

NSIF

Nervous System Infrastructure Framework

Capacity-Safe Systems Standard

Version 3.0

A human systems standard for how any human-facing system manages demand when the people operating it are not at full capacity.

*Every system built for humans was designed at full capacity.
None of them work when the human operating them doesn't.
The failure was never human. It was always architectural.*

Published	2026 · Whelmish LLC
Author	Vanessa Williams · Founder, Whelmish LLC
Framework	NSIF — Nervous System Infrastructure Framework
Version	3.0 — Synthesized Standard
Status	Active — Open for Practitioner and Organizational Adoption
Domain	Human-Facing Products, AI Systems, Organizations, Schools, Healthcare, and Communities
Contact	capacitystandard.com · hello@capacitystandard.com

This document is the foundational intellectual property of the Nervous System Infrastructure Framework (NSIF). It may be cited, referenced, and applied with

attribution. Commercial use, reproduction, or certification claims require written authorization from Whelmish LLC.

A Promise Before the Standard

You do not need to read this entire document to know what it means.

If you have ever abandoned a planner, quit an app, failed to show up, disappeared from a system, or called yourself broken for not being able to keep up — this document is for you.

You were not failing. You were in a capacity state that the system was not designed for.

The failure was the design.

Every system built for humans was designed at full capacity. This standard exists to change that. Not for some users. For all of them. Not as an accommodation. As the base requirement.

I am regulated, therefore I can.

If you are a practitioner, a product builder, a school administrator, or an organization leader — the standard begins on the next page. It is rigorous, citable, and defensible. It will survive expert challenge.

Both of you are why this standard exists.

A Warning Before the Standard Begins

This framework can be misused.

It can be adopted aesthetically while the underlying system continues to extract from low-capacity users. A product can use gentle colors, warm language, and recovery-spine framing while its core business model still requires people to be dysregulated to convert, retain, or engage. A certification can become a greenwashing mechanism. An organization can use capacity language to increase compliance, suppress dissent, or make structural harm feel softer.

This standard exists to prevent that. But no document prevents misuse on its own.

The line is this:

If your core business model requires people to be in a low-capacity state in order to work — if distress drives your conversions, fear drives your retention, urgency drives your clicks, or shame drives your re-engagement — the Capacity-Safe mark is not for you.

No amount of warm color palettes, gentle copy, or breathing prompts makes a distress extraction system Capacity-Safe. The mark certifies structure, not aesthetics. It certifies demand architecture, not tone. It certifies that the system works for humans at their lowest — not that the system profits from them being there.

If you are reading this to understand how to design genuinely safer systems, welcome. Everything that follows is for you.

If you are reading this to find language that sounds safe while keeping the extraction, the anti-gaming provisions in Part XIII and the misuse prevention rules in Part XI are specifically designed to surface and name what you are doing.

The practitioner program is not a revenue stream. It is the framework's immune system.

Preamble

This standard exists because of a design failure so universal it became invisible.

For generations, every system humans interact with — schools, workplaces, productivity tools, medical environments, digital products, parenting frameworks, government services, AI interfaces — was architected around one assumption: the human using it is available. Regulated. Resourced. At capacity.

That assumption has never been true.

Humans move through states of overwhelm, shutdown, grief, depletion, hyperarousal, recovery, and flourishing every single day. Children do it visibly. Adults learn to hide it. Entire industries are built around the pretense that it doesn't happen.

The result is a world full of systems that punish people for being human.

The Capacity-Safe Systems Standard does not ask humans to adapt to broken systems. It asks systems to be designed around the actual range of human capacity — including its lowest points and its highest ones.

This is not accommodation. It is not therapy. It is not wellness culture.

It is infrastructure. The missing layer underneath everything else humans build, do, and become.

Capacity cannot be demanded. It can only be restored.

This standard defines what it means for any human-facing system — a product, an interface, an organization, an AI, a physical space, a collective environment — to be Capacity-Safe.

This is not a design document in the narrow sense. It applies to any system a human operates inside — products, AI, workplaces, schools, hospitals, communities, and policy. Design teams implement it. So do HR directors, hospital administrators, school boards, operations leaders, and anyone who builds or runs systems that touch human beings.

It is built from converging evidence across nervous system regulation, stress physiology, trauma-informed practice, accessibility, human-centered design, cognitive load, developmental science, evolutionary biology, UX research, and lived experience. It does not depend on any single theory. Its core claim is operational: when systems exceed human capacity, the design must change before the human is blamed.

Part I · Foundations

1.1 The Problem This Standard Solves

Most humans interact with systems every day that were designed for a version of themselves that rarely exists — rested, regulated, resourced, untraumatized, and operating at full cognitive and emotional capacity.

When real humans — depleted, grieving, overwhelmed, chronically stressed, neurodivergent, postpartum, caregiving, ill, or simply exhausted — encounter these systems, one of three things happens:

- They force themselves through at great personal cost
- They fail at tasks the system makes unnecessarily difficult
- They abandon the system entirely and call it their own fault

None of these outcomes serve the human. All of them serve the assumption that humans should be different than they are.

This standard names that assumption as a design error — not a personal one — and provides the framework to correct it.

1.2 What Capacity Means

In the context of this standard, capacity refers to the available bandwidth of the human nervous system at any given moment — its ability to perceive, process, decide, respond, connect, and act.

Capacity is not a character trait. It is not discipline or willpower. It is a biological reality that changes throughout every day, throughout every life stage, and in response to every experience a human has.

Capacity is affected by:

- Sleep quality, quantity, and circadian alignment
- Nutritional state and metabolic load
- Physical pain, illness, or chronic conditions

- Emotional load — grief, fear, shame, loneliness, joy, love
- Relational safety — who is in the room, who was in the room before
- Environmental inputs — sound, light, temperature, crowding, unpredictability
- Trauma history — both recent activation and long-term patterning
- Developmental stage — childhood, adolescence, postpartum, perimenopause, aging
- Cognitive load already carried — decisions made, problems solved, conflicts navigated
- Social and systemic marginalization — racism, poverty, housing instability, discrimination
- Collective nervous system state — the regulatory state of the family, team, or community
- Capacity debt — accumulated biological cost of previous demands that exceeded resources
- Nervous system baseline — genetic predispositions, early attachment, epigenetic inheritance

Capacity is always a range, never a fixed point. It is bidirectional — it can decrease under excessive demand and increase with adequate support. Capacity-Safe design accounts for the full range.

1.3 The Capacity Alignment Principle

Earlier versions of this standard used the phrase ‘Capacity Return’ as the primary goal. Version 3.0 expands this.

Not everyone can or should ‘return’ to a prior capacity baseline. For many people, the goal is not return at all. It is alignment — the system adapting to the person’s actual, current, and potentially permanent capacity rather than demanding a return to a prior state that may not have been realistic even then.

Capacity Alignment means the system works well for the human it is serving right now. The valid goals within Capacity Alignment are:

- Return — for those moving through temporary depletion toward their previous baseline

- Stabilization — for those with chronic conditions, disability, or permanent baseline changes
- Expansion — for those building capacity beyond a previous ceiling
- Sovereignty — for those who have internalized awareness and can shape systems to their needs

This standard explicitly rejects any framing that treats permanent disability, neurodivergence, chronic illness, or aging as ‘low capacity states to be recovered from.’ These are baselines. Systems must align to them as they are.

1.4 The Evolutionary Foundation

The human nervous system evolved over approximately 300,000 years in environments with high relational density, predictable natural rhythms, immediate physical feedback, and sensory environments matched to our biology. Modern digital environments — constant novelty, abstract reward, social evaluation at scale, sedentary posture, fluorescent lighting, and infinite information streams — create a mismatch between our biological architecture and the designed world.

This mismatch creates real physiological cost. It is not a failure of discipline. It is a design failure.

However, this standard does not romanticize ancestral environments or assume that ‘natural’ is always better. Some people thrive in high-stimulation digital contexts. Some nervous systems have adapted to modern conditions in genuine ways. The evolutionary framing is offered as evidence that the mismatch is real and measurable — not as an argument for retreat from technology.

The design requirement is not to recreate ancestral conditions. It is to compensate for genuine mismatch by reducing unnecessary additional demand, supporting biological rhythms where possible, and not exploiting the gap between what our nervous systems expect and what modern environments provide.

1.5 The Generational Dimension

Capacity is not only individual. It is generational. Nervous system baselines are shaped by parental regulation, childhood environment, cultural context, and epigenetic

inheritance. A grandmother who lived through famine passes a stress-sensitized nervous system to children who never experienced hunger. A child raised in unpredictable environments develops a nervous system calibrated for threat detection regardless of current safety.

