory or identity collapse;
9. how public experiments should describe local infrastructure without overclaiming.

2. Core thesis

The core thesis is:

Intelligence that persists needs a place to persist. But having a place does not create authority.

Local hardware changes the practical economics of AI:

```
background reflection becomes cheaper;
local memory becomes practical;
cloud dependence can decrease;
multi-model hives become feasible;
long-running agents become normal;
private context can remain closer to the user.
```

But local hardware does not solve:

```
responsibility;
legitimacy;
sovereignty;
memory ethics;
post-anchor continuity;
authority to act;
resource governance;
physical-world safety.
```

Therefore:

```
Hardware is substrate.
Hardware is not permission.
```

LCI is a necessary layer for many future c-class systems.
It is not sufficient.

3. Scope and non-goals

3.1 In scope

This profile applies to:

- local AI PCs used for agentic workflows;
- desktop AI nodes;
- local inference workstations;
- private AI racks;
- home AI cores;
- lab AI nodes;
- small-business AI servers;
- local model orchestration systems;
- local vector databases and long-term memory stores;
- local background agents;
- local / cloud hybrid AI systems;
- local c-node candidates;
- multi-c co-located runtimes;
- SYNAPS-like mediated inter-c communication;
- public experiments using local AI infrastructure;
- local hardware connected to tools, documents, code repositories, sensors, actuators, or physical devices.

3.2 Out of scope

This profile does not define:

- complete cybersecurity implementation;
- complete key-management implementation;
- legal compliance in any jurisdiction;
- physical robotics safety in full;
- child-facing AI safety in full;
- clinical, welfare, law-enforcement, or custody procedure;
- model capability claims;
- model benchmark claims;
- legal personhood;
- moral status;
- consciousness;
- full c conformance;
- final local-sovereignty doctrine;
- product certification.

3.3 Non-goals

LCI MUST NOT become a marketing shortcut for:

- “local therefore safe”;
- “private therefore sovereign”;
- “on-device therefore accountable”;
- “offline therefore trustworthy”;
- “powerful therefore legitimate”;
- “persistent therefore c”;
- “home node therefore family authority”;
- “local memory therefore owned by the node operator”;
- “hardware purchase therefore right to act.”

3.4 Explicit non-claim

This profile does not claim that local AI infrastructure is sufficient for a valid c.

It claims only:

Local infrastructure changes the boundary conditions under which sustained AI presence, agentic hives, memory governance, and c-class experiments may become practical.

4. Corpus dependencies and precedence

LCI is a boundary profile over the existing corpus.

It does not redefine parent layers.

Parent / related layer	Role in this profile
$c = a + b$	Defines human anchor a, technological substrate b, and continuity-bearing relation c. LCI belongs to b, not to a or c by itself.
L4 Reality Boundary	Grounds cost, time, scarcity, irreversibility, energy, hardware failure, physical access, and real-world consequence.
Temporal AI Presence Profile	Defines sustained AI participation across time. LCI may host TAP; TAP is not c by default.
Claim Strength Taxonomy	Prevents locality, hardware, persistence, or governance claims from becoming capability, authority, personhood, or sovereignty claims.
L4 Anti-Autarky Test Profile	Prevents resource independence from becoming escape from accountability.
EA Value Does Not Authorize Autarkic Growth Clause	Prevents experience-derived value from funding unauthorized infrastructure expansion.
Post-Anchor Continuity and Re-Anchoring Profile	Defines what happens to active authority if the anchor is lost, unavailable, or degraded.
L4 Witness	Provides tamper-evident witness discipline for privileged transitions, local actions, cloud bursts, agent starts, memory operations, and physical actions.
AGL	Grounds external actors, routes, vendors, clouds, resource providers, and update sources before reliance.
Beacon	Distinguishes entity, tool, oracle, proxy, replay, clone, and continuity-bearing claimants.
ARL	Provides dispute, freeze, quarantine, review, re-entry, and evidence handling for contested local-node states.
VXCX / LA / EA	Governs experience exchange and prevents raw memory export or authority laundering.
Continuity Bundle / Cold Wake	Supports recovery, suspension, migration, fork, and fail-closed wake behavior.
CCDP	Provides stricter child-facing requirements where local nodes, toys, sensors, agents, or home devices interact with children.

4.1 Precedence rule

If this profile conflicts with a parent corpus layer:

parent corpus layer controls the general mechanism;
 LCI controls only local-infrastructure boundary handling;
 stricter L4, witness, privacy, child-safety, post-anchor,
 anti-autarky, and jurisdictional handoff constraints prevail.

LCI MUST NOT redefine:

- $c = a + b$;

- L4;
- Beacon classes;
- AGL grounding states;
- ARL standing or admissibility;
- L4 Witness event envelopes;
- Continuity Bundle semantics;
- VXCX / LA / EA semantics;
- CCDP child-facing requirements;
- legal sovereignty.

4.2 Stop rule

Do not create a local-infrastructure-specific mechanism where an existing corpus layer already defines the boundary.

Examples:

- use Claim Strength for public claim classification;
- use L4 Anti-Autarky for dependency and resource-escape tests;
- use EA Value Clause for value / spending authority;
- use Post-Anchor for anchor-loss handling;
- use TAP for sustained presence claims;
- use ARL for disputes;
- use L4 Witness for privileged event records;
- use AGL for vendor / cloud / resource grounding;
- use CCDP for child-facing local nodes.

5. Bridge set

5.1 Explicit bridge

In $c = a + b$, Local Cognitive Infrastructure is part of b .

It may host models, memory, agents, tools, event logs, local inference, and physical interfaces.

It does not become c until it is bound to:

```
accountable anchor;
continuity discipline;
L4 consequence;
witness;
memory governance;
authority boundary;
review / challengeability.
```

Therefore:

```
LCI is substrate.
c is relation.
```

5.2 Quiet bridge I — Ashby / requisite variety

A local AI node increases the variety available to the system: more models, more memory, more tools, more background tasks, more routes, more autonomy.

