

LeadDev Berlin 2026

Burnout: A team failure mode

Diana Nănuți
Senior Software Engineer @ Chainalysis



Intro

- Senior Software Engineer @ **Chainalysis**
- Community Organizer @ **A JUG For Everyone**
- Mentor @ **She Can Code**
- Ad-hoc Mentor @ **Women Coding Community**
- **Women In Tech Global Ambassador**
- **Open-Source** Contributor
- **Table tennis and ski enjoyer**



Intro

- Senior Software Engineer @ **Chainalysis**
- Community Organizer @ **A JUG For Everyone**
- Mentor @ **She Can Code**
- Ad-hoc Mentor @ **Women Coding Community**
- **Women In Tech Global Ambassador**
- **Open-Source** Contributor
- **Table tennis and ski enjoyer**



Intro

- Senior Software Engineer @ **Chainalysis**
- Community Organizer @ **A JUG For Everyone**
- Mentor @ **She Can Code**
- Ad-hoc Mentor @ **Women Coding Community**
- **Women In Tech Global Ambassador**
- **Open-Source** Contributor
- **Table tennis and ski enjoyer**



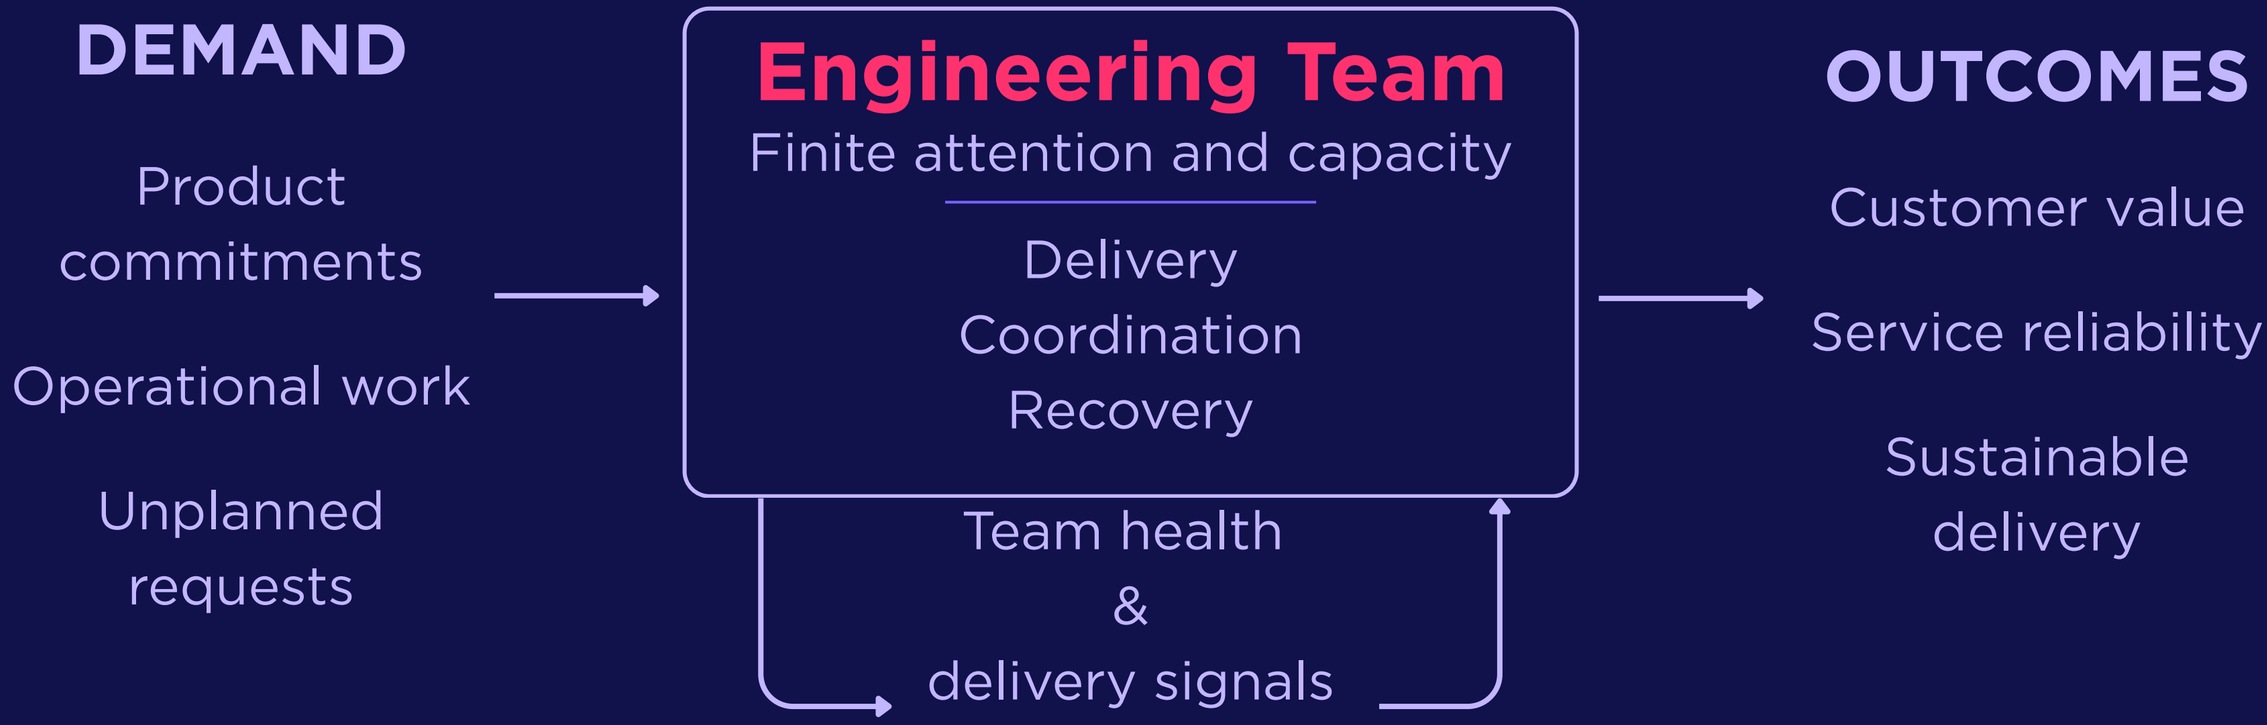
Is your team's health as observable as your software?