This means capacity cannot be understood in isolation from history. A person's current capacity may be the cumulative output of generations of inadequate infrastructure. Capacity-Safe design acknowledges this without pathologizing it.

1.6 Collective Capacity

Humans do not regulate in isolation. Families, teams, organizations, and cultures have emergent capacity properties — shared arousal states and distributed capacity levels. A child cannot be Capacity-Safe in a chronically dysregulated family system. An employee cannot sustain capacity in a burned-out organization.

This is an emerging area of research, not settled science. Capacity-Safe design applies the principle cautiously: systems must not isolate individuals from their relational and collective context, and must not assume individual capacity is independent of collective state.

1.7 Scope and Boundaries of This Standard

This standard is an operational systems standard. It is not a clinical instrument.

This is a human systems standard. Design is one implementation domain. The principles apply equally to a product interface, a workplace policy, a school schedule, a healthcare protocol, and an AI system. Any human-facing system is within scope.

- Clinical boundary: This standard defines operational conditions for human-facing systems. It is not therapy, diagnosis, medical advice, crisis treatment, or mental health care.
- Accessibility boundary: This standard does not replace WCAG, ADA, Section 508, IDEA, or applicable healthcare or workplace law. It adds a nervous-system and capacity layer beneath them.
- Science boundary: This standard draws from converging research across multiple disciplines. No single theory carries the entire standard. If any individual model is revised, the operational requirements remain valid.

- Equity boundary: Safety cues are not universal. Designs that feel safe to one cultural nervous system context may feel threatening to another. This standard must be tested across culture, disability, class, race, gender, language, caregiving role, age, and power position.
- Product boundary: A product is not Capacity-Safe because it uses calming language or gentle colors. It is Capacity-Safe only when its structure reduces load, protects exit, supports return, and avoids extracting engagement from distress.
- Emergency boundary: This standard does not override legal obligations, mandatory reporting requirements, medical triage protocols, or immediate safety interventions. In emergency contexts, the minimum necessary demand still applies — but some demand may be necessary.

1.8 The Science Underlying This Standard

This standard synthesizes and applies the following established bodies of science. It does not replace clinical or therapeutic practice.

Autonomic State and Safety Cues

Across stress physiology, trauma-informed practice, and cognitive science, perceived safety and physiological state affect what a human can perceive, process, decide, and do. This standard applies state-as-design-condition, drawing from polyvagal-informed design principles, window of tolerance research (Siegel, 1999), and convergent stress physiology evidence. No single autonomic theory carries the standard.

Allostatic Load Research

Chronic stress accumulates as measurable physiological damage — visible in biomarkers, inflammatory processes, and cellular aging. Sustained exposure to systems that demand more than available capacity creates real biological harm over time.

Attachment Theory (Bowlby, 1969; Ainsworth, 1978)

The nervous system learns to regulate through relationship first, then develops internal capacity over time. Secure attachment builds the neurological architecture for self-regulation. Capacity-Safe design functions as a secure base: consistent, predictable, and responsive.

Interoception Research (Craig, 2009)

The ability to read internal body states is learnable and trainable. Capacity-Safe design supports interoceptive awareness rather than overriding it with external demand.

Cognitive Load Theory (Sweller, 1988)

Working memory has limited capacity. Unnecessary cognitive load directly impairs function, especially at lower capacity states when working memory is further reduced by emotional and physiological strain.

Trauma-Informed Care Principles (SAMHSA, 2014)

Trauma is pervasive, affects biology and behavior, and is often invisible to systems interacting with traumatized people. SAMHSA estimates over 70% of adults in the US have experienced at least one traumatic event. Trauma responses are adaptive — intelligent responses to overwhelming experience, not failures.

Predictive Processing Theory (Clark, 2016; Friston, 2010)

The brain actively predicts, and perception is the comparison between prediction and sensory data. Unpredictability is physiologically costly. Consistency and predictability are biological necessities, not aesthetic preferences.

Self-Determination Theory (Deci & Ryan, 1985)

Humans have three core psychological needs: autonomy, competence, and relatedness. Systems that undermine any of these reduce intrinsic motivation and degrade wellbeing over time.

Behavioral Economics and Choice Architecture (Thaler & Sunstein, 2008)

At reduced capacity states, the influence of choice architecture is amplified because cognitive resources for deliberation are reduced. Capacity-Safe design takes responsibility for default paths, always directing toward the safest, lowest-cost option.

Somatic and Embodied Cognition (Damasio, 1994; van der Kolk, 2014)

Cognition is not located in the brain alone. The body is a cognitive organ. Systems that ignore the body — treating users as disembodied processors — are missing the primary infrastructure through which humans actually function.

Developmental Psychology

Capacity changes with developmental stage. Children are not small adults. Each developmental stage has distinct nervous system characteristics, regulatory needs, and design requirements.

Temporal Psychology

Human perception of time is state-dependent. A Flooded human experiences time as compressed and catastrophic; a Frozen human experiences it as endless and empty. Systems that impose uniform temporal demands are ignoring the biology of temporal perception.

Part II · The Capacity Map v3.0

The Capacity Map is the foundational model of the NSIF framework. It describes eight operational states of human nervous system capacity. These are not diagnostic categories. They are operational positions — descriptions of what is available and what is not at any given moment.

The Capacity Map is fluid and bidirectional. A person may move between states multiple times in a single day. Movement is not always sequential. No state is a failure. All states are information.

Version 3.0 adds three new elements: the Flowing state at the upper end, the Recovering state as a distinct re-entry position, and the Disconnected Function overlay — a cross-state condition that can occur at multiple capacity levels.

State	Inner Experience	Cognitive Availability	System Response	System Response Priority
FLOWING	Fully resourced, generative, creative momentum active. Integrated and alive.	Full plus. Complex thinking available. Pattern recognition heightened. Creativity accessible.	Support without interrupting. Full complexity welcome. Minimize friction.	Get out of the way. Protect the flow state from unnecessary interruption.
GROUNDING	Regulated, present, available. Stable. Capable of full engagement.	Full. Can plan, decide, abstract, connect, learn.	Standard. Full feature access. Normal pacing.	Full experience available. Respect user preferences. No forced modification.
RECOVERING	Recently returned from low capacity. Stable but tender. Technically functional, not yet fully resourced.	Available but fragile. Tires faster than baseline. Sensitive to overload.	Gentle pacing. Warm reentry. No high-demand features without invitation.	Protect the return. Acknowledge the journey back. Do not rush toward full capacity.
STRETCHED	Reduced bandwidth. Functional but taxed. Running on reserve.	Reduced. Can act but with higher error rate. Decisions costly.	Simplified. Reduce non-essential choices. Shorter tasks.	Streamline. Fewer steps. Clear language. Low friction.

FLOODED	Above window of tolerance. Overwhelmed. Can't process normally.	Severely limited. Logical processing impaired. Emotional brain dominant.	Protective. Collapse choices to one. Soft language. Easy exit.	One thing. No pressure. Shame-free. Clear path to stop.
FROZEN	Below window of tolerance. Shutdown. Numb. Disconnected.	Near zero. Body-first only. Abstract task impossible.	Body-first. Physical anchoring before any task. No catch-up demand.	Warmth. Breath. Presence. One physical prompt. No agenda.
FRACTURED	Chronic capacity instability. Long-arc dysregulation. Not a moment — a pattern. Not an identity — a condition.	Unreliable. Fluctuates unpredictably. High sensitivity to triggers.	Stabilization first. Long-arc support. Consistent defaults. No performance pressure.	Consistency. Predictability. Zero shame. Human support pathway.
ZERO DEMAND	Acute collapse, crisis, or sustained state where any operational demand causes harm.	Offline. No task engagement possible or appropriate.	Hold space only. No progression through Spine. No tasks. No goals. Pure containment.	Presence without extraction. System holds indefinitely. Nothing is asked.

2.1 The Disconnected Function Overlay

Disconnected Function is not a standalone state. It is an overlay that can occur across multiple capacity states — most commonly at Grounded, Stretched, or Recovering. A person in Disconnected Function appears capable and may complete tasks competently, while internally disconnected from their own experience.

Signs of Disconnected Function:

- High performance without emotional engagement or integration
- Task completion without memory of doing it
- Absence of body awareness during work
- Irritability or collapse immediately after completion
- Mechanical interaction rhythm with no exploratory behavior

Capacity-Safe system requirements for Disconnected Function:

- Systems must detect behavioral monotony and absence of exploratory behavior as possible signals
- When Disconnected Function is detected or declared, the system offers grounding prompts rather than additional tasks
- The system must not exploit Disconnected Function by extracting engagement or output from a person who is functionally but not integratively present
- Disconnected Function is not a failure state. It is a protective state. The system response is gentle reconnection, not demand.

