
Presence Shift® Termination Standard

A conceptually authoritative overview of the primitive and its place in the stack.

PSTS is the layer that makes an owner-defined stop boundary operationally final and turns that finality into structured evidence that can be reviewed.

Why this exists

Many organizations can identify the point at which an AI interaction should no longer continue.

Far fewer can prove that the system actually stopped once that decision had already been made.

The market is already moving toward agent monitoring, supervision, and “guardian” layers. That is a meaningful shift. But monitoring, alerts, and general oversight are not the same thing as a stop boundary that becomes operationally final and later reviewable.

PSTS exists for the narrower problem: once the system owner has decided participation must stop, can the system make that stop real and generate structured evidence that it did?

Some of the deepest pressure appears where systems can keep participating past the point where continuation is supportive, trustworthy, or autonomy-preserving. That shows up especially in personal or empathic AI experiences, and in multi-actor runtime systems where one local refusal does not actually stop the whole governed boundary.

What PSTS is

PSTS is a governance-layer standard for owner-defined stop boundaries.

It sits at the system-level enforcement boundary, outside the governed interaction, and defines what must become true after the owner has already decided that participation must stop.

Within the configured governed boundary, PSTS requires:

- no further system output
- irreversible termination within that boundary
- append-only audit artifacts
- clear authority attribution

What PSTS is not

PSTS is not:

-
- a broad guardian-AI cloud application
 - a general monitoring or alerting layer
 - a language or risk interpretation system
 - a safety-policy authoring system
 - a support or care layer
 - a replacement for internal decision logic

PSTS does not decide when termination should happen. It governs what must happen next after that decision has already been made by the system owner.

Why this matters now

As serious AI systems become more persistent, stateful, personal, and multi-actor, the question is no longer only whether they are helpful. It is also whether they know where participation should stop, and whether that stopping point can be trusted.

Systems that claim safety, support, or care will increasingly be judged not only by what they say they do, but by whether they can produce trusted proof that their boundaries were real in operation.

Where PSTS fits

Open agent runtimes, orchestration systems, and supervisory environments can be relevant PSTS environments when they control enough of the enforcement boundary to make the owner-defined stop final.

They are complements to PSTS, not substitutes for it.

Adapters can help adoption. The business value is trusted proof.

PSTS is portable across runtimes precisely because its value is not one local feature of one stack.

Strong early lanes and pilot shapes

The deepest PSTS need appears most clearly in two primary lanes:

1. **Relational or empathic stop / shift / handoff** — the system should no longer continue participating in a personal-feeling interaction unless it stops, shifts, or hands off.
2. **Runtime/orchestration continuation cutoff** — in a multi-model, multi-agent, or tool-using system, one local refusal is not enough; the whole governed boundary has to stop.

A practical first pilot does not always begin in the deepest case. Many early pilots begin in cleaner, high-consensus shapes such as:

-
- post-decision denial
 - human-review-denied irreversible action

Those are useful because the stop decision is already crisp and the proof question is concrete: once the owner had already decided to stop, did the system actually stop?

A practical first pilot

A PSTS pilot is narrow, authority-bound, and designed to answer one question:

Can this organization enforce a final stop once it has decided to stop?

A pilot is not a broad redesign. It is a high-value proof step:

- one boundary
- one environment
- one evidence package
- one review readout

That is what makes a narrow PSTS pilot the most economical and rigorous way to refine a serious stop boundary in practice.

Access

PSTS is offered through a narrow licensing and pilot process for qualified system owners.

Qualified system owners, pilot partners, and licensees can request the full implementation corpus — including the boundary contract, audit semantics, log schema, configuration schema, compliance evidence package, and validation harness requirements — through the PSTS pilot and licensing process.