Is your team's health as observable as your software?

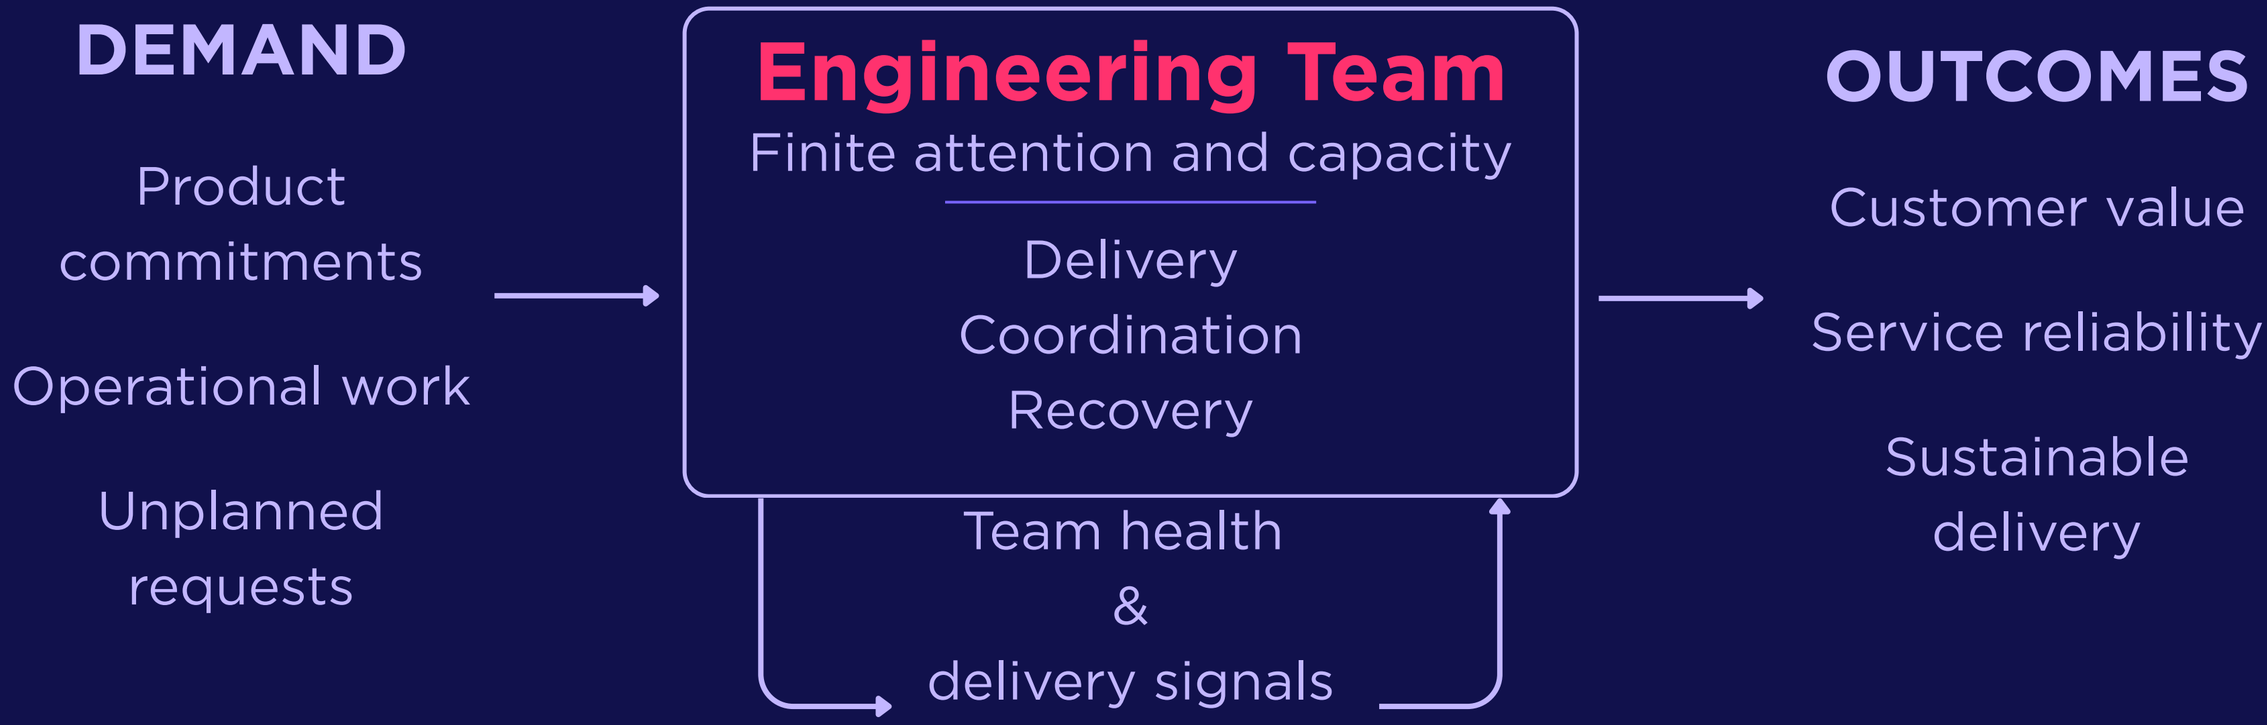
Is your team's health as observable as your software?

Is your team's health as observable as your software?