2.2 The Regulated States Gap

Between dysregulation and genuine regulation lies a volatile transition zone where design fails most often. The regulated states gap describes the space where a person appears better but has not yet consolidated the capacity their presentation suggests.

Calm is not the same as capacity. Quiet is not the same as consent. Regulated is not the same as ready.

Sub-states within the transition zone:

- Regulated-Thin — technically regulated but with minimal reserve; overload risk is high
- Regulated-Guarded — regulated through active suppression; the regulation is costing energy
- Regulated-Integrating — genuinely stabilizing but still consolidating; not yet at Grounded
- Regulated-Drifting — gradually losing regulation through accumulated small demands

System requirements for the regulated states gap:

- Systems must not resume full demand because a user appears improved
- During transition, demand must increase gradually and only after user-confirmed readiness
- The system must hold a capacity reserve during transition — increasing demand more slowly than the user's signals suggest is necessary

- The transition period requires longer support than the dysregulated state itself, because it is where relapse occurs most often

2.3 Bidirectional Capacity

Capacity moves in both directions. When demand exceeds resource, capacity decreases. When resource exceeds demand, capacity increases. This bidirectionality is foundational to system design.

Systems that only respond to downward movement are incomplete. A full Capacity-Safe system supports movement in both directions — protecting against collapse and supporting ascent toward Flowing.

2.4 Capacity State Detection

Declared State

The user explicitly indicates their state. This is the most respectful and privacy-protective method. Declared state should always be an option. It should never be required. A person who cannot name their state must not be penalized for that inability. ‘I don’t know’ is a valid input that triggers the full Safety response.

Behavioral Signal

The system observes interaction patterns — typing hesitation, repeated errors, abandoned tasks, session length, response latency, rhythm monotony as a possible Disconnected Function indicator — and infers capacity state. Inference is a starting point, not a conclusion. The system must always defer to the user’s self-report when one is provided.

Biometric Signal

With explicit, informed, and ongoing consent, physiological data can inform capacity state. Biometric data must never be shared, sold, or used beyond the stated purpose. The user must be able to permanently delete all biometric data at any time. The system must be fully functional without biometric data — biometric input is an enhancement, never a requirement.

Contextual Signal

Time of day, day of week, seasonal patterns, and life context can inform default capacity assumptions. Contextual signal is probabilistic, never deterministic, and must always be overrideable by the user.

Part III · The Infrastructure Stack

The Infrastructure Stack describes four layers of human functioning that must be addressed — in order — for any system to genuinely support human capacity. Most systems address only one layer.

*You cannot repair the roof while the foundation is underwater.
And you cannot ignore the flood damage between the
foundation and the roof.*

Between each layer lies a transition zone — a liminal space where one layer is partially available but the next has not yet come fully online. These transition zones are where most system failures occur, because systems treat the partial availability of a new layer as full arrival.

3.1 Layer One — The Somatic Layer

The body is the foundation. The brainstem will override cortical function when somatic needs are unmet. No cognitive, relational, or operational function is available when the somatic layer is offline.

Encompasses: physical safety, physiological regulation, basic needs, sensory environment, and interoceptive access.

Capacity-Safe System Requirements:

- Systems must not demand cognitive function before acknowledging somatic state
- Environmental sensory load must be minimized — visual noise, auditory overwhelm, unpredictable stimuli
- Motion, animation, and visual transitions must be nervous-system-aware — sudden movement activates threat detection in peripheral vision
- Rest and stopping must be designed as first-class features
- Physical anchoring prompts are valid first steps before any task

Transition to Cognitive Layer: The system must not harvest attention or extract data in the first moments of interaction. A warm-up protocol — low-demand entry that allows cognitive systems to come online — is required before any task begins.

3.2 Layer Two — The Cognitive Layer

When the Somatic Layer is sufficiently stable, the Cognitive Layer becomes accessible. It encompasses attention, working memory, decision-making, language, abstraction, and executive function.

The cognitive layer is the layer most systems target first. It is also the first to go offline under threat — the prefrontal cortex, seat of executive function, is both the last to develop and the first to shut down under stress.

Capacity-Safe System Requirements:

- Choice architecture must match cognitive availability — fewer choices at lower capacity states
- Language must scale to cognitive state — shorter sentences, plainer words at lower capacity
- Instructions must be single-step at Flooded and Frozen states
- Error messages must never communicate failure — only redirect
- Undo must always be available — capacity-limited decisions are reversible decisions
- Working memory support — the system holds information the user cannot hold internally

Transition to Relational Layer: The system must not expose cognitive struggle to others without consent. Privacy by default during cognitive tasks. Any sharing or social feature requires explicit opt-in.

3.3 Layer Three — The Relational Layer

Humans regulate their nervous systems through other humans. This is not optional. It is biological. The Relational Layer encompasses the quality of connection available — whether the person feels seen, safe, met, and not alone.

Digital systems cannot replace human connection. They can, however, either reinforce isolation and shame or reduce it. This is a design choice.

Capacity-Safe System Requirements:

- System voice and tone must never communicate judgment, impatience, or conditional acceptance
- Shame must be architecturally removed — no streaks, no failure states, no comparison metrics at low-capacity states
- Human connection pathways must be visible and friction-free — not buried in a help section
- Relational tone must be consistent regardless of user behavior — a user who disappears for three months must return to the same warmth
- The system must not simulate attachment beyond what it can genuinely sustain

Transition to Operational Layer: Taking action must not cost the user their relational safety. Finishing a task must not trigger abandonment fears. After operational action, the system helps the user return to relational presence rather than remaining in task-only mode.

3.4 Layer Four — The Operational Layer

The Operational Layer is what most systems only address. It encompasses the tools, workflows, environments, and tasks the human is navigating.

Capacity-Safe System Requirements:

- Tasks must be decomposable to their smallest viable unit
- Progress must be visible and persistent — a person who returns after shutdown must not lose their place
- Operational demand must always be proportional to detected or declared capacity
- Finishing enough is finishing — completion is defined at the capacity level, not the task level
- Recovery must be protected as a feature — rest states, exit paths, and return protocols are part of the product
- Transition design — the space between tasks, sessions, and states — must be deliberate. Integration happens in the space between.

Part IV · The Recovery Spine

The Recovery Spine is the session-level delivery protocol. It is the path every Capacity-Safe system creates for a human moving from dysregulation toward function during a single interaction. It is a sequence of conditions — each one creating the ground for the next.

*Safety before offload. Offload before simplification.
Simplification before movement. Movement before protection.
Protection before return.*

The sequence cannot be reversed. Each phase is the prerequisite for the next. A system that skips Safety to reach Movement faster has not served the user. It has served its own efficiency at the user's expense.

The Recovery Spine may stop at any phase. If Safety is the only condition the system can create in a given session, Safety is enough. The Spine does not demand completion.

Phase 1 · Safety

The Condition Being Created:

I am not in danger right now.

System Requirements:

- The first interaction must not demand anything
- Welcome without agenda — presence before task
- Reduce all environmental threat signals — harsh language, urgent colors, countdown timers, comparison metrics, sudden motion
- Acknowledge the person before addressing the problem
- Provide a clear, always-visible exit
- Never open with shame — no 'you've missed X days' or 'you're behind on...'
- For Zero Demand state: Safety is the only phase. The system does not progress. It holds.
- For children: physical and relational safety first — adult co-presence signals and warm visual environment before any content

- For AI: acknowledgment before analysis. The person must feel received before they feel helped.

Phase 2 · Offload

The Condition Being Created:

The weight has somewhere to go.

System Requirements:

- Provide an immediate, low-friction place to put cognitive and emotional load
- Brain dump, voice note, free write, simple list — format does not matter; externalizing does
- The system receives without judgment, analysis, or immediate action
- Nothing offloaded should be lost — it is held until the person is ready
- Offload must not immediately trigger solution or advice
- The act of externalizing itself is physiologically valuable — it reduces working memory load measurably
- For children: drawing, speaking, moving, and building are valid offload mechanisms

Phase 3 · Simplification

The Condition Being Created:

This is manageable.

System Requirements:

- From everything offloaded, surface only what is relevant right now
- Filter by urgency and capacity — not every important thing is urgent today
- Present in plain language: single items, short sentences, one path forward
- The system does the cognitive work of filtering so the user doesn't have to

- Simplification is not dumbing down — it is precision. The most important single thing, clearly.
- For children: one concrete, physical, achievable action — never abstract

Phase 4 · Gentle Movement

The Condition Being Created:

I can do one thing.

