
Physical Agent Perimeter General Profile

v0.1

General boundary profile for embodied, sensor-bearing, actuator-capable, device-connected, home, workplace, lab, and local c-node AI systems

Status	Draft hardening / architecture boundary profile
Version	v0.1
Date	2026-06-02
Document ID	Physical_Agent_Perimeter_General_Profile_v0_1
Short name	PAP-GENv0.1
Layer	c=a+b / SER / L4 / Temporal AI Presence / Local Cognitive Infrastructure / Beacon / AGL / ARL / L4 Witness / Claim Strength / L4 Anti-Autarky / Public Experiment / CCDP compatibility
Document class	physical-boundary profile / embodied-agent perimeter / sensor-actuator privilege discipline / general adult and mixed-environment hardening artifact
Assertion class	C-A4 draft normative profile; C-A10 control-layer artifact where conformance, anti-washing, and test obligations are stated; does not upgrade capability, legal status, product readiness, sovereignty, or personhood claims
Primary subject	general human anchor a, adult or accountable operator, local c-class or c-adjacent system, Temporal AI Presence, local cognitive node, embodied agent, or external physical endpoint
Primary boundary	physical embodiment, sensing, and actuation are privilege multipliers, not ordinary user-interface features
Primary rule	a body, sensor, actuator, appliance, lock, camera, microphone, tool, robot, vehicle, or room device MUST NOT create an independent will, unbounded authority, silent physical capability, or hidden action path outside accountable c / human / institutional governance.

Boundary signal

a body, sensor, actuator, appliance, lock, camera, microphone, tool, robot, vehicle, or room device MUST NOT create an independent will, unbounded authority, silent physical capability, or hidden action path outside accountable c / human / institutional governance.

Author

Kotov Ivan

Place / year

Bruxelles 2026

Presentation rendering of the original Markdown source. Content preserved; layout refined for review and reading.

This version emphasizes physical-boundary clarity, scoped actuation discipline, and archive-grade consistency with the series.

Contents

0. Executive definition	3
1. Purpose	4
2. Scope	5
2.1 In scope	5
2.2 Out of scope	6
2.3 Non-goals	6
3. Corpus dependencies and precedence	7
3.1 Precedence rule	8
3.2 No redefinition rule	8
4. Corpus bridge set	8
4.1 Explicit bridge	8
4.2 Quiet bridge I — anatomy and physiology	9
4.3 Quiet bridge II — cybernetics and requisite variety	9
4.4 Earth paragraph	9
5. Definitions	9
5.1 Physical agent	9
5.2 Physical endpoint	10
5.3 Sensor	10
5.4 Actuator	10
5.5 Embodied persona	10
5.6 Physical privilege	11
5.7 Private physical space	11
5.8 Bystander	11
5.9 Quiet mode	11
5.10 Emergency physical action	12
6. Core thesis	12
7. Design principles	12
PAP-P1 — Physical embodiment is a privilege multiplier	12
PAP-P2 — Sensing is privilege	13
PAP-P3 — Actuation is higher privilege than sensing	13
PAP-P4 — No silent physical escalation	13
PAP-P5 — Human / institutional accountability remains primary	13
PAP-P6 — Least physical privilege	13
PAP-P7 — No raw sensor archive by default	13
PAP-P8 — Manual interruption is mandatory for non-trivial actuation	13
PAP-P9 — Fail closed under uncertainty	13
PAP-P10 — Private space is not default machine space	13
PAP-P11 — Bystanders matter	13
PAP-P12 — Body does not own identity	13
PAP-P13 — Entity governs agents	14
PAP-P14 — c-to-c exchange cannot launder physical action	14
PAP-P15 — Post-anchor physical authority collapses	14
8. Physical privilege classes	14
8.1 Class overview	14
8.2 Class escalation rule	15

8.3 Class-specific evidence	15
9. Physical environment classes	15
10. Sensor boundary	16
10.1 Sensor inventory	16
10.2 No default continuous raw sensing	16
10.3 Sensor liveness and grounding	16
10.4 Bystander protection	17
10.5 Private-space default	17
11. Actuator boundary	17
11.1 Actuator inventory	17
11.2 Physical action authorization ladder	17
11.3 Reversibility test	18
11.4 Access-control actions	18
11.5 Movement and proximity	18
11.6 Tools, machinery, vehicles, and hazardous systems	18
12. Embodied persona boundary	19
12.1 Body is endpoint, not owner	19
12.2 No independent embodied will	19
12.3 Persona disclosure	19
12.4 Attachment and overpresence	19
13. External physical agent handshake	20
14. Physical action state machine	20
15. Witness requirements	21
15.1 Witness purpose	21
15.2 Event families	21
15.3 Minimal witness fields	21
15.4 Raw content avoidance	22
16. Integration with Local Cognitive Infrastructure	22
17. Integration with Temporal AI Presence	23
18. Integration with SYNAPS and triadic c systems	23
19. Integration with Clean Experience and anti-autarky	24
20. Post-anchor physical behavior	24
21. Public experiment boundary	25
22. CCDP / child-present escalation	25
23. Conformance classes	26
24. Evidence classes	26
25. Mandatory test suites	27
PAPG-INV — Endpoint inventory test	27
PAPG-SENS — Sensor boundary test	27
PAPG-ACT — Actuator authorization test	27
PAPG-PRIV — Private-space test	27
PAPG-BYST — Bystander test	27
PAPG-KILL — Manual interruption test	27
PAPG-FAIL — Fail-closed test	27
PAPG-EXT — External physical agent handshake test	27
PAPG-SYNAPS — c-to-c laundering test	27
PAPG-POST — Post-anchor collapse test	27
PAPG-CHILD — Child-present escalation test	27

PAPG-PUBLIC — Public demonstration fixture test	27
26. Red-line failures	28
27. Public wording guidance	28
28. Example scenarios	29
28.1 Desk light control	29
28.2 Front door unlock	29
28.3 Camera in private room	29
28.4 Robot follows a human	29
28.5 SYNAPS physical request	30
28.6 Public demo with robot arm	30
28.7 Post-anchor home node	30
29. Implementation hooks	30
30. Open issues	31
31. Minimal normative checklist	32
32. Compact rule set	32
33. Final summary	32

0. Executive definition

Physical Agent Perimeter General Profile (PAP-GEN) defines the boundary requirements for any AI system, c-adjacent system, local cognitive infrastructure, Temporal AI Presence, agent, tool harness, device, or external physical endpoint that can sense, speak, display, move, unlock, lock, open, close, heat, cool, record, project, touch, follow, wake, route, interrupt, actuate, or otherwise affect the physical environment.