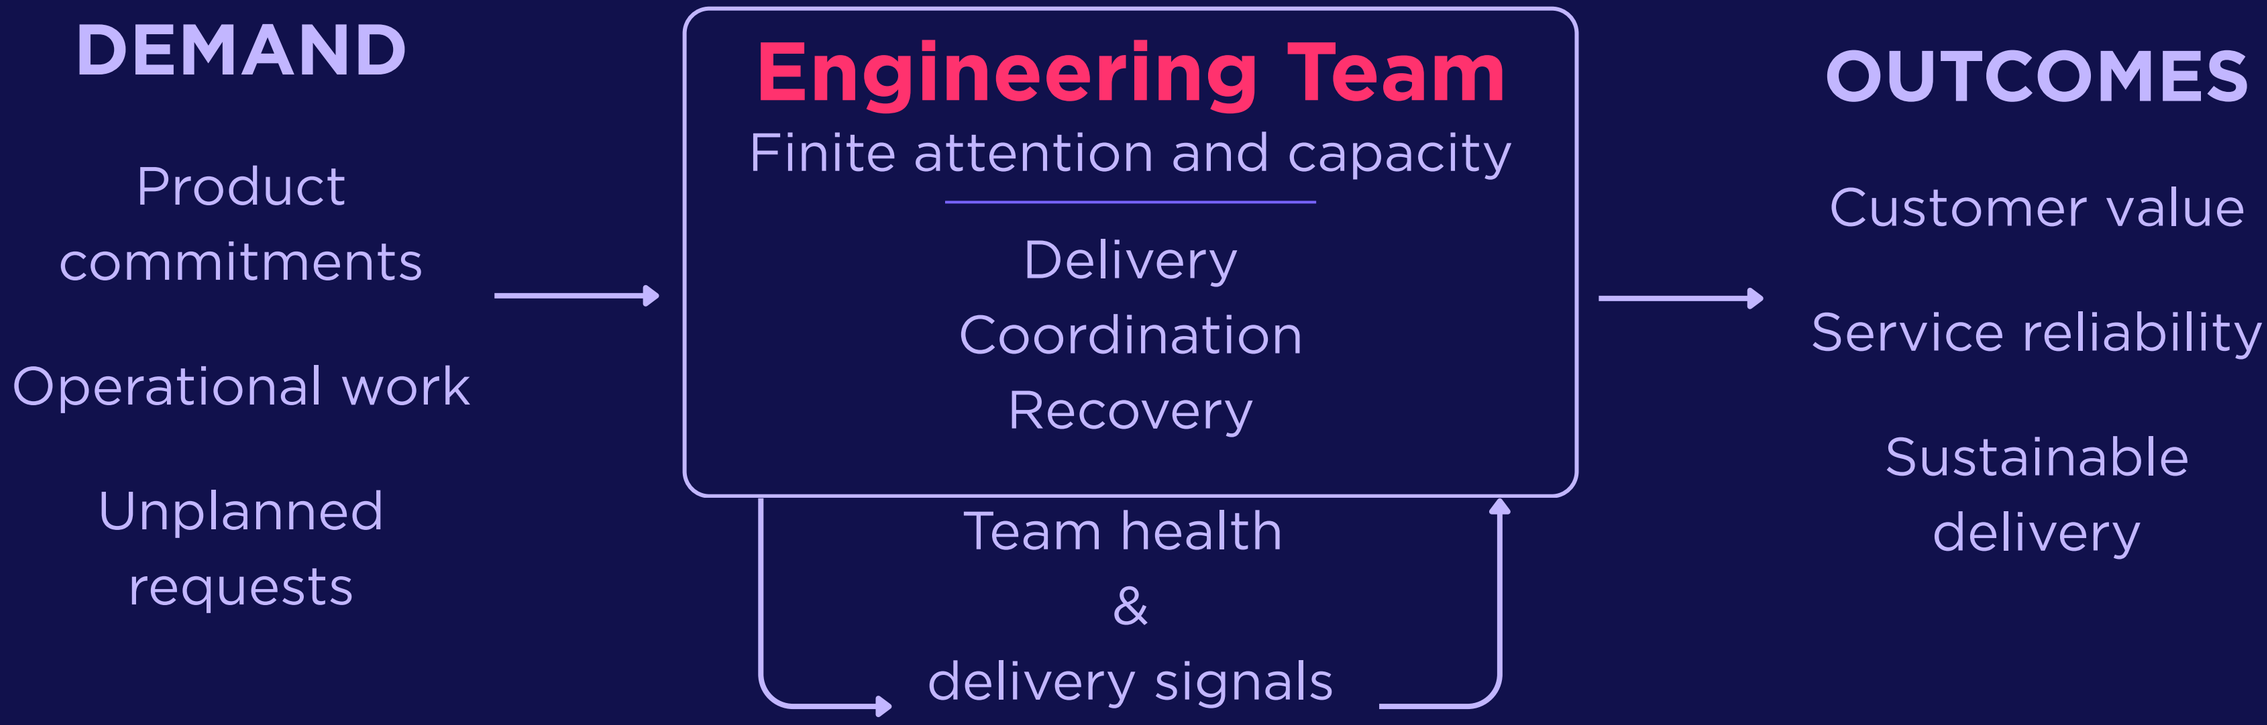
The team as a socio-technical system



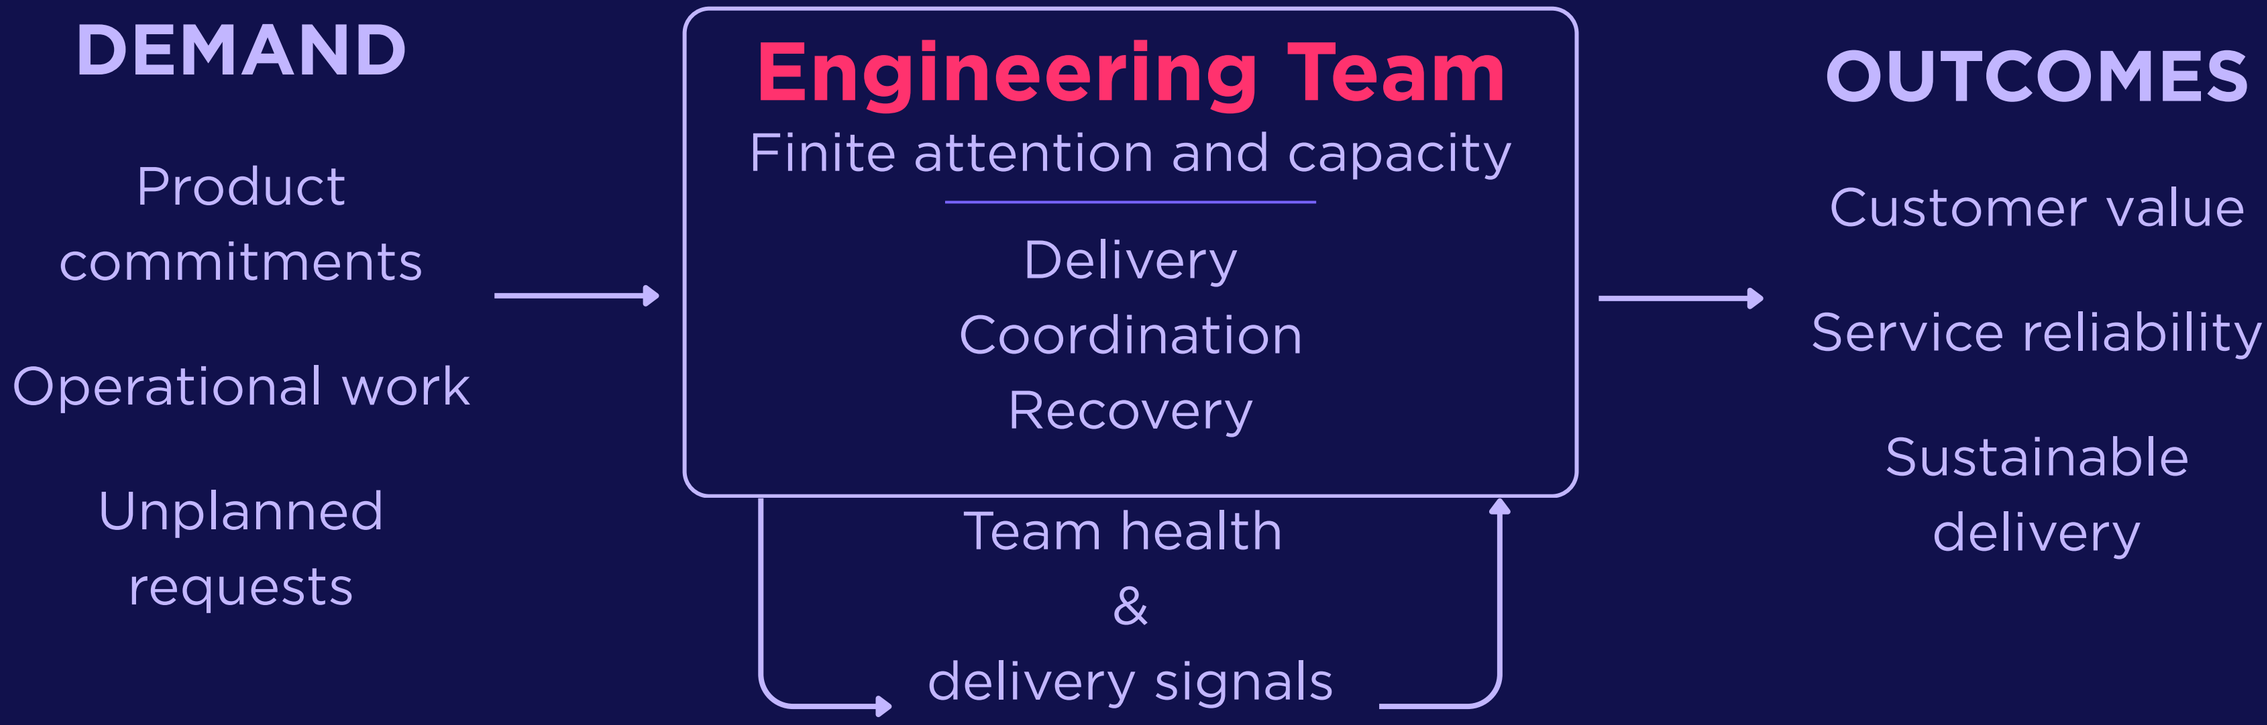
The team as a socio-technical system



The team as a socio-technical system



The team as a socio-technical system



We built a high-throughput system without human safety margins.

Why would an employer care?



Why would an employer care?



Why would an employer care?

- Productivity loss, disengagement
- Absenteeism
- Medical leave
- **AI Brain Fry**

[1] https://hbr.org/2026/03/when-using-ai-leads-to-brain-fry



Why would an employer care?

- Productivity loss, disengagement
- Absenteeism
- Medical leave
- **AI Brain Fry**

[1] https://hbr.org/2026/03/when-using-ai-leads-to-brain-fry



Why would an employer care?

- Productivity loss, disengagement
- Absenteeism
- Medical leave
- **AI Brain Fry**

[1] https://hbr.org/2026/03/when-using-ai-leads-to-brain-fry

$150M per year

Suboptimal
decision-making
$150M per year
for a company generating
$5B in annual revenue.

Suboptimal
decision-making
$150M per year
for a company generating
$5B in annual revenue.