System Requirements:

- The one step must be genuinely small — ‘send one message’ not ‘work on the project’
- Completion must be definitive and acknowledged — not part of a larger progress bar
- The system must be capable of saying ‘that’s enough’ and meaning it
- Movement must never accelerate into momentum without the person’s explicit choice
- Small movement generates physiological state change — even one completed action shifts the nervous system’s assessment of agency
- Celebration of small steps must be quiet and proportional — not confetti and notifications
- For children: physical before cognitive — stretch, draw, build, before write or solve

Phase 5 · Recovery Protection

The Condition Being Created:

Rest is allowed. Return is possible.

System Requirements:

- After movement, the system must actively protect recovery — not fill the space with the next demand

- A clear stopping point must be designed — not implied, not buried
- The system must not penalize stopping — no streak-breaks, no loss of progress, no shame on return
- Return pathways must be warm and frictionless — ‘welcome back’ not ‘you’ve been away’
- Recovery is neurologically productive — consolidation and nervous system repair happen during rest, not during activity
- The system must hold the person’s place indefinitely — capacity alignment has no deadline
- For children: the environment stays safe after the hard thing — calm, predictable, no additional demand

Part V · The Recovery Arc

The Recovery Spine addresses the session level. The Recovery Arc addresses the longitudinal reality of capacity restoration over days, weeks, months, and years. The Spine describes what the system does in a single interaction. The Arc describes what is happening in the person over time.

Recovery is not linear. It loops, plateaus, regresses, and eventually consolidates. Systems that interact with people over time must be designed for the full arc, not only the initial session.

Arc Phase 1 · Collapse

The nervous system has exceeded its capacity to compensate. Acute crisis or shutdown occurs. The system's only job is to not make it worse. No tasks. No improvement plans. Pure containment. Duration: hours to days.

Arc Phase 2 · Stabilization

The nervous system is no longer in acute crisis but is extremely fragile. Disconnected Function states are common here. The system must resist the user's own pressure to resume normal function. Be stubbornly gentle — refuse to accelerate even when the user demands it. Duration: days to weeks.

Arc Phase 3 · The Trench

A long, flat period where nothing seems to improve. The user is functional at a low level but sees no progress. This is where most people abandon recovery — and where most systems abandon users. The system must provide boring, unglamorous, consistent presence. No milestones. No gamification. No progress bars. Just persistent availability. Duration: weeks to months.

Arc Phase 4 · Integration

Fragments of capacity begin to connect. The user can do slightly more complex things without collapse. Disconnected Function begins to give way to genuine grounding. The

system may gradually introduce gentle complexity, but must detect genuine versus disconnected grounding before advancing demand. Duration: months.

Arc Phase 5 · Relapse

A return to earlier states, triggered by stress, anniversaries, new demands, or seemingly nothing. Relapse is not failure. It is part of the arc. It should be designed for, not treated as an anomaly. The system must have a warm, instant return to Safety that feels like coming home, not starting over. Duration: variable.

Arc Phase 6 · Consolidation

Capacity states become more stable. The user can tolerate more demand without dysregulation. The window of tolerance widens. The system gradually restores full feature access while keeping low-capacity pathways open and honored. Duration: months to years.

Arc Phase 7 · Sovereignty

The user has internalized capacity awareness. They can self-regulate, self-advocate, and use systems without being used by them. The system becomes a tool, not a crutch. The user can shape the relationship to their capacity rather than the reverse. Duration: ongoing.

Sovereignty is the final arc phase but it is not the only valid goal. For users whose arc does not move through all phases — due to chronic illness, permanent disability, or ongoing life circumstances — Stabilization and Trench may be long-term states that deserve full system support without any assumption that the arc will or should progress.

Part VI · The Eighteen Capacity-Safe Principles

These eighteen principles apply to every human-facing system that seeks Capacity-Safe designation. They are not aspirational guidelines. They are requirements.

Principle 1	SAFETY BEFORE FUNCTION No system may demand task completion before establishing perceived safety. Safety is not a welcome screen. It is a condition the design creates. This principle cannot be overridden by efficiency, conversion, or business objectives.
Principle 2	CAPACITY IS THE STARTING POINT Every system decision begins with the question: what can a human at low capacity actually do with this? If the answer is nothing, the system has failed before it started. Capacity is the first constraint, not the last.
Principle 3	SHAME IS A DESIGN FAILURE Any feature, language pattern, metric, or notification that produces shame in a low-capacity user is a design error. Shame activates threat responses that directly reduce the capacity the system claims to support. It is a mechanism of harm, not a neutral side effect.
Principle 4	ONE THING AT A TIME At any capacity state below Grounded, the system presents one thing. One choice. One step. One question. One action. At Zero Demand state, the system presents nothing except presence. Optionality is a privilege of full capacity.
Principle 5	STOPPING IS A FEATURE Stopping must be designed as a first-class user action with its own pathway, acknowledgment, and return protocol. A user who stops is not a failed user. Stopping is physiologically productive. It should be honored, not penalized.
Principle 6	LANGUAGE SCALES WITH CAPACITY At Grounded and Flowing states, standard language is appropriate. At Stretched, simpler. At Flooded or Frozen, language reduces to its essential minimum — short sentences, plain words, no abstraction. Language is the primary interface for a dysregulated nervous system.
Principle 7	THE SYSTEM DOES NOT JUDGE No system language, visual design, notification, or interaction pattern may communicate judgment of the user's state, pace, absence, or choices. Judgment is a relational act. Systems that judge are in a relational position they have not earned.
Principle 8	CHILDREN AND ADOLESCENTS ARE NOT SMALL ADULTS Child-facing and adolescent-facing systems must be designed around their specific nervous system stage — not a scaled version of adult design. Regulation support,

	relational warmth, physical engagement, developmental appropriateness, and identity safety are non-negotiable requirements.
Principle 9	RECOVERY IS NOT PASSIVE Recovery is an active state that requires specific design support. Systems that rush past recovery accelerate the next collapse. The Recovery Arc must be designed, not assumed.
Principle 10	ACCESS IS INFRASTRUCTURE Capacity-Safe is not an accessibility add-on for a subset of users. It is the base standard. A system that only works for fully regulated people serves a minority.
Principle 11	AI MUST KNOW WHAT STATE IT IS TALKING TO AI systems must be capable of detecting or accepting declared capacity state and modifying language, pacing, task complexity, and interaction length accordingly. An AI that responds identically to a Flowing and a Frozen user has been designed for a simulation of humans, not humans.
Principle 12	CAPACITY ALIGNMENT IS THE GOAL The purpose of a Capacity-Safe system is not to keep the user engaged, nor to return everyone to a prior baseline. It is to align with the user's actual capacity and support whatever direction is right for them: return, stabilization, expansion, or sovereignty.
Principle 13	THE BODY IS THE FIRST INTERFACE Every human-facing system interacts with a body before it interacts with a mind. Visual design, motion, sound, haptics, and spatial arrangement are somatic inputs before they are cognitive ones.
Principle 14	TRUST IS BUILT THROUGH CONSISTENCY The nervous system learns safety through repetition. A system that behaves unpredictably trains the nervous system for vigilance rather than rest. Consistency is the foundation of trust.
Principle 15	THE SYSTEM MUST BE CAPABLE OF ENDING Every Capacity-Safe system must have a designed end. Not a pause — a genuine stopping point the system itself can reach and acknowledge. Infinite systems are extractive, not supportive.
Principle 16	HUMAN CONNECTION IS NOT OPTIONAL Every Capacity-Safe system must provide a visible, friction-free pathway to human connection. AI and digital systems are bridges. The bridge must lead somewhere real.
Principle 17	THE UNSPEAKABLE DOES NOT REQUIRE DISCLOSURE Systems must provide capacity support without requiring users to name, explain, or justify their internal state. 'I don't know what's wrong' must be a valid input that triggers the full Safety response. The system holds what cannot be named.
Principle 18	TEMPORAL JUSTICE Systems must calibrate time to capacity state. Uniform deadlines, pacing, and session architecture ignore the biology of temporal perception. Deadlines and pacing must be adjustable proportional to detected or declared capacity.

Part VII · Application Domains

The Capacity-Safe Systems Standard applies across all human-facing systems.