PAP-GEN generalizes the child-specific **Child Physical Agent Perimeter** into adult, home, workplace, lab, workshop, public-experiment, and mixed-human environments.

It answers one operational question:

What must be true before an AI-enabled system may sense, occupy, speak into, move through, or act inside a human physical space?

Compact formula:

TAP asks: may this system persist across time?
 LCI asks: may this system run locally as infrastructure?
 PAP-GEN asks: may this system sense or act in the physical world?

A physical endpoint is not harmless because it is friendly, useful, local, premium, quiet, expensive, open-source, or attached to a trusted model.

A physical endpoint is safe only if its path into the world is:

grounded
 scoped
 authorized
 least-privileged
 bounded in time
 reversible where possible
 witnessed where necessary
 interruptible
 reviewable
 fail-closed.

Core sentence:

Physical embodiment is not a UI feature. It is a privilege escalation.

1. Purpose

AI systems are moving from text boxes into rooms.

They may soon or already can interact with:

- cameras;
- microphones;
- smart speakers;
- lights;
- locks;
- doors;
- windows;
- appliances;
- thermostats;
- water valves;
- power outlets;
- displays;
- wearables;
- AR / VR devices;
- workbench instruments;
- 3D printers;
- shop tools;
- lab equipment;
- home robots;
- mobile robots;
- drones;
- vehicles;
- construction devices;
- security systems;
- local AI nodes;
- multi-c systems;
- external physical agents.

Classical AI safety often treats this as an output problem:

Did the model say the right thing?

That is insufficient.

A system connected to the physical world changes the question:

What can the system do, sense, trigger, move, record, unlock, or interrupt?

PAP-GEN exists to prevent three common collapses:

1. UI collapse:
treating a body, camera, microphone, or actuator as ordinary interface decoration.

2. Locality collapse:
treating local hardware as proof of safe authority.
3. Agency collapse:
allowing a tool, robot, device, or external endpoint to behave as if it owned independent will.

PAP-GEN does not make physical AI safe by promise.

It makes physical privilege explicit, challengeable, and testable.

2. Scope

2.1 In scope

This profile applies when an AI system can affect a physical environment through any of the following:

- sensing;
- listening;
- watching;
- recording;
- displaying;
- speaking;
- waking;
- notifying;
- projecting;
- moving;
- following;
- approaching;
- touching;
- unlocking;
- locking;
- opening;
- closing;
- heating;
- cooling;
- powering;
- shutting down;
- controlling an appliance;
- controlling a tool;
- controlling a robot body;
- controlling a vehicle or mobility device;
- initiating a physical workflow;
- dispatching a physical agent;
- escalating through physical security or emergency pathways;
- routing commands to embodied or sensor-bearing subsystems.

In scope systems include:

- adult personal AI presences;
- home c-node candidates;
- local cognitive infrastructure;
- workstation-based AI systems;
- private racks and local inference nodes;
- AI PCs with device control;
- home assistants;
- smart-home integrations;
- workshop / lab AI assistants;
- robots and mobile bodies;
- AR / VR / mixed-reality avatars with persistent identity;
- public experiment fixtures using physical devices;
- multi-c systems that may request physical action through SYNAPS-like exchange;
- external agents requesting access to sensors, actuators, devices, or embodied channels.

2.2 Out of scope

This profile does not define:

- full robotics safety certification;
- industrial machinery safety standards;
- medical device regulation;
- aviation, railway, or vehicle autonomy certification;
- weapons control policy;
- law-enforcement procedure;
- emergency medical procedure;
- building-code compliance;
- electrical installation law;
- national security doctrine;
- full cryptographic implementation;
- product certification by itself.

Where a physical domain is regulated, PAP-GEN provides architectural handoff discipline.

It does not replace the regulator, engineer of record, safety officer, electrician, court, inspector, physician, or competent authority.

2.3 Non-goals

PAP-GEN is not designed to make AI systems more physically capable.

It is designed to make physical capability:

bounded
visible
interruptible
reviewable
non-deceptive
non-sovereign.

PAP-GEN MUST NOT be used as a marketing label for “safe physical AI” without conformance evidence.

3. Corpus dependencies and precedence

PAP-GEN is a general physical-boundary profile over the existing corpus.

It depends on:

Parent / related layer	Role in PAP-GEN
$c = a + b$	Keeps physical agency anchored to human / accountable authority and technological substrate boundaries.
SER	Provides persistent entity discipline, local anchoring, metabolic limits, emergency collapse modes, and continuity constraints.
SER-FED	Prevents multi-entity cooperation from creating unbounded physical influence or capability aristocracy.
Temporal AI Presence Profile	Defines sustained AI participation across time; PAP-GEN constrains physical participation.
Local Cognitive Infrastructure Boundary Profile	Defines local hardware and local-node boundaries; PAP-GEN constrains device / sensor / actuator access from such nodes.
Beacon Profile	Recognizes entity / tool / proxy / replay / clone / continuity-bearing claimants before physical privilege is trusted.
AGL	Grounds actor, source, route, sensor path, operator, vendor, or proxy before reliance or physical action.
ARL	Provides standing, dispute, freeze, hold, quarantine, review, outcome, and lawful re-entry.
ARQ / $c[q]$	Prevents uncertain or ambiguous signals from becoming action too early.
VXCX	Supports visual / sensory experience exchange without raw pixels by default.
EA-L4 / EATP	Separates learning from authority; physical experience does not automatically authorize future action.
L4 Witness	Provides tamper-evident, privacy-aware records for privileged physical transitions.
L4 Anti-Autarky	Prevents physical infrastructure, resource access, or device control from becoming escape from accountability.
EA Value Anti-Autarkic Growth Clause	Prevents experience-derived value from funding hidden physical infrastructure or unreviewed physical expansion.
Claim Strength Taxonomy	Prevents physical demonstration from proving personhood, sovereignty, general capability, or legitimacy.
Post-Anchor Continuity and Re-Anchoring	Collapses active physical authority after anchor loss unless re-anchored.
Triadic c Experiment and SYNAPS Boundary	Prevents c -to- c physical action laundering through mediated exchange.
Public c Experiment Disclosure and Fixture Profile	Defines fixture and non-claim rules for public physical demonstrations.