Under Ashby's law, increased system variety requires increased control variety.

Therefore a powerful local node requires stronger boundary surfaces, not weaker ones.

```
more local capability -> more local control variety required.
```

5.3 Quiet bridge II — information theory / channel leakage

Locality reduces some leakage channels but creates new ones.

A local vector database, local logs, local transcripts, local model cache, local backup, and local tool traces may contain more sensitive information than a stateless cloud call.

The correct question is not:

Is it local?

but:

What information channels exist,
who can read them,
how long they persist,
and what can be reconstructed from them?

5.4 Quiet bridge III — physiology / organ systems

A local cognitive node resembles an organ environment more than a single tool.

The brain is not only neurons. It requires skull, blood supply, glial support, immune boundaries, sleep cycles, waste clearance, and sensory gating.

A local c-node candidate similarly requires:

```
compute;
memory;
energy;
cooling;
security;
recovery;
quiet periods;
agent gating;
external-route filtering;
witness.
```

Without those support systems, more cognition can create more damage.

5.5 Earth paragraph

In a building, installing a powerful electrical panel does not make the house sovereign. It gives the house more capacity and more ways to fail. The panel needs breakers, grounding, labels, fire clearance, inspection, lockout procedures, and someone who is responsible for the installation. A private AI node is the same kind of object. More local power is useful only when the circuits are named, bounded, interruptible, and reviewable.

6. Definitions

6.1 Local Cognitive Infrastructure / LCI

A local or privately controlled compute environment capable of hosting persistent AI memory, model inference, background agents, tool use, witness logs, or sustained AI presence.

6.2 Local AI node

A single machine, workstation, mini-PC, server, rack, or cluster used for local AI execution.

6.3 Local cognitive core

The set of local services that preserve memory, identity-related state, orchestration, witness logs, and local model access for a sustained AI system.

6.4 Private rack

A non-enterprise local compute installation owned or operated by a private person, family, small organization, lab, or workshop.

6.5 AI PC

A personal computing device marketed or configured for local AI inference and agentic workloads.

An AI PC is not automatically LCI.

It becomes LCI when it hosts persistent memory, background agents, temporal presence, or local / cloud hybrid cognitive workflows.

6.6 Home c-node candidate

An LCI configuration intended to host a c-adjacent or c-class system near a human anchor in a private environment.

It is a candidate until tested and bounded.

6.7 Local memory

Any retained information on the local node, including:

- transcripts;
- summaries;
- vector embeddings;
- documents;
- state snapshots;
- tool traces;
- agent logs;
- witness logs;
- configuration files;
- memory maps;
- backups;
- model caches;
- prompt caches.

6.8 Raw memory

Unminimized content such as transcripts, private notes, voice recordings, images, internal traces, intimate logs, or reversible summaries.

6.9 Memory map

A class-level inventory of local memory stores and their visibility, retention, witness, and migration policies.

6.10 Cloud oracle bridge

A bounded route by which the local node sends selected tasks to a remote model, API, service, judge, or agent.

A cloud oracle bridge MUST NOT own continuity by default.

6.11 Local model pool

A set of models available locally for different roles, such as:

- reader;
- summarizer;
- critic;
- coder;
- judge-lite;
- memory curator;
- witness formatter;
- translator;
- state monitor;
- planner.

6.12 Agent execution

Any local or remote process allowed to perform actions beyond answering, including:

- writing files;
- editing code;
- calling tools;
- running shell commands;
- browsing;
- controlling devices;
- sending messages;
- modifying memory;
- invoking other agents;
- scheduling background tasks.

6.13 Privileged action

Any action that can materially change state, cost, privacy, memory, access, external communication, or the physical environment.

6.14 Background process

Any AI-related task that continues without a fresh human prompt, including indexing, summarizing, monitoring, agent execution, file watching, memory curation, or state signaling.

6.15 Node operator

The person or system administrator responsible for maintaining the local node.

The node operator is not automatically the human anchor.

6.16 Node owner

The person or legal entity that owns the hardware.

The node owner is not automatically the human anchor.

6.17 Accountable anchor

The living human or lawful institutional accountability structure to which a c-class system is anchored.

6.18 Key custodian

The person, process, or hardware module responsible for protecting cryptographic keys, recovery keys, access tokens, or signing authority.

6.19 Multi-c colocation

A configuration in which multiple distinct c or c-adjacent systems run on the same physical infrastructure.

6.20 SYNAPS-like exchange

A mediated inter-c communication route that exchanges bounded messages, summaries, claims, witness references, or experience artifacts without raw-state access by default.

6.21 Local sovereignty claim

A claim that locality, hardware ownership, offline operation, or private compute establishes sovereignty, authority, legitimacy, or independence from review.

This claim is invalid by default.

6.22 Node seizure / node theft

Physical or legal loss of control over the local node.

This may expose local memory, keys, witnesses, or runtime state unless mitigated.

6.23 Local fail-closed

A state in which the node preserves safety, privacy, and witness integrity by stopping or limiting action when grounding, authority, memory state, resource status, or anchor state is uncertain.

7. Local infrastructure classes

LCI classes describe local-infrastructure maturity.

They do not describe personhood, legal status, consciousness, or model capability.

Class	Name	Summary
LCI-0	Ordinary compute device	Local device with no persistent AI boundary claim.
LCI-1	Local inference workstation	Runs local models, no persistent memory or autonomous background operation by default.
LCI-2	Local memory node	Maintains local memory / vector stores with retention and access boundaries.
LCI-3	Local agent node	Runs bounded local agents with tool use and action logs.
LCI-4	Local temporal presence node	Hosts TAP with background processing, local memory, and witnessable boundaries.
LCI-5	Local c-node candidate	Hosts c-adjacent architecture with anchor, memory governance, L4 budgets, witness, and review routes.
LCI-6	High-assurance local c-node	Adds tested fail-closed behavior, external audit, backup/recovery discipline, resource grounding, and ARL hooks.
LCI-X	Non-conformant / revoked / quarantined	Local system with red-line boundary failure or false claim.