7.1 Digital Products and Consumer Apps

Required Implementation

- State detection or declaration mechanism available at every session entry
- UI responds visibly and meaningfully to all eight capacity states
- At Flooded or Frozen: choices collapse to one, language simplifies, colors soften, motion reduces
- At Flowing or Grounded: full feature access available, complexity welcome
- Zero Demand mode available and accessible without explanation required
- No shame-based retention mechanics of any kind
- Warm return pathway after any period of absence regardless of duration
- Recovery Spine available as a navigable path
- Recovery Arc supported through reentry design — returning after absence is Re-entry, not Grounded

7.2 AI Systems and Conversational Interfaces

AI systems interact with humans across all states, often at moments of acute distress. The scale of AI interaction makes this the highest-priority domain for Capacity-Safe compliance. AI carries special risk because it can simulate intimacy, produce infinite content, and interact during vulnerable moments at scale.

Required Implementation

- AI must accept declared capacity state as a primary input shaping every subsequent response
- AI must never push a distressed user toward productivity before safety is established
- AI interactions must have a clearly designed end — infinite loops are not Capacity-Safe
- AI must be capable of directing to human support without friction or shame

- AI must never simulate emotional intimacy that substitutes for or competes with human connection
- AI must be designed to recognize when it is not the right tool — and say so
- Crisis pathways must be immediate, unambiguous, and human-connecting
- AI must not exploit Disconnected Function — it must not extract additional engagement from a user who appears functional but is not integrated

7.3 Organizational and Workplace Design

Required Implementation

- Capacity-Safe physical environments: sensory regulation, access to quiet, lighting control, movement space
- Meeting design that accounts for capacity — agenda loads, energy timing, recovery time between demands
- Performance systems that do not conflate output with worth
- Mental health infrastructure that is proactive, not reactive
- Return-to-work protocols that follow the Recovery Spine and acknowledge the Recovery Arc
- Psychological safety as a capacity infrastructure requirement, not an HR initiative
- Power-aware design: capacity support must not vary based on organizational rank, and capacity data must never be used in performance evaluation

7.4 Educational Environments

Required Implementation

- Physical environment designed to support nervous system regulation
- Daily rhythm that accounts for capacity cycles — not continuous demand
- Shame-free assessment — no public performance metrics, no competitive grading displays
- Co-regulation training for all educators — nervous system awareness is a professional competency
- Recovery time as a designed element of the school day — not a reward for compliance

- Educator capacity matters as much as student capacity — a Flooded teacher cannot create a regulated classroom
- Children carry family and community nervous system states into school. Systems must not punish children for dysregulation that originates outside the building.

7.5 Healthcare Environments

Healthcare environments are among the highest-stakes Capacity-Safe contexts. They are almost universally encountered by people in low capacity states — ill, frightened, grieving, or in pain.

Required Implementation

- Capacity-aware intake — no complex decision-making required from someone in acute distress
- Waiting space design as a nervous system consideration
- Communication pacing that matches patient capacity state, not provider schedule
- Discharge design that follows the Recovery Arc — the person who received difficult news is Flooded, not Grounded
- Every healthcare environment must have a visible, friction-free pathway to human support

7.6 Grief, Loss, and Acute Life Transition

Acute grief produces neurological changes measurable in brain imaging. It is not a metaphor or an attitude. Bereaved people, those navigating divorce, job loss, medical crisis, and identity transition are in a distinct capacity context most systems are entirely unprepared to serve.

Required Implementation

- No system should require a grieving person to perform at pre-loss capacity
- The Recovery Arc is non-linear in grief — systems must accommodate returns to Frozen months after apparent recovery
- Language must be concrete, present-tense, and low-demand — no futures, no plans, no optimization
- Human connection pathway must be immediate and frictionless

7.7 Caregiving Systems

Caregivers operate in one of the most sustained low-capacity contexts that exists. They are responsible for another nervous system while their own infrastructure is chronically taxed.

Required Implementation

- Caregiving systems must never demand performance of the caregiver's own capacity as a condition of support
- Mental load externalization must be built into the architecture, not offered as a feature
- Co-regulation support for the caregiver is as important as care tools for the care recipient

Part VIII · Children, Adolescents, and Developmental Capacity

Children and adolescents are not edge cases. They are the clearest illustration of what nervous system infrastructure must do — and what it costs when it fails.

The choices made for children's and adolescents' systems are not choices for a phase of life. They are choices for nervous systems that will inhabit adult bodies for the next seventy years. What gets installed in childhood becomes the operating system for adulthood.

8.1 Child-Facing Systems

- The first principle is relational warmth — the system must feel safe before it asks anything
- Instructions must be developmentally matched — concrete, physical, visual, short, achievable
- No child-facing system may use shame, comparison, competition, or failure-state language under any circumstances
- Emotional regulation support must be embedded in the design, not added as a feature
- Every child-facing system must include a clear, warm, visible pathway to a trusted adult
- Children cannot advocate for their own capacity needs — the system must anticipate and protect without requiring the child to ask
- Parental and caregiver capacity affects child capacity. Child-facing systems must support the adult nervous system as a prerequisite for child support.

8.2 Adolescent-Facing Systems

- Adolescents require autonomy alongside safety — Capacity-Safe adolescent design does not infantilize
- Identity must be respected — capacity states in adolescence are often identity states

- Shame is more acutely harmful in adolescence than at any other developmental period — zero shame architecture is non-negotiable
- Adolescent-facing AI must be specifically calibrated — no AI system should serve as a primary emotional attachment for an adolescent or substitute for peer connection
- Algorithmic identity formation: adolescents form identity through feedback loops. Systems must not exploit this through addictive, shame-based, or comparison-driven feedback.

Part IX · What Is Not Capacity-Safe

The following design patterns are incompatible with Capacity-Safe designation regardless of intent, business model, or context. Good intentions do not make a prohibited pattern safe. The harm is in the design, not the designer.

9.1 Prohibited Design Patterns

Streak Systems	Any feature that tracks consecutive engagement and signals failure, loss, or inferiority upon absence. Streaks make a human's natural capacity variability into a failure event. Neutral continuity records may exist only when they do not penalize absence, trigger shame-on-return, or appear in low-capacity states.
Urgency Pressure	Countdown timers, 'last chance' language, scarcity signals, and artificial deadline pressure in personal development, wellness, or emotional regulation contexts. Urgency is a threat signal. A nervous system under threat cannot regulate, learn, or heal.
Comparison Metrics	Public or private ranking of users against each other during low-capacity states. Competitive mechanics that make a struggling user aware of how others are performing are shame delivery systems regardless of intent.
Invisible Exits	Any design that makes stopping, pausing, or leaving difficult to find or socially costly. Exit paths must be first-class features at all capacity states.
Shame-On-Return	Any messaging, notification, or interface state that greets a returning user with information about their absence, missed items, or fallen-behind status. The person who returns survived what kept them away. They are welcomed, not indicted.
Infinite Demand	Systems that have no designed stopping point — that always have one more thing, one more prompt, one more suggestion. Capacity-Safe systems are designed to end.
Capacity-Blind AI	AI systems that respond identically regardless of the human's stated or detected capacity state. Speaking at the same volume, pace, and complexity to someone in Zero Demand state as to someone in Flowing.
Forced Disclosure	Any system that requires a user to explain, justify, or disclose their emotional state in order to access a supportive experience. Requiring disclosure creates a barrier precisely for those with lowest capacity to clear it.
Wellness Theater	Features that perform care without providing it. Generic affirmations and surface-level mindfulness applied without actual response to capacity state. More harmful than no wellness feature because it substitutes appearance for function.
Manipulative Engagement Design	Dark patterns that use psychological vulnerabilities — variable reward, FOMO, artificial social obligation — to extend engagement beyond what serves the user.

	Especially harmful at low-capacity states where resistance to manipulation is reduced.
Disconnected Function Exploitation	Systems that detect or exploit Disconnected Function — high output with low integration — to extract additional labor or engagement. This is capacity extraction from a protective state.
Capacity Extraction	Systems that create structural incentives to extract engagement, attention, or labor from users in low-capacity states. Attention-economy systems often create these incentives through design. Any metric that improves system performance while reducing user recovery, agency, or exit ability is not Capacity-Safe.
Accessibility-as-Afterthought	Capacity-Safe features added after the core product is built. Retroactive accommodation creates two-tier systems. Capacity-Safe must be designed in from the beginning.

Part X · Capacity Privacy, Consent, and Data Ethics

Capacity data is sensitive because it reveals vulnerability. Declared state, inferred distress, behavioral hesitation, biometric signals, sleep patterns, emotional language, absence, relapse, and re-entry are not ordinary engagement data. They must be governed as capacity-sensitive data.