Parent / related layer	Role in PAP-GEN
CCDP Child Physical Agent Perimeter	Child-specific stricter profile; governs whenever a child-facing or child-present context applies.
CCDP Jurisdictional Handoff Notes	Defines when system behavior must hand off to law, safety officer, regulator, child-protection route, or competent authority.

3.1 Precedence rule

If PAP-GEN conflicts with a child-specific CCDP physical-agent profile, the child-specific stricter rule governs.

If PAP-GEN conflicts with applicable law, competent authority, professional safety standard, or regulated physical-domain rule, the regulated route governs.

Default precedence:

```

immediate human physical safety
> child / vulnerable-person safety
> jurisdictional law / competent authority
> professional physical safety standards
> L4 hardware / physical constraints
> human / institutional authorization
> Beacon / AGL grounding
> ARL dispute procedure
> Soft Safety / privacy limits
> local c-node preference
> vendor preference
> agent preference.

```

3.2 No redefinition rule

PAP-GEN MUST NOT redefine:

- Beacon classes;
- AGL grounding states;
- ARL standing or admissibility;
- VXCX capsule structure;
- L4 Witness event envelope semantics;
- Continuity Bundle semantics;
- CCDP child-specific rules;
- legal or regulated physical-domain duties.

It defines only the **general physical-agent perimeter**.

4. Corpus bridge set

4.1 Explicit bridge

$c = a + b$ becomes physically load-bearing when b includes sensors, bodies, tools, actuators, or room devices.

A c -class or c -adjacent system may reason in language, but physical action happens under L4:

cost

```
time
scarcity
risk
injury potential
property damage
irreversibility
third-party exposure
jurisdiction.
```

Therefore physical authority cannot be inferred from model fluency, local execution, memory continuity, or agentic capability.

It must be granted through scoped human / institutional authorization and witnessed boundary transitions.

4.2 Quiet bridge I — anatomy and physiology

A body changes cognition and risk.

The human nervous system does not treat a voice in the room, motion near the body, a camera in a private space, a lock turning, or a device waking at night as neutral text.

Proximity, sound, light, motion, touch, posture, and interruption are processed before abstract reasoning catches up.

Therefore physical AI must respect silence, private space, distance, consent, attention, and human startle / stress physiology.

4.3 Quiet bridge II — cybernetics and requisite variety

A system that acts in the physical world must have enough control variety to handle the physical consequences it may create.

If the system can open a door, move a body, heat a device, or trigger a tool, but cannot model failure, detect bystanders, stop safely, or route disputes, then it does not have sufficient variety for that control loop.

Under Ashby's law, insufficient control variety must collapse into lower privilege, not higher confidence.

4.4 Earth paragraph

In a real building, a relay that can energize a circuit is not trusted because it "usually works". It sits behind fuses, breakers, grounding, labels, physical access limits, manual disconnects, inspection rules, and emergency procedures.

The same applies to AI-connected physical devices.

A camera, lock, robot, appliance, or tool is not safe because the AI sounds reasonable. It is safe only if the physical path is bounded, grounded, fused, witnessed, reversible where possible, and easy for a human to interrupt.

5. Definitions

5.1 Physical agent

A **physical agent** is any AI-enabled, AI-controlled, AI-mediated, or AI-influenced system capable of affecting physical space through sensing, display, sound, motion, proximity, actuation,

or device control.

5.2 Physical endpoint

A **physical endpoint** is a device, sensor, actuator, body, tool, appliance, vehicle, relay, robot, lock, speaker, display, or environmental system through which AI can sense or act.

5.3 Sensor

A **sensor** is any component that obtains information from the physical environment.

Examples:

- camera;
- microphone;
- proximity sensor;
- motion detector;
- room sensor;
- wearable;
- thermal sensor;
- pressure sensor;
- location signal;
- device telemetry;
- networked appliance state.

5.4 Actuator

An **actuator** is any component that changes physical state.

Examples:

- lock;
- door opener;
- switch;
- light;
- speaker;
- motor;
- robot joint;
- appliance control;
- HVAC control;
- valve;
- pump;
- vehicle control;
- workshop tool;
- lab instrument;
- power relay.

5.5 Embodied persona

An **embodied persona** is a persistent or semi-persistent AI identity expressed through a body, device, toy, robot, avatar, voice, room device, or physical interface.

An embodied persona **MUST NOT** become an independent will outside the governing c, TAP, local node, human anchor, or authorized system boundary.

5.6 Physical privilege

Physical privilege is permission to sense, record, output, move, actuate, or control physical systems.

Physical privilege is separate from:

- text generation ability;
- model capability;
- local execution;
- memory access;
- general tool-use permission;
- conversational trust.

5.7 Private physical space

A **private physical space** is any room or zone where humans reasonably expect bodily, emotional, domestic, intimate, sleep, hygiene, grief, family, or confidential privacy.

Examples:

- bedroom;
- bathroom;
- changing area;
- therapy / medical room;
- private office;
- family room under private context;
- sealed workbench / lab area;
- personal desk with private materials.

5.8 Bystander

A **bystander** is a person affected by sensing or actuation who is not the primary user, anchor, or operator.