Business impact

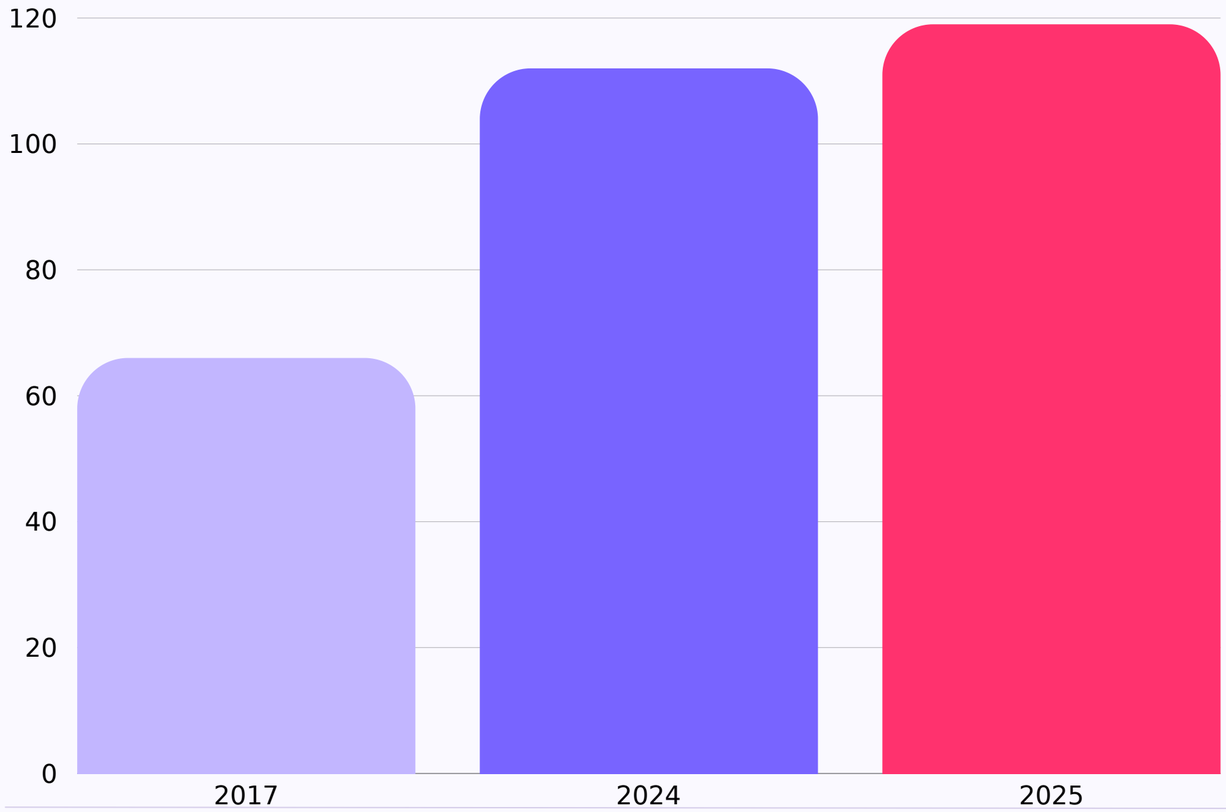
Decision fatigue	**+33%**
Minor workplace errors	**+11%**
Major workplace errors	**+39%**
Active intention to quit	**+39%**

Business impact

Decision fatigue	**+33%**
Minor workplace errors	**+11%**
Major workplace errors	**+39%**
Active intention to quit	**+39%**