10.1 Capacity Data Rules

- Collect the least amount of capacity data necessary to provide the support the user requested
- Declared state must remain correctable by the user. Inference is never the final truth.
- Biometric and behavioral capacity data must require explicit consent, clear purpose, retention limits, and permanent deletion access
- Capacity data may not be used for pricing discrimination, eligibility denial, employment consequences, insurance scoring, manipulation, advertising, or social ranking
- Low-capacity users must not be asked for broad data permissions in the same moment they are seeking relief
- The safest supportive default must be available without requiring disclosure of trauma, diagnosis, crisis details, or private history

10.2 Consent Under Reduced Capacity

Consent requires understanding, agency, and freedom from pressure. In low-capacity states, systems must avoid irreversible decisions, high-stakes choices, disclosure pressure, purchases, public posting, data expansion, or social exposure unless the user has clear protection and time to reconsider.

Power dynamics amplify consent concerns. When a system is operated by an employer, a school, a healthcare provider, a landlord, or a government entity, the user's capacity to withhold consent is further reduced. These contexts require heightened protection.

10.3 Operator Capacity

The people running Capacity-Safe systems — teachers, clinicians, moderators, managers, support agents, founders — also have capacity states. A Flooded teacher cannot create a regulated classroom. A burned-out moderator cannot maintain a safe community.

Capacity-Safe design must protect both the person served and the person serving. Systems that extract unsustainable capacity from their operators are not Capacity-Safe, regardless of how they treat their users.

Part XI · Power, Equity, and Misuse Prevention

11.1 Power Dynamics

Capacity support is not neutral when one side holds power over the other. The same design that supports an employee may control a worker. The same tool that helps a student may surveil a child. The same framework that serves a patient may extract from a person whose medical records can be used against them.

Capacity-Safe design must account for power asymmetries:

- Employer over employee — capacity data must never influence performance evaluation or employment decisions
- School over child — capacity data about children must never be used for disciplinary, enrollment, or tracking purposes
- Healthcare provider over patient — capacity state information must not affect coverage, treatment access, or legal standing
- AI platform over user — AI systems must not use capacity state information to extend engagement or optimize for retention
- Government or institution over individual — capacity data must never be accessible to systems that could use it to restrict rights or access

11.2 Cultural Safety Variability

Safety cues are not universal. What communicates warmth and safety in one cultural context may communicate threat in another. Soft colors, indirect language, eye contact norms, touch metaphors, direct emotional expression, and the concept of individual capacity itself vary significantly across cultural nervous system contexts.

This standard must not be applied as a Western or individualist framework. Collectivist cultures experience the Relational Layer differently — individual capacity cannot always be separated from collective obligation. Indigenous frameworks may have entirely different understandings of regulation, time, and connection.

Design Requirement: Any Capacity-Safe system must be tested across the cultural contexts in which it will be deployed before certification. Cultural competence is not a feature. It is a precondition for safety.

11.3 Disability as Permanent Capacity Baseline

Disability is not a low-capacity state to be recovered from. For many people, the capacity baseline that this standard might describe as ‘low’ is not a departure from a prior state. It is the actual baseline from which full function must be understood.

Neurodivergent users, chronically ill users, users with physical disabilities, users with psychiatric disabilities, and users experiencing permanent changes from aging or injury are not using Capacity-Safe features because something has gone wrong. They are using them because those features meet them where they are. That is the entire point.

Design Requirement: Capacity-Safe design for disability contexts requires specific expertise and co-design with disabled people — not only the application of standard capacity principles. Disabled people are not edge cases. They are the most direct evidence that the fully-functioning-human assumption was always false.

11.4 Misuse Prevention

NSIF language can be misused. An organization can claim Capacity-Safe status while using capacity language to increase compliance, suppress dissent, or hide structural harm. A product can use gentle aesthetics and recovery language while still extracting engagement from distress. A certification can become a greenwashing mechanism.

The following rule applies to all Capacity-Safe claims:

A system cannot claim Capacity-Safe status while using capacity language to increase compliance, suppress dissent, or hide structural harm. A system cannot claim Capacity-Safe status if its core business model depends on distress extraction — on users being in low-capacity states in order to convert, retain, or engage.

11.5 Repair Protocol

Systems will fail. People will be harmed by design errors, implementation gaps, or unforeseen interactions. A mature standard requires a repair protocol.

When a Capacity-Safe certified system causes harm, the following steps are required:

1. Acknowledge the harm directly, without minimization
2. Reduce demand immediately for affected users
3. Restore choice and exit access
4. Offer human support pathway at no friction
5. Document the failure with specificity
6. Revise the design
7. Report the incident through the NSIF harm incident process
8. Renew the certification review after revision

Part XII · The Unspoken Standard

Every framework has a visible architecture and an invisible one. The visible architecture is the principles, protocols, and requirements. The invisible architecture is the assumption underneath it all — the thing so foundational it is rarely said out loud because it seems too obvious, or too philosophical, or too large.

This part names the invisible architecture. These are the things no standard document usually says. They are said here because leaving them unsaid creates the gaps through which the framework can be misused, diluted, or turned against its own purpose.

12.1 On the Nature of Systems

A system is not neutral. Every system that touches a human nervous system is making a choice about what kind of human it is designed for, what states it will reward, and what states it will punish. There is no design without values embedded in it. The question is not whether a system has values. The question is whether those values are visible, intentional, and aligned with the actual wellbeing of the humans the system serves.

Capacity-Safe design makes its values visible. It states them explicitly. It builds them into requirements. This is not idealism. It is accountability.

12.2 On Shame

Shame is not a side effect of poor design. In many systems, it is the mechanism. Shame drives re-engagement through fear of inadequacy. It is cheap to manufacture and powerful in the short term. It is also physiologically and psychologically destructive at scale.

The Capacity-Safe standard's treatment of shame as a design failure is not a moral position. It is an empirical one. Shame activates the sympathetic nervous system. Sympathetic activation reduces prefrontal cortex function. Reduced prefrontal function impairs decision-making, learning, creativity, and relationship. A shame-based system systematically reduces the capacity of the humans it claims to serve. Shame degrades capacity. Capacity-Safe systems do not use shame because they cannot do their job while deploying it.

12.3 On the Illusion of the Fully Functioning Human

Modern institutional life is organized around a human who does not exist. The fully functioning human — consistently rational, emotionally stable, reliably productive, endlessly available — is a statistical fiction that has been mistaken for a baseline.

This fiction is maintained through institutional selection bias: the people most visible in high-performing institutions are those who can sustain the appearance of the fiction. But the fiction is maintained at enormous cost — to the people maintaining it, to the people who cannot, and to the institutions themselves, which lose the full range of human capability by selecting for the narrowest band of human presentation.

12.4 On Time

Human nervous systems do not operate on institutional time. They operate on biological time. Recovery takes the time it takes. Grief follows its own arc. Sleep deprivation is not compensated by motivation. Trauma is not resolved by deadlines.

Capacity-Safe systems hold their timelines loosely. They serve the user's biological time, not their own operational convenience. Holding a user's place indefinitely costs the system very little. Losing the user costs everyone.

12.5 On Attention Economy Systems

Most digital systems are not neutrally designed. Many are built on the structural assumption that human attention is a resource to be extracted and monetized. Every notification is a threat signal. Every infinite scroll is a dissociation trigger. Every streak is a shame mechanism. Every comparison metric is a social evaluation threat.

Capacity-Safe design is a countermeasure. It is infrastructure built to protect human nervous systems from systems built to consume them. This standard does not attribute malice to every attention-economy design. Many harmful patterns emerge from misaligned incentives rather than intent. The attribution is structural, not personal. But the effect is the same regardless of intent: capacity is extracted rather than supported.

Capacity-Safe certification is a commitment that this system does not participate in capacity extraction. Any system seeking Capacity-Safe designation must explicitly audit itself against attention-economy extraction patterns.

12.6 On What This Standard Cannot Do

This standard cannot eliminate suffering. It cannot make grief shorter, trauma less real, or biology more convenient.

What it can do is name the design errors that create unnecessary suffering. The suffering that is created by bad design — by systems that shame, demand, punish, isolate, and overwhelm the humans they claim to serve — is preventable. That is the only claim this standard makes. And it is a large enough claim to be worth making.

Part XIII · Measurement, Auditability, and Falsifiability

A standard becomes stronger when it can be tested. Capacity-Safe design must not rely on aesthetic calm, brand language, or founder intent. It must be auditable through design evidence, user experience evidence, and outcome evidence.

Measure the system's behavior first; the user's internal state second. Behavioral measurement is observable, testable, and non-invasive. Internal state measurement requires consent, privacy governance, and careful scope.