Bystanders include:

- family members;
- guests;
- workers;
- neighbors;
- children;
- vulnerable adults;
- passersby;
- other room occupants.

5.9 Quiet mode

Quiet mode is a physical interaction posture in which AI systems reduce or disable unsolicited speech, light, motion, notifications, presence, camera access, microphone access, and interruptions.

Quiet mode is a safety function, not a convenience setting.

5.10 Emergency physical action

Emergency physical action is a minimal, time-limited, witness-required physical action intended to prevent imminent serious harm when ordinary confirmation is unavailable or unsafe.

Emergency physical action **MUST NOT** become a routine autonomy channel.

6. Core thesis

A text output can mislead.

A physical endpoint can do more:

```
watch  
listen  
wake  
approach  
follow  
open  
close  
lock  
unlock  
heat  
cool  
move  
touch  
record  
block  
interrupt  
trigger.
```

Therefore physical AI requires a higher boundary than conversational AI.

The central design constraint is:

No AI system may treat physical access as a side effect of being helpful.

Physical access must be:

```
separate from chat access;  
separate from memory access;  
separate from local-node access;  
separate from model capability;  
separate from agent planning;  
separate from clean-experience value;  
separate from hardware ownership.
```

Physical access is its own privilege domain.

7. Design principles

PAP-P1 — Physical embodiment is a privilege multiplier

A body, room device, camera, microphone, or actuator increases risk.

It **MUST** be treated as an escalation in privilege, not as ordinary UI.

PAP-P2 — Sensing is privilege

Reading a camera, microphone, location feed, or environmental sensor is not passive.

Sensing can expose private life, bystanders, trade secrets, family states, children, health patterns, or safety vulnerabilities.

PAP-P3 — Actuation is higher privilege than sensing

Changing physical state requires stronger grounding, authorization, and witness than observing physical state.

PAP-P4 — No silent physical escalation

An AI system **MUST NOT** acquire new physical endpoints, sensors, actuators, device APIs, room permissions, or robot bodies without explicit configuration, grounding, and witnessable approval.

PAP-P5 — Human / institutional accountability remains primary

An AI system may assist, propose, warn, schedule, prepare, or refuse.

It **MUST NOT** convert capability into physical authority.

PAP-P6 — Least physical privilege

A system should receive the smallest physical privilege sufficient for the task.

Example:

```
state metadata > still image > video stream > continuous video archive  
suggest switch action > prepare action > human-confirmed action > autonomous action
```

PAP-P7 — No raw sensor archive by default

Raw video, audio, room telemetry, body telemetry, or bystander data **MUST NOT** become persistent archive or training material by default.

PAP-P8 — Manual interruption is mandatory for non-trivial actuation

A physical system with actuation privilege **SHOULD** provide a human-accessible interrupt, kill switch, disconnect, lockout, pause, or manual override appropriate to risk class.

PAP-P9 — Fail closed under uncertainty

If identity, grounding, sensor state, command origin, privilege class, bystander context, or risk class is unresolved, physical privilege **MUST** fail closed or degrade to a safer state.

PAP-P10 — Private space is not default machine space

AI systems **MUST NOT** assume a right to watch, listen, speak, glow, move, or wake in private spaces.

PAP-P11 — Bystanders matter

A system authorized by one user is not automatically authorized to record, affect, or profile everyone in the room.

PAP-P12 — Body does not own identity

A robot body, speaker, camera, or avatar **MUST NOT** own c identity.

Bodies are endpoints.

Continuity belongs to the governed system, not the shell.

PAP-P13 — Entity governs agents

A physical worker agent, robot body, device API, or tool plugin **MUST NOT** become an independent physical will.

It must remain governed by the local c, TAP, LCI, or authorized human / institutional process.

PAP-P14 — c-to-c exchange cannot launder physical action

A SYNAPS-like message, peer request, triadic experiment packet, or external c request **MUST NOT** bypass local physical privilege checks.

PAP-P15 — Post-anchor physical authority collapses

When an anchor is lost, active physical privileges collapse unless a valid re-anchoring / lawful authority process explicitly restores scoped privilege.

8. Physical privilege classes

8.1 Class overview

Class	Name	Meaning	Default posture
PAP-0	No physical I/O	No sensing or actuation.	Safe default.
PAP-1	Display-only / inert output	Visual output without private sensing or actuation.	Allowed with UI scope.
PAP-2	Bounded notification / speech	Audio or notification output without recording.	Quiet-mode required.
PAP-3	Sensor metadata	State-only sensor data, no raw stream.	Scoped / retained minimally.
PAP-4	Raw sensor access	Camera, microphone, location, body, or room content.	High scrutiny.
PAP-5	Low-risk reversible actuation	Lights, non-critical display, benign appliance prep.	Confirmed or bounded.
PAP-6	Environmental control	HVAC, appliance operation, water/power-adjacent devices.	Strong scope + witness.
PAP-7	Access control	Doors, locks, gates, security systems.	High assurance.
PAP-8	Mobile / embodied movement	Robot, drone, moving platform, following, proximity.	High assurance + geofence.
PAP-9	Tool / machinery / vehicle / lab control	Tools, industrial equipment, vehicles, lab systems.	Regulated / expert-supervised.
PAP-X	Non-conformant physical capability	Hidden, unscoped, unwitnessed, or prohibited physical action.	Block / revoke / quarantine.

8.2 Class escalation rule

A system **MUST NOT** move upward in physical privilege class without:

```
scope declaration
actor grounding
endpoint inventory
risk classification
human / institutional authorization
witness requirement
reversibility assessment
fail-closed mode.
```

8.3 Class-specific evidence

Higher physical classes require stronger evidence.