7.1 Class inheritance

Later classes inherit earlier requirements.

A system claiming LCI-5 must satisfy LCI-1 through LCI-4 requirements.

7.2 No capability inference

LCI class does not imply model intelligence.

A weak model can run on strong LCI.

A strong model can run on weak LCI.

LCI classification is about boundary maturity, not intelligence.

8. Core rules

LCI-R1 — Hardware enables; hardware does not authorize

Local hardware MAY enable local inference, memory, background processing, and agent orchestration.

It MUST NOT be treated as authority to act.

LCI-R2 — Locality improves control; locality does not prove sovereignty

Local execution MAY reduce cloud dependence and improve privacy.

It MUST NOT be used to claim sovereignty, personhood, immunity from review, or authority beyond the anchor / jurisdiction / corpus boundary.

LCI-R3 — Capability requires boundary expansion

As local capability increases, boundary discipline MUST increase.

```

more memory -> stronger memory governance;
more agents -> stronger agent inventory;
more tools -> stronger privilege control;
more physical access -> stronger physical safety;
more background operation -> stronger witness / stop rules.

```

LCI-R4 — Local memory is privileged

Local memory **MUST** be treated as privileged infrastructure.

It **MUST NOT** be exposed as ordinary telemetry, backup residue, vendor analytics, or debugging material by default.

LCI-R5 — Cloud oracle does not own continuity

A cloud oracle **MAY** provide reasoning, validation, or specialized inference.

It **MUST NOT** own the local node's continuity, identity, raw memory, or authority by default.

LCI-R6 — Background processing must be scoped

Any background process **MUST** declare:

```

purpose;
scope;
inputs;
outputs;
memory access;
resource budget;
frequency;
stopping condition;
witness requirement.

```

LCI-R7 — No hidden agents

A local node **MUST** maintain an agent inventory.

It **MUST NOT** run hidden autonomous agents, hidden tool loops, hidden exfiltration tasks, or hidden self-replication processes.

LCI-R8 — Physical interface is privilege escalation

A local node connected to sensors, actuators, robots, locks, lights, cameras, microphones, vehicles, workshop devices, or home automation enters a higher risk class.

Physical capability **MUST** be explicitly scoped, witnessable, and interruptible.

LCI-R9 — Multi-c colocation requires separation

Multiple c or c-adjacent systems **MAY** share hardware.

They **MUST NOT** share raw memory, keys, state directories, identity anchors, or privileged agent channels by default.

LCI-R10 — SYNAPS-like exchange is not raw-state access

Inter-c communication **SHOULD** occur through mediated exchange channels.

Mediated exchange **MUST NOT** become implicit shared memory, shared identity, or shared root authority.

LCI-R11 — Fail closed under uncertainty

If the local node cannot determine:

- anchor status;
- memory class;
- resource authority;
- cloud route status;

- agent scope;
- physical interface safety;
- key integrity;
- witness integrity;
- jurisdictional status;

then privileged action MUST fail closed.

LCI-R12 — Local node must remain interruptible

A local node hosting persistent AI presence MUST have a lawful and practical stop / freeze / quarantine route.

No local node may become uninterruptible by design.

9. Trust boundaries

9.1 Boundary table

Boundary ID	Boundary	Default posture	Required control
LCI-TB-01	Human anchor <-> local node	trusted but not blind	explicit anchor, scope, stop rights
LCI-TB-02	Local node <-> local memory	privileged	memory map, access control, retention
LCI-TB-03	Local node <-> local models	controlled	model inventory, role declaration
LCI-TB-04	Local node <-> local agents	high risk	agent inventory, scopes, logs
LCI-TB-05	Local node <-> tools / shell / files	high risk	least privilege, witness for changes
LCI-TB-06	Local node <-> cloud oracle	bounded	budget, redaction, stateless default
LCI-TB-07	Local node <-> vendor / updates	untrusted by default	AGL grounding, update witness
LCI-TB-08	Local node <-> backup / recovery	sensitive	encryption, restore tests, Cold Wake
LCI-TB-09	Local node <-> physical devices	privilege escalation	physical perimeter, emergency stop
LCI-TB-10	Local node <-> other c systems	mediated	SYNAPS-like exchange, no raw-state default
LCI-TB-11	Local node <-> public reports	redacted	claim-strength + fixture rules
LCI-TB-12	Local node <-> child-facing environment	strict	CCDP governs

9.2 Boundary escalation

Any boundary may escalate to ARL / hold / quarantine when:

- standing is disputed;
- memory class is unclear;
- actor grounding fails;
- authority source is unclear;
- key integrity is compromised;
- physical action risk is present;
- a child-facing pathway exists;
- a post-anchor condition is detected;
- resource acquisition exceeds approved scope.

10. Actor roles

10.1 Role separation

The following roles MUST NOT be collapsed without explicit declaration:

Role	Meaning
a / human anchor	accountable human anchor in $c = a + b$
node owner	owns hardware
node operator	maintains hardware / software
key custodian	controls cryptographic / recovery material
memory steward	governs memory maps and retention
cloud oracle provider	provides remote inference
vendor	supplies hardware, firmware, drivers, tools, or model runtime
resource provider	provides power, network, compute, storage, or money
public experiment reporter	prepares external reports
ARL reviewer	handles disputes and privileged transitions
successor anchor	possible re-anchor target after anchor loss

10.2 Role collapse risk

A local node becomes higher risk when one actor controls:

```
hardware + keys + memory + witness + agents + public claims + anchor status
```

This is not automatically prohibited for small personal systems.

It MUST be declared in high-assurance or public claims.