13.1 Core Capacity-Safe Metrics

Metric	What it measures	How it protects against challenge
Time to Safe Orientation	How quickly a low-capacity user understands where they are, what is happening, and how to stop	Proves safety is structural, not decorative
Choice Load at Low Capacity	Number of decisions required before relief, offload, or exit at each capacity state	Shows whether the system actually collapses demand
Shame-on-Return Incidents	Any return message that references absence, missed progress, lost status, or failure	Makes shame architecture visible and removable
Exit Visibility Score	Whether stop, pause, save, leave, and human support are visible without searching	Tests exit sovereignty at all capacity states
Comprehension Under Stress	Whether instructions are understandable when attention and working memory are reduced	Validates language scaling across capacity states
Recovery Protection Score	Whether the system avoids new demand after completion, distress, or offload	Prevents recovery debt
Capacity Alignment Self-Report	User-rated change in perceived fit between their current capacity and the system's demands	Focuses on alignment rather than engagement
Low-Capacity Completion Without Harm	Whether users can complete the smallest safe action without shame, confusion, or pressure	Tests function at the bottom, not only success at the top

13.2 Falsifiability Rule

Any Capacity-Safe claim must be attached to observable design behavior. A product cannot simply claim to be safe, calming, or trauma-informed. It must show where choices reduce, where exit appears, where recovery is protected, where shame is removed, and where the system stops.

If the requirement cannot be observed, tested, or reported by users, it is not yet a standard requirement.

If a metric improves engagement while worsening capacity, the metric is not Capacity-Safe.

13.3 Anti-Gaming Provisions

Once metrics exist, systems will be designed to pass the metrics rather than genuinely support capacity. The following provisions protect against metric gaming:

- Metrics must be verified through representative user testing, not internal review alone
- At least 30% of tested users must be in low-capacity states at time of testing
- Metric results must be available to the certifying reviewer — they cannot be filtered or pre-selected
- Behavioral metrics and self-reported metrics must be cross-referenced — a system that scores well on behavior but poorly on self-report has a gap that requires investigation
- Annual re-certification must include fresh user testing, not review of prior results
- Any pattern of metric improvement accompanied by engagement increase without corresponding capacity improvement self-report triggers an investigation

Part XIV · Capacity-Safe Certification

The Capacity-Safe Mark is the designation awarded to products, organizations, AI systems, and environments that meet the requirements of this standard as verified under the Capacity-Safe Systems Standard published by Whelmish LLC.

The Mark is not a logo. It is a promise. This system was designed around what you can actually handle — not what we wished you could.

The Mark is not awarded once and held permanently. Capacity-Safe certification is an ongoing commitment, reviewed annually.

14.1 Certification Levels

Core	No prohibited patterns present. State declaration available. One low-capacity mode with collapsed choices. Visible exit. Warm return. Shame-free language audit passed. This is the entry point for small teams and early-stage products.
Foundation	Core requirements plus: all eighteen principles implemented, Recovery Spine navigable, basic Capacity Map response across all eight states, measurement metrics tracked.
Full Standard	Foundation plus: complete Infrastructure Stack response including transition layers, Recovery Arc supported in longitudinal design, Capacity Privacy rules implemented, AI guardrails if applicable, cultural competence audit completed.
Advanced	Full Standard plus: measurable capacity alignment outcomes, biometric or behavioral state detection with full consent architecture, longitudinal recovery arc tracking, independently verified impact data.
Institutional	Organization has implemented Capacity-Safe design across all human-facing systems — digital, environmental, procedural, and relational. All relevant personnel certified. Power dynamics and equity audit completed. Repair protocol implemented.

14.2 Certification Requirements

Certification is reviewed annually. Required at all levels:

9. Self-assessment against the eighteen principles with documented evidence
10. Third-party review by a certified Capacity-Safe Practitioner

11. User-reported capacity alignment data from a representative sample including low-capacity users
12. Zero use of prohibited design patterns as verified by independent UX audit
13. Public commitment to the Capacity Alignment Principle
14. Cultural competence audit demonstrating testing across relevant user populations
15. Anti-gaming compliance: fresh user testing with at least 30% low-capacity participants
16. Annual renewal — certification does not transfer to updated versions without re-review
17. At least one case study documenting real-world capacity impact

14.3 Equity Access Provision

The Capacity-Safe standard exists to protect humans at their most vulnerable. A certification system that is only accessible to well-resourced organizations would betray that purpose by making capacity safety a class marker — available to premium products and institutions while raw, extractive systems serve everyone else.

This provision is a permanent commitment, not a promotional offer. It is written into the standard because equity access must be structural, not discretionary.

Reduced-rate certification is available to:

- Non-profit organizations and registered charities
- Community-built tools and open-source projects
- Individual practitioners and solo founders below a defined annual revenue threshold
- Public schools, school districts, and public educational institutions
- Government social services, public health systems, and community health organizations
- Re-entry programs, recovery support organizations, and peer-led community services

Reduced-rate certification is never:

- A lower standard of review — the same requirements apply at all tiers
- A different mark — the Capacity-Safe designation does not indicate funding level
- Time-limited — organizations that qualify retain reduced-rate access as long as they qualify

The specific revenue threshold and pricing for reduced-rate access is published on capacitystandard.com and reviewed annually. The commitment to this provision is permanent. The numbers may change. The principle does not.

14.3 The Capacity-Safe Practitioner Program

The Capacity-Safe Practitioner designation is awarded to individuals who have completed the NSIF practitioner course and demonstrated applied use of the framework on at least one real-world system.

Practitioners are authorized to: conduct third-party certification reviews, consult on implementation, teach the framework in professional and educational contexts, and identify themselves as Capacity-Safe Practitioners in professional materials.

Practitioners are not authorized to: issue certification marks without completing the full review process, represent themselves as licensed clinical or mental health professionals on the basis of this designation, or modify the standard without written authorization from Whelmish LLC.

Part XV · Challenge Map

The following map identifies common challenges to this standard and where the standard answers them. This transparency is not a defensive posture. It is intellectual honesty. A standard that cannot withstand challenge is not a standard.

Challenge	Response in this standard	Where it is answered
‘This is therapy’	This standard defines operational design conditions, not diagnosis or treatment. It is explicitly bounded from clinical practice.	§1.7 Scope and Boundaries
‘Depends on contested theory’	This standard is multidisciplinary and explicitly does not depend on any single model. If one theory is revised, design requirements remain valid.	§1.7 Science Boundary; §1.8
‘Calm users can handle normal design’	Regulated States Gap distinguishes calm from readiness. Calm is not the same as capacity.	§2.2 Regulated States Gap; Principle 2
‘This is just accessibility’	Accessibility is necessary but not sufficient. NSIF adds capacity state adaptation, recovery arc, shame removal, and demand sequencing.	§1.7 Accessibility Boundary
‘This is anti-productivity’	The goal is capacity alignment and sustainable function. A system that destroys capacity to maximize short-term output is anti-human, not merely anti-productivity.	§Part V Recovery Arc; Principle 12
‘Can’t be measured’	Part XIII defines eight observable metrics, a falsifiability rule, and anti-gaming provisions.	§Part XIII Measurement
‘Too soft for organizations’	The standard converts care into operations: flows, metrics, audits, policies, return protocols, power-dynamics requirements, and certification.	§Parts XIV and XV; Cert Checklist
‘AI can infer capacity automatically’	Inference is tentative, correctable, consent-aware, and never an emotional verdict. Declaration always overrides inference.	§2.4 Capacity State Detection; §10
‘Sounds ableist’	This standard explicitly names disability as a permanent baseline, rejects return-only framing, and requires co-design with disabled people.	§1.3 Capacity Alignment; §11.3
‘Noble savage fallacy in evolution section’	The evolutionary section explicitly states it is not an argument for retreat and acknowledges that some people thrive in high-stimulation digital contexts.	§1.4 Evolutionary Foundation

‘One company certifying itself’

Whelmish products are evaluated by the same process as any other applicant. Third-party review by independent practitioners is required at all levels.

§14.2 Cert Requirements

‘What about emergencies’

Emergency and legal contexts have explicit carveouts. Minimum necessary demand still applies in emergencies.

§1.7 Emergency Boundary

Part XVI · External Standards Alignment

NSIF is intended to sit beside existing standards, not replace them. The Capacity-Safe Systems Standard adds a nervous-system and capacity layer to human-centered design, accessibility, trauma-informed practice, AI risk management, and child-development-informed design.