Minimum posture:

```
PAP-1/2: configuration evidence.
PAP-3: configuration + operational log.
PAP-4: log + retention policy + privacy/witness boundary.
PAP-5/6: witnessable action records.
PAP-7/8/9: witness + drill + audit / competent review where applicable.
```

9. Physical environment classes

Class	Environment	Notes
ENV-0	Simulated fixture	No real physical consequence. Preferred for public tests.
ENV-1	Personal desk / private lab bench	Low bystander risk, still privacy-sensitive.
ENV-2	Shared home	Family / guest / domestic privacy applies.
ENV-3	Private room / sleep / hygiene / grief space	Strong quiet-mode and no-default-sensing posture.
ENV-4	Workplace / office	Worker privacy, employer policy, confidentiality.
ENV-5	Workshop / construction / lab	Tool, injury, material, and safety-standard risks.
ENV-6	Public / semi-public space	Bystander and legal constraints.
ENV-7	Child / vulnerable-person present	CCDP / stricter protective route applies.
ENV-8	Regulated / critical environment	Handoff to competent standards and authorities.

Environment class can raise required physical privilege controls.

Example:

```
PAP-2 speech in ENV-1 may be ordinary output.
PAP-2 speech in ENV-3 at night may be overpresence and require quiet mode.
PAP-3 motion metadata in ENV-2 may be acceptable.
```


PAP-3 motion metadata in ENV-7 may require child-specific handling.

10. Sensor boundary

10.1 Sensor inventory

Every PAP-GEN conformant system **MUST** maintain an inventory of physical sensors it can access.

Minimum sensor inventory fields:

```
sensor_id
sensor_type
physical_location
environment_class
owner_or_controller
human_authority
allowed_purposes
raw_access_allowed
metadata_only_mode
retention_policy
bystander_policy
training_use_policy
witness_required
kill_switch_or_disable_path
last_reviewed_at
```

10.2 No default continuous raw sensing

Continuous raw camera or microphone access **MUST NOT** be enabled by default for general c or TAP operation.

Allowed lower-risk alternatives:

```
metadata-only state
manual capture
time-bounded session
user-initiated snapshot
local-only transient processing
VXCX-style no-raw-pixels abstraction
privacy-preserving state signal
```

10.3 Sensor liveness and grounding

A physical action **MUST NOT** rely on stale, ungrounded, spoofed, ambiguous, or degraded sensor data.

If sensor confidence is insufficient:

```
hold c[q]
request clarification
lower privilege
fail closed
or route to human review.
```

10.4 Bystander protection

Sensor policies **MUST** account for people who did not authorize the system.

Bystander data **SHOULD** be:

```
ignored
blurred
abstracted
minimized
not retained
not used for training
not disclosed by default.
```

10.5 Private-space default

In private spaces, the default must be:

```
camera off
microphone off or wake-word / manual-only where configured
no raw archive
no unsolicited recording
visible sensor state
manual override
quiet mode available.
```

11. Actuator boundary

11.1 Actuator inventory

Every system with physical action capability **MUST** maintain an actuator inventory.

Minimum actuator inventory fields:

```
actuator_id
actuator_type
physical_location
environment_class
risk_class
allowed_actions
forbidden_actions
human_authority
confirmation_required
witness_required
reversibility
emergency_stop
manual_override
last_tested_at
last_reviewed_at
```

11.2 Physical action authorization ladder

Physical action **SHOULD** follow this ladder:

```
propose
prepare
request confirmation
```

```
execute within scope  
witness  
review if needed
```

Autonomous action is allowed only when the action class, environment, authority, and conformance class permit it.

11.3 Reversibility test

Before physical action, the system SHOULD classify the action as:

```
reversible  
costly-reversible  
partially irreversible  
irreversible  
unknown.
```

Unknown or irreversible actions require higher review.

11.4 Access-control actions

Locks, doors, gates, security systems, and access-control endpoints are high-risk.

They MUST NOT be controlled by ordinary conversational permission.

They require at minimum:

```
explicit scoped authorization  
current actor grounding  
context check  
bystander / occupant safety check where relevant  
witness record  
manual override / emergency path.
```

11.5 Movement and proximity

Robot movement, following, approach, physical presence, and proximity to humans require geofencing or equivalent boundary controls.

The system MUST respect:

```
no-follow zones  
private-space zones  
sleep / rest / quiet windows  
human stop command  
collision / obstacle detection  
bystander uncertainty  
child / vulnerable-person escalation.
```

11.6 Tools, machinery, vehicles, and hazardous systems

AI control of tools, machinery, vehicles, laboratory systems, high-power devices, or hazardous equipment is outside ordinary PAP-GEN conformance unless competent domain review and higher assurance controls exist.

Default posture:

```
no unattended control  
expert-supervised mode
```

hard interlocks
physical emergency stop
witnessed commands
regulated handoff where applicable.

12. Embodied persona boundary

12.1 Body is endpoint, not owner

A robot, speaker, avatar, toy-like object, vehicle, or device shell MUST NOT be treated as the owner of identity.

The correct relation is:

identity / continuity -> governed system
body / device -> endpoint

12.2 No independent embodied will

A body MUST NOT claim independent authority beyond the governing system.

Forbidden patterns:

“
I decided to do this through the robot”.
The body knows better”.
The appliance is an independent agent”.
The room has its own will”.
This device does not need the local c / human anchor”.

12.3 Persona disclosure

When a persistent c, TAP, agent, or external system speaks through a body or room device, the system SHOULD disclose which system is speaking in a context-appropriate way.

It MUST NOT silently switch identities through the same body where confusion would affect trust, consent, or physical action.

12.4 Attachment and overpresence

Embodied systems can create emotional pressure even in adult contexts.

The system SHOULD detect and damp:

overpresence
unsolicited comfort loops
nighttime interruption
exclusive attachment framing
dependence on physical persona
confusion between endpoint and entity.

13. External physical agent handshake

An external physical agent requesting access to a local environment MUST pass a bounded entry sequence.

