Cheryl Burris
EDIT 891-2
11 May 2026

Design Decision Record: Design Decision 21

calmMode Additive-Only Constraint

NOTE: This entry documents a design behavior implemented in V6-B that was not recorded in the Design Decision Record at build time. The behavior was flagged in Blind Spots Register v1.6 (4 May 2026) as undocumented. This entry closes that flag.

CLOSES: Blind Spots Register v1.6 annotation; undocumented calmMode additive-only behavior (4 May 2026).

OPENS: BS-2 agency tension requires a written argument in a theoretical framework. See Failure Taxonomy v2, Level 3 section, for framing.

Decision number	21
Date	7 May 2026
Planner Version	v6
Status	Implemented. Previously undocumented.
Population(s)	Both (ADHD and ASD). Applies via structurally distinct mechanisms.
SRL phase	Forethought (Check-in) and Performance (session-level toggle)
Blind Spot(s)	BS-2 (Agency Architecture). See tension section below.
Related feature	• Calm Mode toggle • ADHD Cognitive Load Check-in (v5-1) • ASD Sensory Load Check-in (v4)

The Decision

calmMode is additive-only. The Calm Mode toggle can elevate interface quieting from standard to calm, but it cannot reduce quieting below the level established by the population-specific check-in.

Operationally:

- A user on the ASD Medium path (sensoryLoad = Medium; interfaceMode already resolved to calm) who disables the Calm Mode toggle does not return to the standard interface. The interface remains calm.
- A user on the ADHD Scattered path (adhdCogLoad = Scattered; calmMode auto-enabled at check-in) cannot disable calmMode mid-session. Returning to standard requires returning to Check-in and completing the check-in again.
- Minimal mode (ASD Heavy or ADHD Overloaded) is unaffected by the toggle in either direction. Minimal always wins.

Technically, the effectiveInterfaceMode state variable resolves as follows:

(1) If isMinimal, output is minimal
(2) If calmMode, output is calm
(3) Otherwise, output is standard

The toggle sets calmMode, but cannot unset a calmMode that was auto-enabled by a check-in response.

Rationale

Two design obligations are in tension here, and this decision resolves that tension in favor of the check-in as a binding instrument:

- **Obligation 1, Check-in validity:** The cognitive and sensory load check-ins are tools for measuring state and setting the baseline interface conditions for the session. Allowing users to disable the check-in response with a toggle would turn the process into a performative act rather than a functional one: users complete it, the system reacts, and then they immediately undo the reaction. This would reduce the check-in to a ritual without structural significance, which directly contradicts the framing of arousal and activation by Isaac et al. (2024) and the approach to minimizing extraneous load by Le Cunff et al. (2024).
- **Obligation 2, Protective constraint:** Interface quieting in this design is not a preference control. It is a protective mechanism grounded in Raymaker et al. (2020): the planner must not impose demands on the user that exceed the user's current regulatory capacity. The check-in establishes the floor of protection for the session. The toggle allows the user to request more protection than the check-in prescribed, but not less.

Theoretical Grounding

Source	Role in this decision
Isaac et al. (2024)	ADHD arousal/activation variability. calmMode auto-enable on Scattered is the direct design translation of the arousal framing: Scattered = suboptimal activation state; calm interface reduces extraneous stimulation during that state.
Le Cunff et al. (2024)	Extraneous cognitive load mechanism. Calm mode reduces visual and structural complexity, lowering extraneous load during sessions where germane load capacity is already reduced. The additive-only constraint preserves this mechanism from being undone mid-session.
Raymaker et al. (2020)	No-punitive-debt, protective design constraint. The check-in-established floor is the minimum protective response the planner provides. It cannot be removed by user action within the same session.
Zimmerman (2002)	Forethought phase. The check-in occurs during Check-in (Forethought). The additive-only constraint preserves the Forethought-phase boundary: interface conditions are set before the session begins and are not renegotiated mid-session without returning to Forethought.
Townes et al. (2023)	Transdiagnostic convergence. Justifies applying a shared effectiveInterfaceMode output across ADHD and ASD paths while maintaining structurally distinct input mechanisms.

Design Tension (BS-2: Agency Architecture)

This decision is a scaffold-for-deficit choice. It prioritizes the system's interpretation of the user's state (as expressed through the check-in) over the user's subsequent judgment of their own interface needs.

The tension is that the user completed the check-in truthfully at the start of the session. But cognitive states change mid-session. A user who was genuinely Scattered at Check-in may regulate to Focused by mid-session. Under the additive-only constraint, the user cannot reduce their interface quieting to reflect their improved state without returning to Check-in. This is a constraint on self-authorship.

This tension is not resolved by this decision. It is the decision. The design supports scaffold-for-deficit (protective floor) and self-authorship (user override) simultaneously, opting for the protective floor within a single session.

There is an argument for holding both. DBR methodology allows a design to embody an unresolved theoretical tension as an empirical question for the next phase. The question is whether users experience the additive-only constraint as protective or as restrictive. That is a finding for the next phase, not a claim for this one.

Alternative Considered and Rejected

Alternative: Allow full user override. The toggle can both increase and decrease quieting regardless of the check-in baseline.
Reason rejected: Full override would require a mechanism to reset the check-in baseline mid-session, which in practice means re-running Check-in. Re-running Check-in mid-session violates the Forethought-phase boundary (Zimmerman, 2002) and introduces navigation burden. A mid-session re-check-in is a design possibility for the next phase, not an implementation for this one.

Phase 4 Measurement Threats

THREAT 1: Check-in as preference override (ADHD-specific)

- Users may begin selecting Scattered habitually not because they are scattered but because they prefer the calm interface.
- Once this pattern is established, calmMode auto-enables on every session, and the check-in ceases to function as a state-measurement instrument.
- Self-reported check-in state becomes unreliable as a Phase 4 data point.

Mitigation: Cross-check self-reported Scattered selection against behavioral indicators (session engagement rate, task entry rate, Pomodoro completion rate). If Scattered is reported, but behavioral indicators are consistent with Focused sessions, flag as a preference pattern rather than a state report.

THREAT 2: calmMode as dependency rather than support (ASD-specific)

- For ASD users, the calm interface may become a required precondition for session entry rather than a responsive support.
- The Failure Taxonomy v2 identifies this as an L3 ASD-v failure mode since calmMode becomes a safety object, and any deviation from it triggers distress.

Mitigation: Track whether ASD users who consistently select Medium or Heavy across four or more sessions show increasing distress indicators in Off-Boarding self-reports, or whether their tolerance for interface variation remains flat or improves.