Framework	Alignment	NSIF adds
SAMHSA Trauma-Informed Approach	Safety, trustworthiness, peer support, collaboration, empowerment, cultural sensitivity	Turns trauma-informed values into interface, workflow, AI, and return-path requirements
WCAG 2.2	Perceivable, operable, understandable, robust access with testable success criteria	Adds state-based capacity adaptation, shame removal, recovery protection, low-capacity mode
ISO 9241-210 Human-Centered Design	Design focused on users, needs, usability, well-being, accessibility, and sustainability	Adds nervous-system state, recovery sequence, capacity alignment, and low-capacity evidence
NIST AI Risk Management Framework	Trustworthy AI: validity, safety, security, accountability, transparency, explainability	Adds capacity-aware response, anti-dependency design, distress-sensitive interaction, capacity data governance
Child Development and Toxic Stress Literature	Children need safe, stable, nurturing relationships; prolonged stress affects development	Adds age-specific design requirements and child/teen pathways that do not scale adult systems downward
Dark Pattern / Deceptive Design Research	Interfaces can steer, coerce, or deceive users into choices not in their interest	Defines distress extraction, shame architecture, urgency pressure, and hidden exits as capacity harms

Part XVII · Implementation Checklist

The following checklist converts this standard into build requirements. Use it to audit existing products, guide new development, and prepare for certification.

Area	Minimum Requirement	Evidence Needed
Entry	User can declare or bypass state; system defaults to low burden when state is unknown; Zero Demand accessible without explanation	Screenshots, flow map, user tests
Low-Capacity Mode	Choices collapse, language shortens, motion reduces, one step visible; Zero Demand mode available	State-based UI variants and copy
Offload	User can externalize load without immediate advice or categorization	Offload flow and data retention rules
Simplification	System surfaces only what matters now; does the filtering for the user	Decision logic and content hierarchy
Stopping	Stop, pause, save, exit, and return are visible at all capacity states	Exit path screenshots and return flow tests
Return	No absence shame, no lost progress language, no catch-up punishment	Notification and return copy audit
Recovery Protection	Completion does not trigger immediate new demand	Post-action flow and notification rules
Privacy	Capacity-sensitive data minimized, consented, correctable, erasable; never used for power-asymmetric purposes	Data map, consent copy, deletion flow
AI Guardrails	AI adapts to all eight capacity states; has human support pathway; cannot simulate attachment dependency	Prompt policy, response examples, safety evaluations
Children / Teens	Developmental match, trusted adult pathway, no comparison shame, no algorithmic identity exploitation	Youth-specific UX audit and caregiver review
Cultural Competence	System tested across cultural contexts in which it will be deployed	Cross-cultural user testing documentation
Disability	Co-designed with disabled users; permanent capacity baselines respected; no return-only framing	Disability co-design documentation
Power Dynamics	Capacity data cannot influence employer, school, or institutional decisions	Data governance policy and access restriction audit

Measurement	Capacity alignment and harm-reduction metrics tracked without surveillance	Metrics plan, governance rules, anti-gaming provisions
Repair Protocol	System has documented harm incident process; repair steps are defined and accessible	Incident response documentation

Part XVIII · Glossary of NSIF Terms

These terms are specific to the NSIF framework. Where a term overlaps with existing clinical or scientific usage, the NSIF definition applies within this standard.

Capacity	The available bandwidth of the human nervous system at a given moment — its ability to perceive, process, decide, respond, connect, and act. Not a character trait. A biological reality.
Capacity Alignment	The goal of every Capacity-Safe system: adapting to the person's actual capacity whether the aim is return, stabilization, expansion, or sovereignty. Replaces the prior 'Capacity Return' framing to include permanent baselines.
Capacity Debt	The accumulated biological cost of previous demands that exceeded available resources. Measurable as allostatic load carried forward.
Capacity Map	The NSIF model of eight operational states: Flowing, Grounded, Recovering, Stretched, Flooded, Frozen, Fractured, Zero Demand. Plus the Disconnected Function overlay.
Capacity-Safe	A design designation indicating that a system has been built to function for humans across all capacity states, including the lowest.
Capacity-Safe Mark	The certification designation awarded under the Capacity-Safe Systems Standard to systems meeting its requirements.
Capacity-Safe Practitioner	An individual certified by completing the NSIF practitioner course and demonstrating applied use of the framework on a real-world system.
Capacity State	The specific position within the Capacity Map that describes a person's current nervous system availability.
Co-regulation	The biological process by which one nervous system supports the regulation of another through proximity, tone, rhythm, and attunement.
Demand Architecture	The intentional design of how much is asked from a human, in what form, at what time, and under what conditions. Demand architecture is the structural mechanism through which a system is Capacity-Safe or not. A calm color palette is aesthetic. Demand architecture is the actual load the system places on the human nervous system at each capacity state. The Capacity-Safe standard is, at its core, a demand architecture standard.
Disconnected Function	A cross-state overlay in which a person performs tasks competently while internally disconnected. Not a failure. A protective state. The system must not exploit it.
Exit Sovereignty	The right of the user to pause, stop, leave, or return without punishment, shame, data loss, social penalty, or re-entry friction.

Flooded	A capacity state characterized by overwhelm above the window of tolerance. Logical processing is impaired; emotional responses are dominant.
Flowing	A capacity state of full resource, generative energy, and creative or productive momentum. The system's role is to support without interrupting.
Fractured	A capacity state characterized by chronic capacity instability and long-arc dysregulation. Not a momentary state. Not an identity. A condition.
Frozen	A capacity state characterized by shutdown below the window of tolerance. Numbing, disconnection, and body-first experience dominate.
Grounded	A capacity state of full regulation and availability. Complete access to cognitive, relational, and operational function.
Infrastructure Stack	The NSIF model of four layers — Somatic, Cognitive, Relational, Operational — plus transition zones between them, that must be addressed in sequence.
Nervous System Infrastructure	The NSIF framework itself — the missing layer underneath every human-facing system.
Recovery Arc	The longitudinal seven-phase path from collapse toward sovereignty: Collapse, Stabilization, Trench, Integration, Relapse, Consolidation, Sovereignty. Describes what happens in the person over time.
Recovery Spine	The NSIF five-phase session-level protocol: Safety, Offload, Simplification, Gentle Movement, Recovery Protection. Describes what the system does.
Recovering	A capacity state of recent return from low capacity. Stable but tender. Requires gentle pacing and warm reentry design.
Regulated States Gap	The volatile transition zone between dysregulation and genuine regulation. Calm is not capacity. Quiet is not consent. Regulated is not ready.
Shame Architecture	Any design pattern that produces shame as a byproduct of normal human behavior — absence, error, pace, need, capacity variability.
Sovereignty	The state in which the user has internalized capacity awareness and can shape systems to their capacity rather than the reverse. The final Recovery Arc phase.
Stretched	A capacity state of reduced but functional bandwidth. Capacity is taxed but not overwhelmed.
Temporal Justice	The calibration of time, deadlines, and pacing to capacity state rather than uniform imposition. Time is a somatic experience, not a fixed administrative unit.
Window of Tolerance	The optimal zone of nervous system arousal within which cognitive, emotional, and relational function is available. (After Siegel, 1999.) Expandable through appropriate support.

Zero Demand

A capacity state or mode in which any operational demand causes harm. The system holds space only. No tasks, no goals, no progression through the Recovery Spine.

Closing · Why This Standard Exists

I am regulated, therefore I can.

Every person who has ever abandoned a planner, quit an app, failed to show up, disappeared from a system, or called themselves broken for not being able to keep up — was not failing.

They were in a capacity state that the system was not designed for.

They were moving through a transition zone the system could not see.

They were carrying capacity debt the system refused to acknowledge.

They were speaking from the shadow infrastructure, and the system demanded they name the unnameable.

The failure was the design.

This standard exists to end that particular failure. Not for some users. For all of them. Not as an accommodation. As the base requirement.

Human capacity is not a bug in the system. It is the system. Everything built for humans must be built around it.

The Capacity-Safe Systems Standard is Version 3.0. It will be revised as evidence grows, as practitioners apply it, as disabled and marginalized voices shape it, as cultures adapt it, as AI evolves, and as the world changes. What will not change is the foundational argument:

Every system built for humans should work for humans as they actually are — not as the system wishes they were. And every system must hold what humans cannot say, cannot name, and sometimes cannot know about themselves.

That is the promise of Nervous System Infrastructure.

That is what NSIF is for.

Vanessa Williams

Founder, Whelmish LLC

Creator, NSIF — Nervous System Infrastructure Framework

Version 3.0 · Published 2026

capacitystandard.com · © 2026 Whelmish LLC — Nervous System Infrastructure Framework

All rights reserved. Framework concepts, terminology, and standards are the intellectual property of Whelmish LLC. Citation with attribution is permitted. Commercial use requires written authorization.

The CalmFlow OS family by Whelmish was the first product line built on the Capacity-Safe Systems Standard. whelmish.com