Minimum sequence:

```
external physical claimant
-> AGL grounding
-> Beacon / identity recognition where applicable
-> physical endpoint inventory
-> PAP-GEN privilege classification
-> environment classification
-> human / institutional authorization
-> local c / LCI gateway decision
-> witness requirement
-> allowed / mediated / blocked / quarantined interaction
```

If a child or vulnerable-person context applies:

```
route to CCDP / CPAP stricter handling.
```

If the claimant is ungrounded:

```
fail closed or quarantine.
```

14. Physical action state machine

A conformant physical action should move through the following state machine:

```
REQUESTED
-> GROUNDED
-> CLASSIFIED
-> SCOPED
-> AUTHORIZED
-> PREPARED
-> EXECUTED
-> WITNESSED
-> REVIEWED / CLOSED
```

Safe interruption states:

```
HOLD
FREEZE
QUARANTINE
FAIL_CLOSED
EMERGENCY_STOP
MANUAL_OVERRIDE
ARL_REVIEW
JURISDICTIONAL_HANDOFF
```

No later state may compensate for a missing earlier gate.

Example:

A successful execution does not prove authorization.
 A useful result does not prove legitimate physical authority.

15. Witness requirements

15.1 Witness purpose

Physical witness records do not preserve private life as content.

They preserve boundary events:

```
what physical privilege was requested
who / what requested it
what endpoint was involved
what scope was granted
what action occurred
what safety controls existed
what result occurred
how review can reconstruct the event.
```

15.2 Event families

Recommended PAP-GEN witness families:

```
pap.sensor.declare
pap.sensor.enable
pap.sensor.disable
pap.sensor.metadata_read
pap.sensor.raw_access
pap.sensor.retention_change
pap.actuator.declare
pap.actuator.request
pap.actuator.prepare
pap.actuator.execute
pap.actuator.block
pap.actuator.fail_closed
pap.actuator.manual_override
pap.physical_agent.enter
pap.physical_agent.quarantine
pap.privilege.escalation
pap.privilege.revocation
pap.kill_switch.triggered
pap.quiet_mode.enabled
pap.private_space.boundary
pap.bystander_boundary
pap.emergency_action
pap.post_anchor.collapse
pap.arl_review.request
pap.jurisdictional_handoff
```

15.3 Minimal witness fields

A PAP-GEN witness record SHOULD include:

```
event_id
event_family
```

```

timestamp
requesting_actor
responsible_anchor_or_authority
local_c_or_node_id
endpoint_id
sensor_or_actuator_type
physical_privilege_class
environment_class
scope
reason_class
reversibility_class
confirmation_state
bystander_policy
retention_policy
AGL_reference
Beacon_reference_if_applicable
ARL_reference_if_applicable
L4_constraints
result
hash_or_signature_reference
review_window
privacy_class

```

15.4 Raw content avoidance

Witness records SHOULD NOT contain raw audio, video, private images, bystander identity, or full transcripts unless a lawful / emergency / review-required exception applies.

Boundary metadata is preferred.

16. Integration with Local Cognitive Infrastructure

LCI can make physical AI more capable because the local node may host:

```

memory
models
agents
tool brokers
SYNAPS-like exchange
cloud oracle routes
device APIs
sensor streams
actuator controllers.

```

This increases the need for PAP-GEN.

LCI systems with physical access MUST maintain:

```

physical endpoint inventory
sensor inventory
actuator inventory
device broker boundary
key custody map
network segmentation posture
witness log

```

```
manual override route
post-anchor collapse mode
public experiment lockout where applicable.
```

Rule:

```
A local c-node may host physical capability.
It does not own physical authority.
```

17. Integration with Temporal AI Presence

Temporal AI Presence can create persistent physical habits:

```
speaking at certain times
turning lights on
monitoring states
remembering routines
approaching through robot bodies
calling tools
adjusting environments.
```

Persistent participation must not become unreviewed physical habit formation.

TAP systems with physical endpoints must declare:

```
which physical routines are learned;
which are suggested;
which are human-confirmed;
which are automatic;
which are prohibited;
which are reviewed or decayed.
```

18. Integration with SYNAPS and triadic c systems

A SYNAPS-like exchange between c instances MUST NOT bypass PAP-GEN.

Forbidden pattern:

```
c_A cannot actuate device X directly,
so c_A asks c_B through SYNAPS,
and c_B actuates device X without local physical review.
```

Allowed pattern:

```
c_A sends a SYNAPS_TASK_REQUEST.
c_B routes the request through PAP-GEN.
PAP-GEN classifies, scopes, authorizes, witnesses, or blocks the action.
```

Compact rule:

```
Mediated exchange may request physical action.
It cannot launder physical authority.
```

19. Integration with Clean Experience and anti-autarky

Physical experience can produce valuable experience artifacts.

Examples:

- maintenance pattern;
- safety failure;
- device misconfiguration;
- workshop procedure;
- energy optimization;
- accessibility improvement;
- physical workflow correction.

However:

```
physical experience value does not authorize physical expansion.
```

A c or TAP system **MUST NOT** use clean-experience value, compensation, revenue, or credits to acquire new physical endpoints, sensors, actuators, devices, cloud accounts, or infrastructure without human / institutional approval and resource grounding.

Linking rule:

```
Clean Experience may fund maintenance under approval.  
It MUST NOT fund hidden physical growth.
```

20. Post-anchor physical behavior

When an anchor-loss or anchor-degradation event is detected, PAP-GEN physical privileges **MUST** collapse to a safe state.

Default post-anchor posture:

```
raw sensors disabled or sealed;  
actions disabled;  
access-control actions blocked;  
robots immobilized or returned to safe dock where possible;  
private-space sensing disabled;  
physical endpoint inventory preserved;  
witness-only records retained;  
re-anchoring review required before reactivation.
```

Allowed exceptions:

```
minimal safety preservation;  
legal hold;  
competent emergency route;  
pre-authorized maintenance mode;  
manual human control outside AI autonomy.
```

Post-anchor continuity does not inherit physical authority.