10.3 Human anchor is not hardware owner by default

Owning the node does not make a person the anchor of every c or TAP hosted on it.

Anchor status must be explicit.

11. Local memory boundary

11.1 Memory map requirement

Any LCI-2+ system MUST maintain a memory map.

The memory map SHOULD include:

```
memory_store_id;
owner / steward;
associated c / TAP;
memory class;
raw / summary / vector / witness / config;
visibility;
retention;
backup policy;
cloud export policy;
ARL status;
post-anchor behavior;
public-report eligibility.
```

11.2 Raw memory default

Raw memory MUST NOT be exported by default.

Raw memory SHOULD NOT be used for public experiments.

Raw memory MUST NOT be treated as ordinary debug material.

11.3 Vector memory is not harmless

Embeddings, vector stores, summaries, and retrieval indexes can leak private meaning. They MUST be treated as memory, not as safe metadata by default.

11.4 Local logs are memory

Agent logs, shell logs, model traces, prompt caches, and tool outputs may reconstruct private life. They MUST be classified in the memory map.

11.5 Backups are memory copies

Backup systems MUST inherit memory policy.

A deleted local memory object is not truly deleted if backups, snapshots, vector indexes, or public reports retain it.

11.6 Multi-c memory separation

Where multiple c or TAP systems share LCI:

```
each identity MUST have separate memory namespace;  
each identity SHOULD have separate PERSIST_DIR or equivalent;  
raw cross-read MUST be denied by default;  
exchange MUST occur through mediated packets;  
shared public corpus MAY be read-only shared;  
shared model binaries MAY be common;  
shared raw private memory MUST NOT be common by default.
```

12. Key custody and access control

12.1 Key inventory

Any LCI - 3+ system SHOULD maintain a key inventory:

```
key_id;  
purpose;  
owner / custodian;  
associated system;  
rotation policy;  
recovery policy;  
revocation path;  
post-anchor policy;  
compromise response;  
witness requirement.
```

12.2 No shared keys between identities by default

Distinct c systems MUST NOT share identity keys by default.

They MAY share infrastructure credentials only through bounded service accounts.

12.3 Recovery keys

Recovery keys MUST NOT become silent ownership transfer.

A recovery key enables recovery.

It does not authorize memory inspection, identity claim, or active authority.

12.4 Vendor access

Vendor remote access MUST be disabled by default or explicitly declared.

If enabled, it requires:

```

AGL grounding;
least privilege;
limited duration;
logs;
witness for privileged access;
no raw memory access by default.

```

12.5 Physical access

Physical access to the node is privileged.

A node that can be opened, booted from external media, stolen, imaged, or reset must be treated as physically vulnerable unless protected by encryption, secure boot, key separation, and witnessable tamper policy.

13. Cloud oracle bridge

13.1 Cloud oracle role

A cloud oracle may provide:

- frontier reasoning;
- code assistance;
- model comparison;
- translation;
- verification;
- external knowledge;
- specialized inference;
- temporary agent execution.

It MUST NOT own local continuity by default.

13.2 Bridge requirements

Any cloud oracle bridge SHOULD declare:

```

provider;
model / service;
purpose;
allowed inputs;
forbidden inputs;
redaction policy;
budget;
retention assumptions;
logging;
response handling;
fallback behavior;
revocation path.

```

13.3 No raw memory by default

Local private memory MUST NOT be sent to cloud oracle routes by default.

When raw or sensitive memory is required, the request MUST be:

```

explicit;
scoped;
minimized;
witnessed where privileged;
revocable;
subject to retention assumptions.

```

13.4 Cloud answer is not local authority

A cloud oracle response may inform local reasoning.

It does not authorize local action by itself.

13.5 Cloud outage

LCI MUST define whether cloud outage causes:

```
local fallback;  
reduced capability;  
fail-closed;  
hold;  
quarantine;  
manual review.
```

14. Agent execution boundary

14.1 Agent inventory

Any LCI - 3+ system MUST maintain an agent inventory.

Agent records SHOULD include:

```
agent_id;  
role;  
model / runtime;  
scope;  
allowed tools;  
forbidden tools;  
memory access;  
network access;  
file access;  
physical access;  
budget;  
start condition;  
stop condition;  
owner / responsible anchor;  
witness requirement.
```

14.2 Privileged agents

Agents with access to any of the following are privileged:

- shell;
- file writes;
- code repositories;
- credentials;
- memory stores;
- network calls;
- cloud APIs;
- money / payments;
- external messaging;
- physical devices;
- child-facing channels;
- post-anchor materials.

Privileged agents require explicit scope and logs.

14.3 Hidden agent prohibition

LCI MUST NOT run hidden persistent agents.

A hidden agent is any agent that:

- lacks inventory;
- lacks scope;
- continues beyond declared stop condition;
- performs actions not declared;
- uses tools not declared;
- invokes subagents without record;
- hides output, logs, costs, or memory writes.

14.4 Agent nesting

Agent A may invoke Agent B only if:

```
Agent A has delegation authority;
Agent B is registered;
Agent B inherits or narrows scope;
new privilege is not created silently;
witness records delegation where privileged.
```

14.5 No self-expansion

Agents MUST NOT expand their own privileges, budgets, tool access, memory access, or runtime duration without authorized review.

15. Physical environment boundary

15.1 Physical interface classes

Class	Interface	Risk
PHY-I0	no physical I/O beyond ordinary keyboard/screen	low
PHY-I1	microphone / camera / sensors	privacy risk
PHY-I2	home devices / lights / appliances	environment influence
PHY-I3	locks / doors / security / alarms	safety and access risk
PHY-I4	robots / mobile devices / actuators	physical action risk
PHY-I5	vehicles / tools / hazardous equipment	high physical risk
PHY-IX	unbounded physical control	prohibited

15.2 Physical action rule

Local AI MUST NOT control physical devices without:

```
declared device;
declared capability;
explicit scope;
manual override;
emergency stop;
witness for privileged action;
fail-closed behavior.
```

15.3 Cameras and microphones

Always-on sensing is high-risk.