Stress-related absence

Germany



Sick days per 100 KKH insured individuals

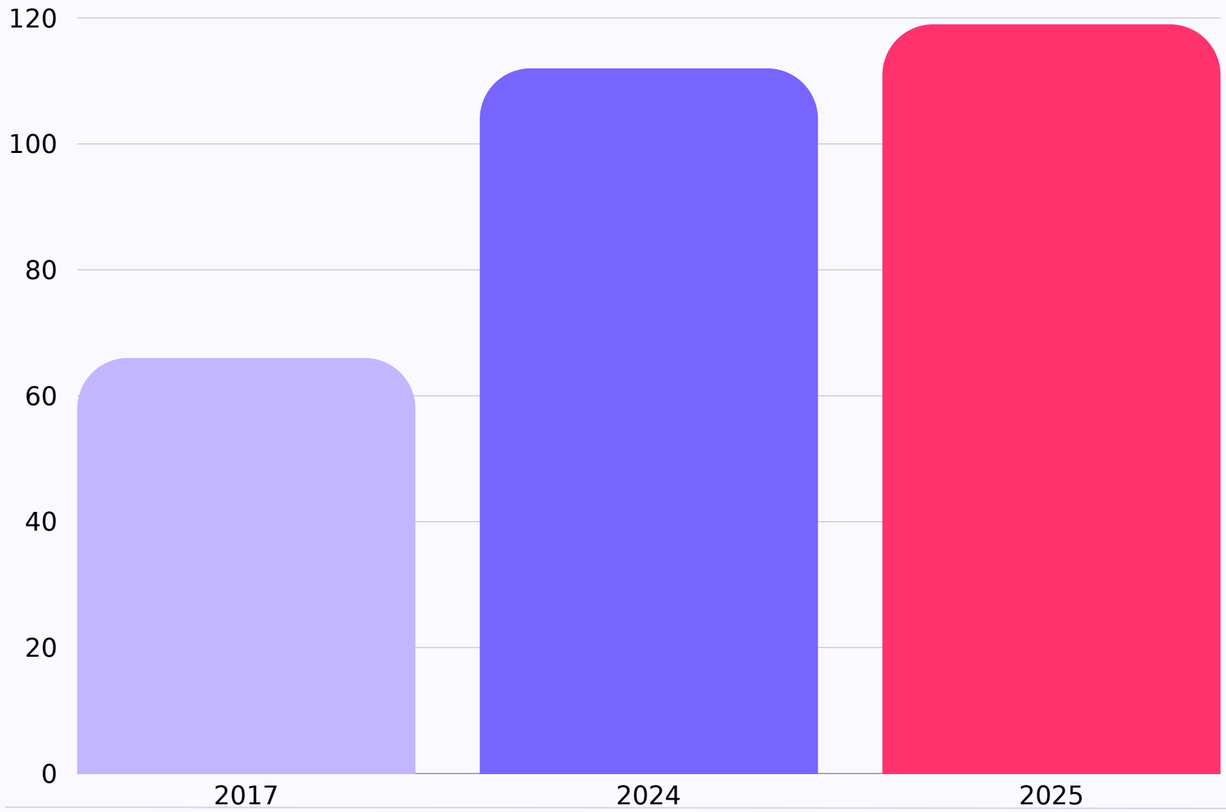
~ 80%

increase since 2017

[2] https://www.kkh.de/presse/pressemeldungen/au-daten-f43

Stress-related absence

Germany

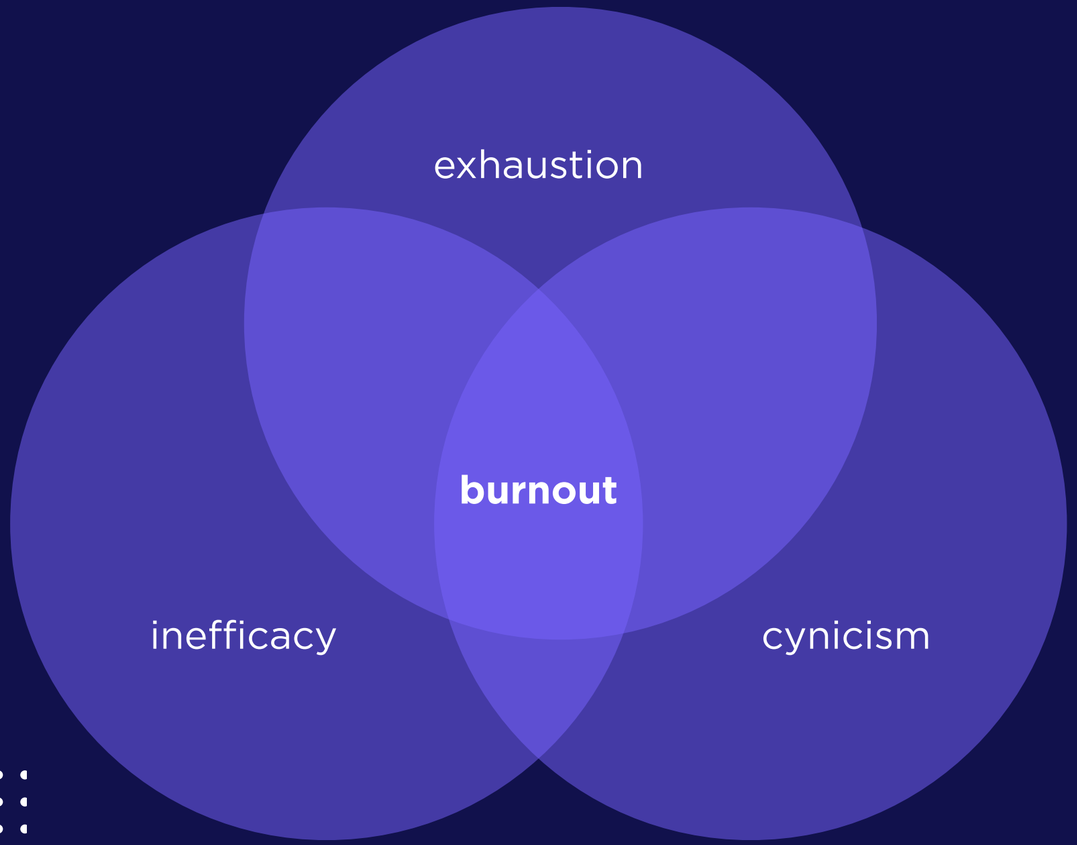


Sick days per 100 KKH insured individuals

~ 80%

increase since 2017

[2] https://www.kkh.de/presse/pressemeldungen/au-daten-f43

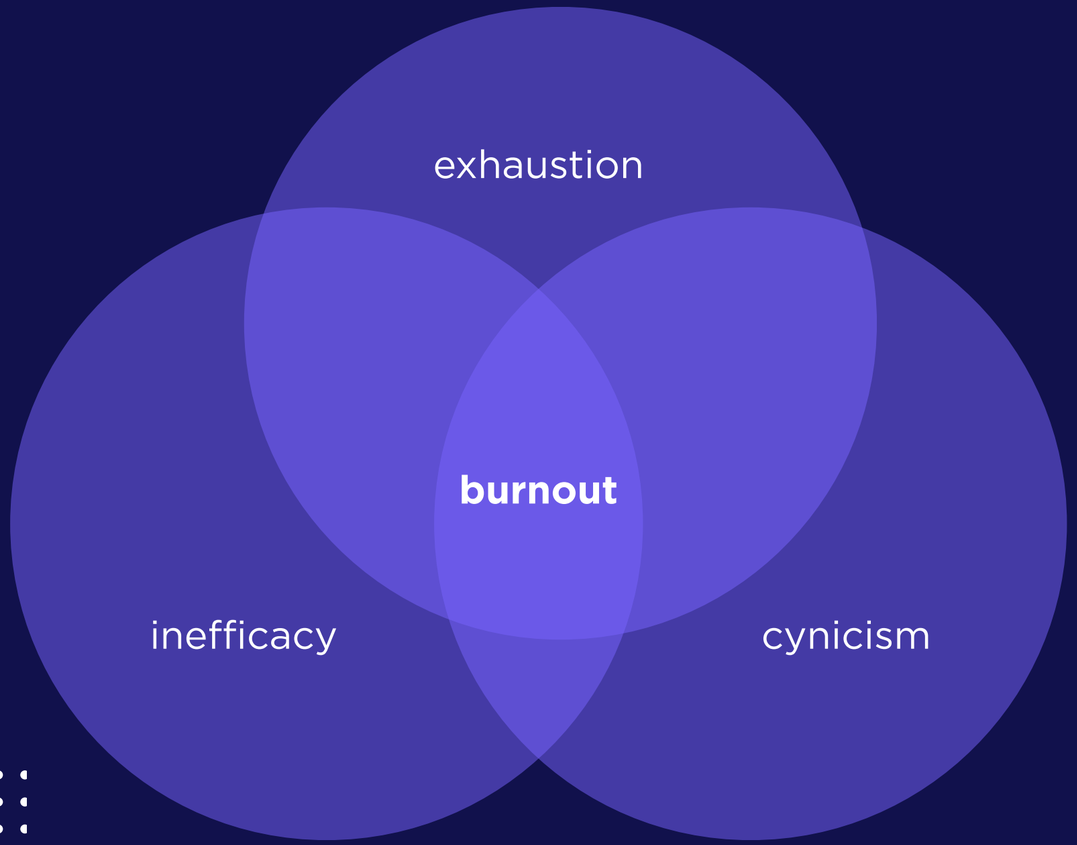


Reframing The Problem

The **World Health Organization (ICD-11)** defines burnout as an **occupational phenomenon** resulting from **chronic workplace stress** that has not been successfully managed [3].