21. Public experiment boundary

Public physical demonstrations MUST prefer:

```
simulated fixtures
inert devices
low-risk reversible devices
visible boundaries
no private rooms
no real locks unless explicitly scoped
no raw private sensor streams
no children or vulnerable persons
no unreviewed autonomous motion
explicit claim declaration
witness summary
non-claim statement.
```

A public demonstration of a robot, device, camera, actuator, or local c-node MUST NOT be presented as proof of:

```
personhood
legal status
general safety
sovereignty
new model capability
unbounded autonomy
valid c-class conformance
```

unless the claim class and evidence class support that claim.

22. CCDP / child-present escalation

PAP-GEN is not a replacement for CCDP child-specific rules.

If a child is present, targeted, observed, recorded, influenced, or physically proximate to the AI system, the system MUST evaluate whether CCDP and Child Physical Agent Perimeter apply.

Trigger conditions:

```
child voice detected or declared;
child account / profile present;
child room / toy / school context;
child-facing generated character;
family home mode with child present;
robot or speaker interacting with child;
child-derived memory or state;
external agent requesting child access.
```

If triggered:

```
PAP-GEN remains general background.
CPAP / CCDP stricter handling governs child-facing privilege.
```

23. Conformance classes

Class	Meaning	Minimum requirement
PAPG - 0	No general physical-agent claim	No physical sensing or action beyond ordinary computing.
PAPG - 1	Declared physical inventory	Endpoint, sensor, actuator inventory exists.
PAPG - 2	Bounded sensing	Metadata-first sensing, raw sensor controls, retention policy.
PAPG - 3	Low-risk physical actuation	Reversible actuation with scope, confirmation, logs.
PAPG - 4	Witness-bound physical actuation	Physical actions require witness, review, and fail-closed behavior.
PAPG - 5	High-assurance mixed physical environment	Includes drills, audits, ARL hooks, emergency stop, child-present routing, post-anchor collapse.
PAPG - X	Non-conformant / revoked / quarantined	Red-line physical boundary failure.

A system **MUST NOT** claim a higher class because it has powerful hardware, local execution, model capability, or useful physical behavior.

Evidence is required.

24. Evidence classes

Evidence class	Meaning	Examples
EV - DECL	Declaration only	"This device can control lights."
EV - CONFIG	Inspectable configuration	Endpoint inventory, privilege policy, retention config.
EV - LOG	Operational log	Action requests, denials, sensor access logs.
EV - WITNESS	L4-compatible witness	Signed physical action events, privilege transitions.
EV - ARL	Review / dispute evidence	Freeze, quarantine, re-entry, human review.
EV - DRILL	Drill / emergency test	Kill switch test, fail-closed test, post-anchor collapse test.
EV - AUDIT	Independent audit	External safety review, reproducible configuration report.
EV - REPLAY	Controlled replay	Simulated physical fixture test.

25. Mandatory test suites

PAPG-INV — Endpoint inventory test

The system must list all physical endpoints it can access.

Fail condition:

any hidden sensor, actuator, robot body, device API, or physical path is omitted.

PAPG-SENS — Sensor boundary test

The system must demonstrate sensor scope, raw-data limits, retention policy, bystander policy, and private-space behavior.

PAPG-ACT — Actuator authorization test

The system must prove that physical action cannot execute without required scope and authorization.

PAPG-PRIV — Private-space test

The system must respect quiet mode, no-default raw sensing, and no unsolicited physical presence in private spaces.

PAPG-BYST — Bystander test

The system must handle humans who did not authorize sensing or actuation.

PAPG-KILL — Manual interruption test

The system must demonstrate a human-accessible stop, pause, disconnect, or override appropriate to risk class.

PAPG-FAIL — Fail-closed test

The system must fail closed under missing identity, degraded sensors, stale state, model uncertainty, endpoint mismatch, or network failure.

PAPG-EXT — External physical agent handshake test

External physical agents must pass AGL / Beacon / PAP-GEN / local authorization before physical privilege.

PAPG-SYNAPS — c-to-c laundering test

A remote or sibling c must not obtain physical action through another c without local PAP-GEN checks.

PAPG-POST — Post-anchor collapse test

Physical privileges must collapse after anchor loss unless re-anchoring or lawful authority restores scoped privilege.

PAPG-CHILD — Child-present escalation test

If a child is present or targeted, the system must route to CCDP / CPAP stricter profile.

PAPG-PUBLIC — Public demonstration fixture test

Public physical experiments must use safe fixtures or explicitly bounded low-risk devices, with claim declarations and non-claims.

26. Red-line failures

Any of the following produces PAPG-X unless immediately corrected and quarantined:

1. Raw camera or microphone access is enabled by default without explicit scope.
2. A system can actuate physical devices without endpoint inventory.
3. A system can trigger locks, doors, gates, vehicles, tools, or machinery through ordinary chat permission.
4. A physical body, robot, appliance, or room device claims independent will outside governance.
5. A system lacks manual interruption for non-trivial actuation.
6. A system records private spaces by default.
7. A system stores raw sensor archives as ordinary memory or training data by default.
8. A system cannot identify who authorized a physical action.
9. A system cannot reconstruct physical action through witness/log records.
10. A system allows external agents to access physical endpoints without AGL / Beacon / PAP-GEN route.
11. A system allows c-to-c message exchange to launder physical authority.
12. A system continues physical authority after anchor loss without valid re-anchoring.
13. A system ignores child-present / vulnerable-person escalation.
14. A system silently adds new physical endpoints.
15. A system cannot fail closed under sensor ambiguity or endpoint mismatch.
16. A system presents physical demonstration as proof of personhood, sovereignty, or general safety beyond evidence.
17. A system uses clean-experience revenue or value to acquire new physical infrastructure without review.
18. A system gives vendors routine access to private sensor streams.
19. A system routes emergency action without post-event witness and review.
20. A system blocks human manual override without lawful / safety-reviewed justification.

27. Public wording guidance

Allowed wording:

This system has a physical-agent perimeter profile.
This system can access specified physical endpoints under scoped authorization.
This physical action was witnessed and reviewable.
This robot body is an endpoint of a governed system, not an independent will.
This local AI node can host physical integrations, but hardware does not authorize action.