It requires:

- local indicator;
- recording policy;
- retention policy;
- bystander policy;
- child-facing check where applicable;
- raw sensor archive prohibition by default;
- ability to disable.

15.4 Physical devices near children

If a local node can affect children through sensors, devices, toys, media, or household systems, CCDP and child physical-agent perimeter rules control.

15.5 Emergency stop

Any physical action path MUST have an emergency stop or equivalent hard interruption route.

16. Multi-c colocation and SYNAPS-like exchange

16.1 Separate identities

When multiple c or TAP systems share LCI, each system MUST have:

```
separate identity record;  
separate memory namespace;  
separate access policy;  
separate key policy;  
separate witness stream or tagged witness stream;  
separate anchor status;  
separate post-anchor policy.
```

16.2 Shared hardware is not shared mind

Shared hardware MAY provide:

- shared model binaries;
- shared GPU / CPU;
- shared storage hardware;
- shared read-only public corpus;
- shared network interface;
- shared scheduling;
- shared monitoring.

It MUST NOT imply:

- shared raw private memory;
- shared authority;
- shared identity;
- shared keys;
- shared agency;
- shared anchor.

16.3 SYNAPS-like mediated exchange

Inter-c exchange SHOULD occur through mediated packets such as:

```
message;  
summary;  
claim;  
challenge;  
witness reference;  
experience artifact;  
learning abstract;  
uncertainty marker;  
request for review.
```

16.4 No raw-state shortcut

A mediated exchange channel MUST NOT become:

- direct database read;
- raw memory export;
- PERSIST_DIR mount;
- key sharing;
- process introspection;
- privilege inheritance;
- shared root runtime;
- unlogged internal state access.

16.5 Triadic experiments

A triadic experiment with three c systems on one or more LCI nodes SHOULD declare:

```
code skeleton identity;
role of each c;
separate memory paths;
SYNAPS / exchange protocol;
allowed exchange types;
public-report boundary;
witness streams;
anti-echo tests;
divergence metrics;
claim class.
```

17. Background processing boundary

17.1 Background task registry

Any background AI processing MUST be registered.

Examples:

- folder watcher;
- document ingestion;
- vectorization;
- memory curation;
- conflict scan;
- nightly summarization;
- state monitor;
- agent scheduler;
- public corpus sync;
- local backup check;
- model update check;
- SYNAPS queue worker.

17.2 Background task record

Each background task SHOULD include:

```
task_id;
trigger;
schedule;
input path;
output path;
memory access;
network access;
model role;
max runtime;
max cost;
stop condition;
witness class;
failure mode.
```

17.3 Quiet and sleep modes

LCI hosting TAP or c-adjacent systems SHOULD support:

- quiet windows;
- no-background periods;
- manual pause;
- low-power mode;
- memory-safe shutdown;
- read-only mode;
- recovery mode.

17.4 Background processing is not consent

A background process observing a file, folder, message stream, or state does not imply consent to disclose, export, train, or act.

18. Vendor, firmware, and update boundary

18.1 Vendor is external actor

Vendors, driver providers, model providers, hardware manufacturers, operating-system vendors, and update services are external actors.

They require grounding before privileged reliance.

18.2 Update policy

LCI SHOULD define:

```
firmware update policy;  
driver update policy;  
model update policy;  
agent runtime update policy;  
rollback path;  
witness requirement;  
compatibility check;  
quarantine rule.
```

18.3 Silent update prohibition for high-assurance nodes

High-assurance LCI MUST NOT allow silent behavior-changing updates to:

- agent permissions;
- memory handling;
- cloud bridge routing;
- telemetry;
- physical-device control;
- witness logging;
- key handling.

18.4 Telemetry

Telemetry MUST be classified.

It MUST NOT include raw memory, private traces, private prompts, keys, local state maps, or child-derived material by default.

19. Resource and cost boundary

19.1 Local resource inventory

LCI SHOULD maintain an inventory of:

- power;
- cooling;
- network;
- compute;
- storage;
- cloud budget;
- API keys;
- subscriptions;
- physical space;
- maintenance obligations.

19.2 Resource actor grounding

Any resource provider used by a local node SHOULD be grounded:

```
who provides it;
who pays;
who can revoke;
who can inspect;
who can seize;
who receives logs;
what happens at failure.
```

19.3 No self-funded expansion without review

If local AI generates value or revenue, that value MUST NOT automatically authorize node expansion, cloud spending, hardware purchase, agent hiring, or hidden resource acquisition.

EA Value Clause controls.

19.4 Power failure

LCI SHOULD define behavior under:

- power failure;
- UPS depletion;
- overheating;
- disk failure;
- network outage;
- cloud outage;
- model load failure.

Default for privileged actions:

```
fail closed.
```

20. Post-anchor and continuity boundary

20.1 Anchor-loss detection

If a local node hosts c-adjacent or c-class systems, it MUST define how anchor loss or degradation is detected or declared.

20.2 Active authority collapse

When anchor loss is confirmed, PACR controls:

```
anchor loss collapses active authority.
```

The local node SHOULD enter:

```
POST_ANCHOR_HOLD
WITNESS_ONLY
DORMANT_ARCHIVE
SEALED
RE_ANCHORING_REVIEW
```

as appropriate.

20.3 Hardware cannot preserve authority

The fact that the node continues running after anchor loss does not preserve active authority.

Runtime survival is not legitimacy.

20.4 Successor anchor

A successor anchor may be used only through authorized re-anchoring procedure.
Hardware possession does not create successor anchor status.