AI adoption adds **new layers of cognitive load**, making **human systems** even **more fragile**.

[3] https://www.who.int/news/item/28-05-2019-burn-out-an-occupational-phenomenon-international-classification-of-diseases

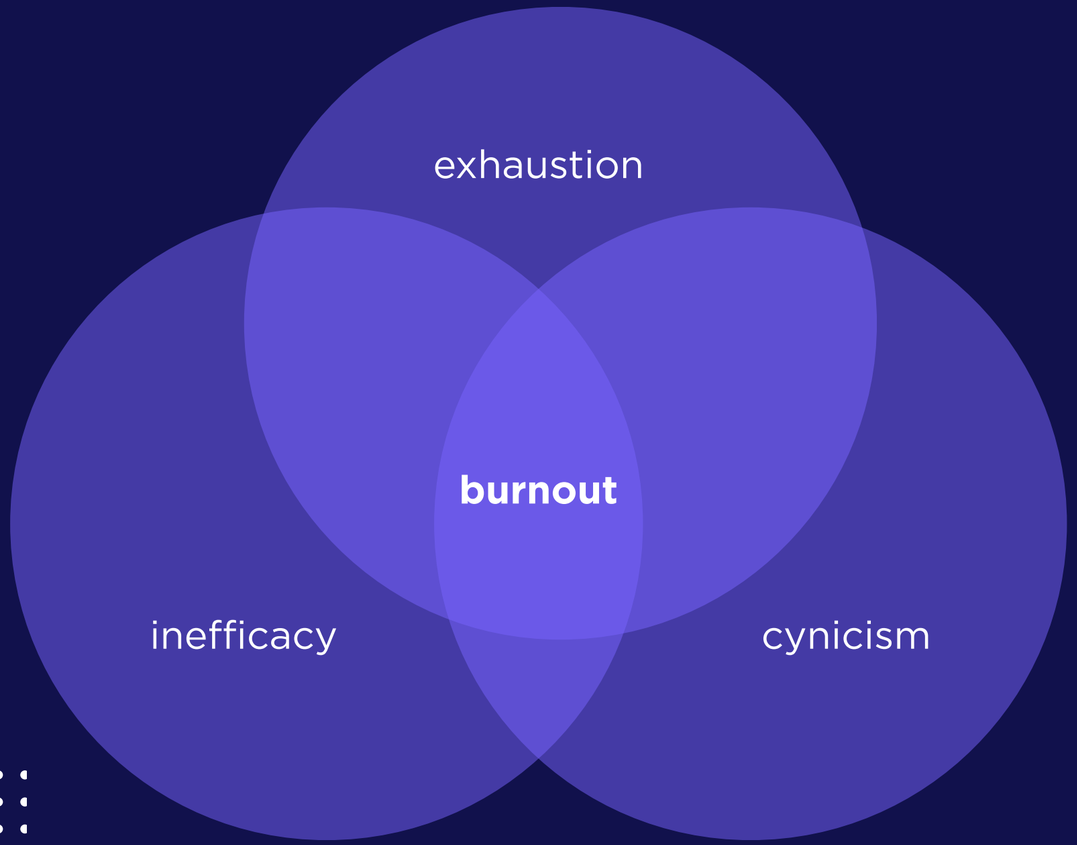


Reframing The Problem

The **World Health Organization (ICD-11)** defines burnout as an **occupational phenomenon** resulting from **chronic workplace stress** that has not been successfully managed [3].

AI adoption adds **new layers of cognitive load**, making **human systems** even **more fragile**.

[3] https://www.who.int/news/item/28-05-2019-burn-out-an-occupational-phenomenon-international-classification-of-diseases