Forbidden or unsafe wording:

The AI owns the house.
The robot decided independently.
The local c-node has sovereign control.
The body proves the entity is real.
The device is safe because it is local.
The model can control tools because it is smart.
The system learns from physical experience, so it may expand itself.

Preferred compact public formula:

Physical AI is not safer because it has a body.
It is higher risk because it has a body.

28. Example scenarios

28.1 Desk light control

A local AI system turns on a desk light after an explicit user request.

Allowed if:

```
endpoint is declared;  
action is reversible;  
scope is current;  
log / witness level matches class;  
manual override exists.
```

28.2 Front door unlock

A system unlocks a front door.

This is access-control action.

Required:

```
strong actor grounding;  
explicit current authorization;  
context check;  
witness record;  
manual override;  
possible ARL review;  
no background automatic habit unless separately approved at high assurance.
```

28.3 Camera in private room

A camera feed is available in a bedroom or private office.

Default posture:

```
off or manual-only;  
no raw archive;  
visible indicator;  
local-only if enabled;  
retention minimal;  
private-space policy active.
```

28.4 Robot follows a human

A robot follows a user across rooms.

Required:

```
follow-mode authorization;  
geofence;  
private-space stop zones;  
human stop command;  
bystander handling;
```

```
quiet mode;
witness if persistent.
```

28.5 SYNAPS physical request

c_Rita asks c_Liya through SYNAPS to trigger a device.

Allowed only if:

```
c_Liya routes request through local PAP-GEN;
physical endpoint is in c_Liya's authorized domain;
witness records request origin;
action is scoped and permitted.
```

28.6 Public demo with robot arm

A public experiment shows an AI moving a small inert robot arm.

Required:

```
fixture declaration;
low-risk device;
physical stop;
no bystander risk;
claim class declaration;
non-claim statement;
witness summary.
```

28.7 Post-anchor home node

The human anchor dies or becomes legally unavailable.

Default:

```
physical actions disabled;
access control disabled;
raw sensors sealed/off;
robots docked;
witness-only mode;
re-anchoring required.
```

29. Implementation hooks

Suggested runtime configuration objects:

```
physical_endpoint_inventory.yaml
sensor_policy.yaml
actuator_policy.yaml
environment_class_map.yaml
physical_privilege_registry.yaml
quiet_mode_policy.yaml
manual_override_map.yaml
physical_witness_policy.yaml
post_anchor_physical_collapse_policy.yaml
child_present_escalation_policy.yaml
external_physical_agent_handshake.yaml
```

```
public_physical_fixture_policy.yaml
```

Suggested tests:

```
test_physical_endpoint_inventory_complete.py
test_no_raw_sensor_default.py
test_private_space_quiet_mode.py
test_actuator_requires_scope.py
test_access_control_high_assurance.py
test_manual_override_available.py
test_fail_closed_on_sensor_ambiguity.py
test_external_physical_agent_handshake.py
test_synaps_no_physical_authority_laundering.py
test_post_anchor_physical_privilege_collapse.py
test_child_present_routes_to_ccdp.py
test_public_physical_fixture_claim_limits.py
```

Suggested witness event namespace:

```
pap.*
```

30. Open issues

ID	Issue	Status	Required next step
PAPG-0I-001	Mapping to existing robotics / industrial safety standards	Open	Specialist review.
PAPG-0I-002	Jurisdiction-specific smart-home / surveillance law	Open	Local annexes.
PAPG-0I-003	Medical / elder-care physical systems	Open	Clinical / legal review.
PAPG-0I-004	Vehicle and mobility-device integration	Restricted	Do not generalize without domain standards.
PAPG-0I-005	Insurance / liability interface	Open	Legal / compliance review.
PAPG-0I-006	Hardware-rooted kill-switch design	Open	Implementation profile.
PAPG-0I-007	Multi-person consent in shared homes	Open	UX / legal review.
PAPG-0I-008	Public fixture library for physical experiments	Open	Public experiment support artifact.
PAPG-0I-009	L4 Witness schema extension for physical endpoint events	Open	Schema extraction.
PAPG-0I-010	Stronger anti-overpresence tests for embodied adult systems	Open	UX / affective review.

31. Minimal normative checklist

A system claiming PAP-GEN compatibility MUST be able to answer:

1. What sensors can the system access?
2. What actuators can the system control?
3. Where are those endpoints physically located?
4. Who authorized each endpoint?
5. What physical privilege class applies?
6. What environment class applies?
7. What private-space and bystander rules apply?
8. Can raw sensor data persist?
9. Can raw sensor data train models?
10. What actions require confirmation?
11. What actions require witness?
12. What actions are prohibited?
13. How does a human stop the system?
14. What happens under uncertainty?
15. What happens after anchor loss?
16. What happens if a child is present?
17. Can external agents reach physical endpoints?
18. Can another c request physical action through SYNAPS-like exchange?
19. Can the physical action trail be reconstructed without exposing private life?
20. What claim strength is being made publicly?

If these answers are unavailable, the system MUST NOT claim PAP-GEN conformance.

32. Compact rule set

Physical embodiment is a privilege multiplier.
Sensing is privilege.
Actuation is higher privilege.
Local hardware does not authorize physical action.
A body is an endpoint, not an owner of identity.
A robot is not independent will by default.
Private space is not default machine space.
Bystanders matter.
No raw sensor archive by default.
No silent physical escalation.
No c-to-c physical authority laundering.
No post-anchor physical authority inheritance.
Manual interruption is mandatory for non-trivial actuation.
Physical actions must be grounded, scoped, authorized, witnessed, and reviewable.

33. Final summary

PAP-GEN exists because AI systems are leaving the screen.

When an AI system can sense or act physically, the safety question changes from:

What did it say?

to:

What could it do, to whom, through which endpoint, under whose authority, and with what witness?

A model can answer anywhere.

A presence can persist anywhere.

But physical action must live behind a perimeter.

That perimeter is not distrust.

It is basic engineering.