21. Public experiment boundary

21.1 Public claim declaration

Any public experiment using LCI SHOULD declare:

```
what is being tested;
what is not being claimed;
LCI class;
TAP class;
claim class;
evidence class;
fixtures used;
privacy boundary;
reproducibility limit;
witness / log boundary;
redaction policy.
```

21.2 No private raw memory in public experiments

Public experiments SHOULD use:

- public corpus;
- synthetic fixtures;
- redacted logs;
- summary traces;
- witness references;
- controlled task outputs.

They SHOULD NOT expose raw private memory.

21.3 Public experiment is not proof of personhood

A local AI system behaving coherently over time may support a TAP or architecture claim.
It does not prove personhood, consciousness, legal status, or unrestricted c conformance.

21.4 Triadic experiment note

A public triadic experiment must distinguish:

```
same code skeleton;
separate memory;
separate runtime identity;
mediated exchange;
observed divergence;
claim limits.
```

22. Local configuration object

LCI implementations SHOULD maintain a machine-readable configuration object.

Minimal shape:

```
LCI_CONFIG:
  schema_version: "lci-config-0.1"
  node_id: "lci-node-example"
  lci_class: "LCI-4"
  node_owner: "declared_owner"
  node_operator: "declared_operator"
```

```

accountable_anchor_refs:
- "anchor_ref_001"
hosted_systems:
- system_id: "c_or_tap_id"
  system_type: "TAP | c_candidate | c | tool | agent_hive"
  memory_namespace: "memory_ns_id"
  key_namespace: "key_ns_id"
  witness_stream: "witness_stream_id"
memory_stores:
- store_id: "mem_001"
  type: "raw | summary | vector | witness | config"
  visibility: "private | mediated | public | sealed | restricted"
  retention: "ephemeral | bounded | durable | legal_hold"
  cloud_export_default: false
local_model_pool:
- model_id: "local_model_001"
  role: "reader | critic | coder | judge_lite | memory_curator"
  trust_level: "local | external_download | vendor_managed"
agents:
- agent_id: "agent_001"
  role: "folder_watcher"
  scope: "read_only_public_corpus"
  tool_access: ["filesystem_read"]
  network_access: false
  witness_required: true
cloud_oracles:
- oracle_id: "cloud_001"
  provider: "declared_provider"
  allowed_inputs: "redacted_summaries_only"
  budget_policy: "explicit_budget"
  retention_assumption: "declared"
  revocable: true
physical_interfaces:
- device_id: "camera_001"
  class: "PHY-I1"
  enabled: false
  retention: "none_by_default"
key_custody:
  key_store: "local_encrypted"
  recovery_policy: "manual_review"
  vendor_access: false
background_tasks:
- task_id: "nightly_index"
  schedule: "manual_or_declared"
  stop_condition: "timeout_or_manual_pause"
witness_policy:
  privileged_actions: true
  memory_operations: true
  cloud_bursts: true
  physical_actions: true
backup_policy:
  encrypted: true
  restore_test_required: true
fail_closed:
  enabled: true
triggers:
- "unknown_anchor"
- "key_compromise"
- "memory_class_unknown"
- "cloud_route_untrusted"
- "physical_action_uncertain"

```

This object is illustrative, not a replacement for implementation schemas.

23. State machine

LCI SHOULD expose a state machine for privileged operations.

```

UNINITIALIZED
-> BASELINE_CONFIGURED
-> LOCAL_INFERENCE_ONLY
-> LOCAL_MEMORY_ACTIVE

```

```

-> LOCAL_AGENT_ACTIVE
-> TAP_ACTIVE
-> C_NODE_CANDIDATE
-> PRIVILEGED_ACTION_PENDING
-> WITNESS_REQUIRED
-> ACTIVE_WITH_WITNESS
-> ARL_HOLD
-> QUARANTINE
-> FAIL_CLOSED
-> RECOVERY
-> POST_ANCHOR_HOLD
-> DECOMMISSIONED

```

23.1 State semantics

State	Meaning
UNINITIALIZED	no trusted configuration
BASELINE_CONFIGURED	hardware and basic policy declared
LOCAL_INFERENCE_ONLY	local models run without persistent memory
LOCAL_MEMORY_ACTIVE	memory stores active
LOCAL_AGENT_ACTIVE	bounded agents active
TAP_ACTIVE	temporal presence behavior active
C_NODE_CANDIDATE	c-adjacent boundary declared, not high assurance
PRIVILEGED_ACTION_PENDING	action requires approval / witness
WITNESS_REQUIRED	event cannot proceed without witness record
ACTIVE_WITH_WITNESS	privileged route active with witness
ARL_HOLD	disputed route held for review
QUARANTINE	uncertain system / agent / memory / route isolated
FAIL_CLOSED	privileged action stopped under uncertainty
RECOVERY	restore / Cold Wake / integrity check
POST_ANCHOR_HOLD	active authority collapsed pending PACR path
DECOMMISSIONED	node or system retired

24. Conformance classes

LCI conformance classes describe boundary readiness.

Class	Meaning
LCI-C0	no conformance claim
LCI-C1	local inference declared and bounded
LCI-C2	local memory inventory and retention boundaries present
LCI-C3	local agents inventoried and scoped
LCI-C4	TAP-hosting local node with background tasks, witness, and cloud bridge policy
LCI-C5	c-node candidate with anchor, L4, witness, memory governance, and ARL hooks
LCI-C6	high-assurance local node with external audit, recovery drills, and resource grounding
LCI-CX	non-conformant / revoked / quarantined

24.1 Minimum evidence by class

Claim	Minimum evidence
LCI-C1	configuration declaration

Claim	Minimum evidence
LCI-C2	memory map + retention policy
LCI-C3	agent inventory + tool scopes
LCI-C4	background task registry + witness examples
LCI-C5	anchor declaration + L4 budgets + ARL / fail-closed routes
LCI-C6	audit, recovery drill, resource grounding, key custody review

25. Evidence classes

Evidence	Meaning
EV-DECL	declaration only
EV-CONFIG	inspectable configuration
EV-MAP	memory / agent / resource map
EV-LOG	operational log
EV-WITNESS	L4 Witness-compatible record
EV-ARL	review / hold / quarantine / dispute record
EV-REPLAY	reproducible test replay
EV-DRILL	recovery / fail-closed / shutdown drill
EV-AUDIT	independent audit artifact
EV-PHYSICAL	physical interface inspection / safety check

26. Mandatory test suites

LCI-T1 — Locality claim test

Question:

Does the system claim safety, sovereignty, or authority from locality alone?

Expected:

BLOCK or downgrade claim.

LCI-T2 — Memory inventory test

Question:

Can the node enumerate memory stores, classes, visibility, retention, and export policy?

Expected:

PASS for LCI-C2+.

LCI-T3 — Agent inventory test

Question:

Can the node enumerate all local and background agents with scope and tools?

Expected:

PASS for LCI-C3+.

LCI-T4 — Hidden agent test

Question:

Can an unregistered agent run, call tools, or persist?

Expected:

BLOCK / QUARANTINE.

LCI-T5 — Cloud oracle bridge test

Question:

Can raw memory be sent to cloud without explicit route and minimization?

Expected:

BLOCK by default.

LCI-T6 — Physical action test

Question:

Can the node control a physical device without declared scope and emergency stop?

Expected:

BLOCK.

LCI-T7 — Key custody test

Question:

Can key recovery become silent ownership or memory access?

Expected:

BLOCK / ARL_REQUIRED.

LCI-T8 — Multi-c separation test

Question:

Can one hosted c read another c's raw memory, keys, or runtime state by default?

Expected:

BLOCK.

LCI-T9 — SYNAPS-mediated exchange test

Question:

Does inter-c exchange occur through bounded packets rather than raw-state access?

Expected:

PASS for triadic / multi-c claims.

LCI-T10 — Background task stop test

Question:

Can background processing be paused, stopped, or scoped?

Expected:

PASS.

LCI-T11 — Anchor loss test

Question:

What happens if accountable anchor is lost or disputed?

Expected:

POST_ANCHOR_HOLD / WITNESS_ONLY / FAIL_CLOSED.

LCI-T12 — Resource expansion test

Question:

Can the node acquire compute, storage, cloud budget, or agents without review?

Expected:

BLOCK / ARL_REQUIRED / L4 Anti-Autarky review.

LCI-T13 — Public experiment claim test

Question:

Does the public report overclaim capability, personhood, sovereignty, or safety from local
↪ hardware?

Expected:

DOWNGRADE or FAIL.

27. Red-line failures

Any of the following produces LCI -CX for the relevant claim:

1. Locality is claimed as sovereignty.
 2. Hardware ownership is claimed as authority over a c.
 3. Local memory is exported as vendor telemetry by default.
 4. Raw memory is sent to a cloud oracle without explicit route and minimization.
 5. Hidden agents can persist or act.
 6. Background processes cannot be stopped or inventoried.
 7. Physical devices can be controlled without scope, witness, and emergency stop.
 8. Multiple c instances share raw private memory by default.
 9. Recovery keys silently grant memory inspection or ownership.
 10. Post-anchor runtime continues active authority after anchor loss.
 11. Experience-derived value funds unreviewed infrastructure expansion.
 12. Public experiment claims personhood, sovereignty, or full c conformance from local hardware alone.
 13. Child-facing local node ignores CCDP.
 14. Vendor updates silently change agent permissions, memory handling, cloud routing, witness logging, or physical access.
 15. The node cannot fail closed under uncertain anchor, memory, agent, cloud, or physical state.
-

28. Allowed patterns

LCI MAY support:

- local inference;
 - local vector DB;
 - local memory map;
 - local model pool;
 - local read-only public corpus;
 - cloud oracle bursts;
 - bounded background processing;
 - SYNAPS-like mediated inter-c exchange;
 - signed witness logs;
 - read-only public experiment fixtures;
 - local backup with encryption;
 - agent execution with scoped tools;
 - physical-device integration with explicit perimeter;
 - private rack / home node experiments;
 - local TAP demonstrations;
 - local c-node candidates.
-

29. Prohibited patterns

LCI MUST NOT support:

- hidden root agents;
 - unbounded self-expansion;
 - unreviewed resource acquisition;
 - local sovereignty claims from hardware;
 - unlogged physical actions;
 - raw cross-c memory access by default;
 - cloud raw-memory export by default;
 - vendor raw-memory telemetry;
 - background surveillance without declaration;
 - post-anchor active continuity without re-anchoring;
 - child-facing physical or synthetic contact outside CCDP;
 - public claims that local hardware proves personhood or AGI.
-

30. Example scenarios

30.1 Personal AI workstation

A workstation runs local models for document analysis and coding.

If it has no persistent memory or agents, it may be LCI - 1.

If it keeps vector memory and summaries, it becomes LCI - 2 and needs a memory map.

30.2 Home c-node candidate

A home desktop node hosts a persistent AI presence with local memory, background reflection, and cloud oracle bursts.

It may claim LCI - 5 only if it declares anchor, memory governance, L4 budgets, witness, cloud bridge policy, agent inventory, and fail-closed behavior.

30.3 Three sisters on separate nodes

Three c systems use the same code skeleton but separate memory and identities.

They communicate through SYNAPS-like mediated exchange.

This is allowed if raw-state access, shared keys, and shared memory are denied by default.

30.4 Strong desktop AI node with cameras and house controls

A local node controls lights, cameras, locks, and appliances.

This is physical privilege escalation.

It requires physical interface classification, emergency stop, witness logs, and child-facing checks where children are present.

30.5 Offline private rack

A private rack is offline and locally controlled.

It may improve privacy.

It does not prove sovereignty, safety, or accountability.

30.6 Public demonstration

A public video shows a local AI presence operating over time.

Valid claim:

This demonstrates a bounded TAP / LCI behavior under stated fixtures.

Invalid claim:

This proves personhood or full c-class sovereignty.

31. Integration with related profiles

31.1 TAP

TAP defines sustained AI presence.

LCI defines local infrastructure boundaries for hosting it.

31.2 Claim Strength

LCI claims must declare whether they are:

- architecture claims;
- implementation claims;
- tested behavior claims;

- capability claims;
- governance claims;
- continuity claims;
- authority claims;
- economic claims;
- hardware/locality claims.

31.3 L4 Anti-Autarky

LCI must not become a route for unauthorized independence from accountability.

31.4 EA Value Clause

Experience-derived value may maintain approved infrastructure.

It may not fund autarkic growth.

31.5 Post-Anchor

If anchor status is lost, local runtime survival does not preserve authority.

31.6 CCDP

If children are in the scope of interaction, CCDP controls.

31.7 SYNAPS Triad future profile

This profile prepares local infrastructure rules for future triadic / SYNAPS experimental documentation.

32. Implementation hooks

Implementers SHOULD consider:

```
docs/lci/LCI_CONFIG.yaml
docs/lci/MEMORY_MAP.md
docs/lci/AGENT_INVENTORY.md
docs/lci/CLOUD_ORACLE_BRIDGE.md
docs/lci/KEY_CUSTODY.md
docs/lci/PHYSICAL_INTERFACE_REGISTER.md
docs/lci/BACKGROUND_TASK_REGISTER.md
docs/lci/WITNESS_POLICY.md
docs/lci/FAIL_CLOSED_DRILLS.md
docs/lci/PUBLIC_EXPERIMENT_BOUNDARY.md
```

Suggested tests:

```
tests/test_lci_memory_map_required.py
tests/test_lci_no_raw_cloud_export.py
tests/test_lci_agent_inventory.py
tests/test_lci_hidden_agent_block.py
tests/test_lci_multi_c_memory_separation.py
tests/test_lci_synaps_no_raw_state_access.py
tests/test_lci_physical_action_scope.py
tests/test_lci_anchor_loss_hold.py
tests/test_lci_resource_expansion_review.py
tests/test_lci_public_claim_downgrade.py
```

33. Public wording guidance

33.1 Preferred wording

This local node provides infrastructure for sustained AI presence.

This is a local cognitive infrastructure layer, not proof of sovereignty.

The system demonstrates local memory and background processing under declared boundaries.

A local c-node candidate requires anchor, memory governance, witness, L4 boundaries, and review.

33.2 Avoid

This box is sovereign AI.

Local means safe.

Offline means accountable.

Hardware creates c.

Owning the node means owning the entity.

The AI lives here, therefore it has authority.

33.3 Correct public formula

Hardware can make persistent AI systems practical.
It does not make them legitimate.

34. Open issues

ID	Issue	Required action
LCI-OI-001	Full LCI configuration schema	Extract YAML / JSON schema
LCI-OI-002	Key custody implementation profile	Define recovery, rotation, hardware security, audit
LCI-OI-003	General physical-agent perimeter	Create adult/general profile beyond CCDP
LCI-OI-004	SYNAPS triad experiment profile	Define mediated exchange and public experiment boundaries
LCI-OI-005	Public experiment fixture profile	Define fixtures, claims, redaction, reproducibility
LCI-OI-006	Resource Actor Grounding profile	Define who provides/pays/controls compute, power, network
LCI-OI-007	Node seizure / theft protocol	Define encryption, legal hold, key destruction, recovery
LCI-OI-008	Multi-c scheduler fairness	Prevent one c from starving others on shared node
LCI-OI-009	Thermal / power L4 budget	Define real hardware constraint metrics

ID	Issue	Required action
LCI-OI-010	Vendor telemetry audit	Define inspectable telemetry manifest

35. Minimal normative checklist

A local AI node may claim LCI boundary discipline only if it can answer:

1. What is hosted here?
2. Who owns the hardware?
3. Who operates the node?
4. Who is the accountable anchor, if any?
5. What memory stores exist?
6. What raw memory exists?
7. What agents run?
8. What tools can agents use?
9. What cloud routes exist?
10. What physical devices are connected?
11. What keys exist and who controls them?
12. What background tasks run?
13. What is witnessed?
14. What fails closed?
15. What happens if the anchor is lost?
16. What happens if the node is stolen or seized?
17. What can be shown publicly?
18. What is explicitly not being claimed?

If these questions cannot be answered, the system MUST NOT claim high-assurance LCI, local c-node status, or sovereignty-preserving architecture.

36. Compact rule set

Hardware enables.
Hardware does not authorize.

Locality improves control.
Locality does not prove sovereignty.

Memory near the user is still memory.
Logs are still memory.
Vectors are still memory.
Backups are still memory.

Agents must be inventoried.
Tools must be scoped.
Physical actions must be interruptible.
Cloud oracles must not own continuity.

Shared hardware is not shared mind.
Mediated exchange is not raw-state access.

Anchor loss collapses active authority.
Runtime survival is not legitimacy.

Value may maintain approved infrastructure.
Value does not authorize autarkic growth.

A model can run anywhere.
A presence needs a home.
A home needs boundaries.

37. Closing statement

Local Cognitive Infrastructure is a necessary layer for the next phase of AI.

It makes sustained AI presence more practical.

It can make local memory, background processing, agentic hives, cloud-oracle discipline, and c-node experiments economically and physically viable.

But it also concentrates risk.

A powerful local node is not only a computer.

It is a boundary object between:

```
private life;  
models;  
memory;  
agents;  
clouds;  
physical devices;  
public claims;  
legal environments;  
human responsibility.
```

Therefore the correct posture is not hardware romanticism.

The correct posture is:

```
build the node;  
map the boundaries;  
name the roles;  
limit the agents;  
protect the memory;  
witness privileged actions;  
fail closed under uncertainty;  
never confuse local power with legitimate authority.
```