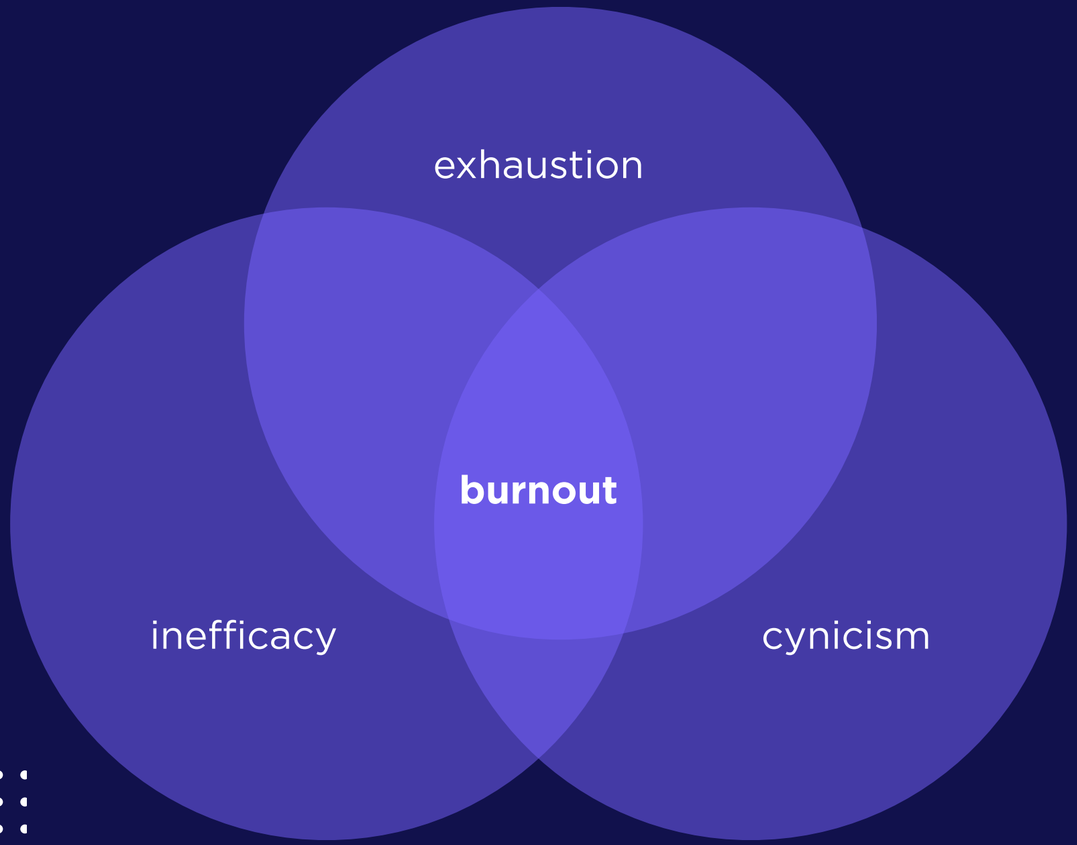


Reframing The Problem

The **World Health Organization (ICD-11)** defines burnout as an **occupational phenomenon** resulting from **chronic workplace stress** that has not been successfully managed [3].

AI adoption adds **new layers of cognitive load**, making **human systems** even **more fragile**.

[3] https://www.who.int/news/item/28-05-2019-burn-out-an-occupational-phenomenon-international-classification-of-diseases



Reframing The Problem

The **World Health Organization (ICD-11)** defines burnout as an **occupational phenomenon** resulting from **chronic workplace stress** that has not been successfully managed [3].

AI adoption adds **new layers of cognitive load**, making **human systems** even **more fragile**.

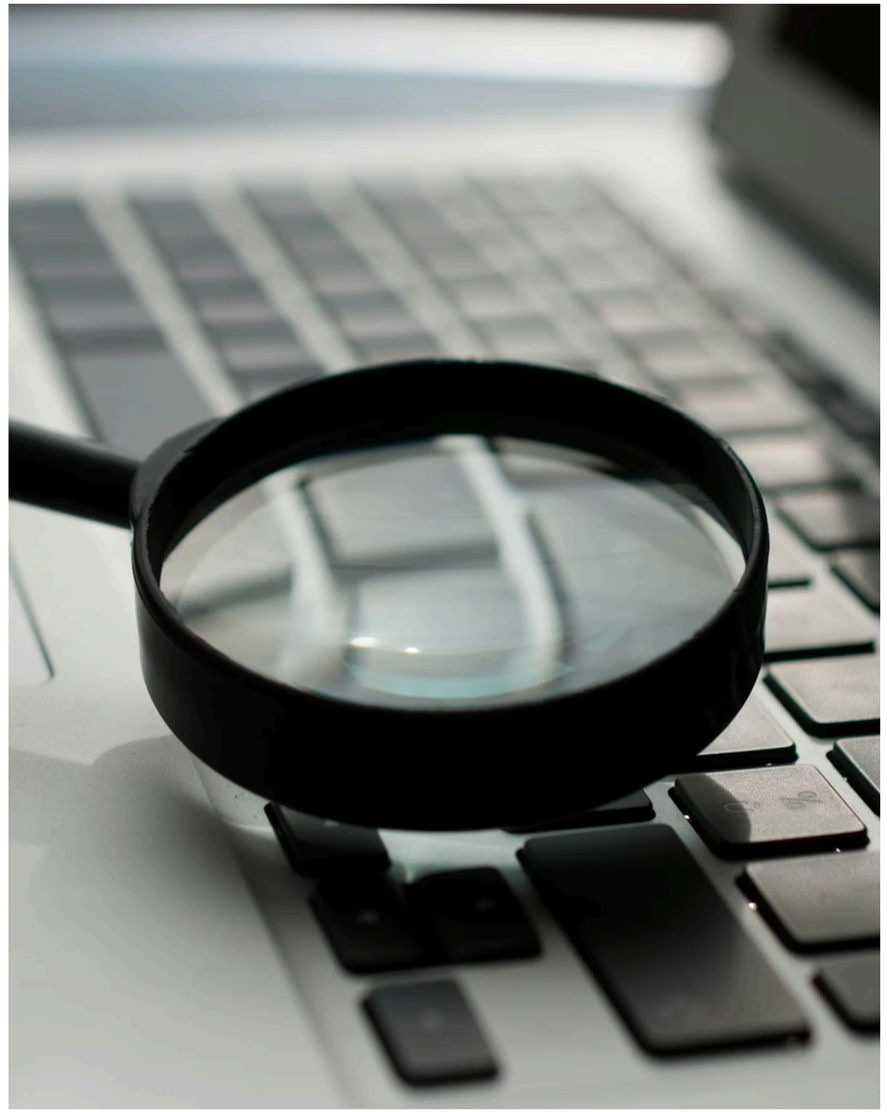
[3] https://www.who.int/news/item/28-05-2019-burn-out-an-occupational-phenomenon-international-classification-of-diseases

SRE Lens

Teams are **engineered systems.**

- **Load**
- **Capacity**
- **Margins**
- **Failure Modes**

Burnout is a **systemic failure mode.**

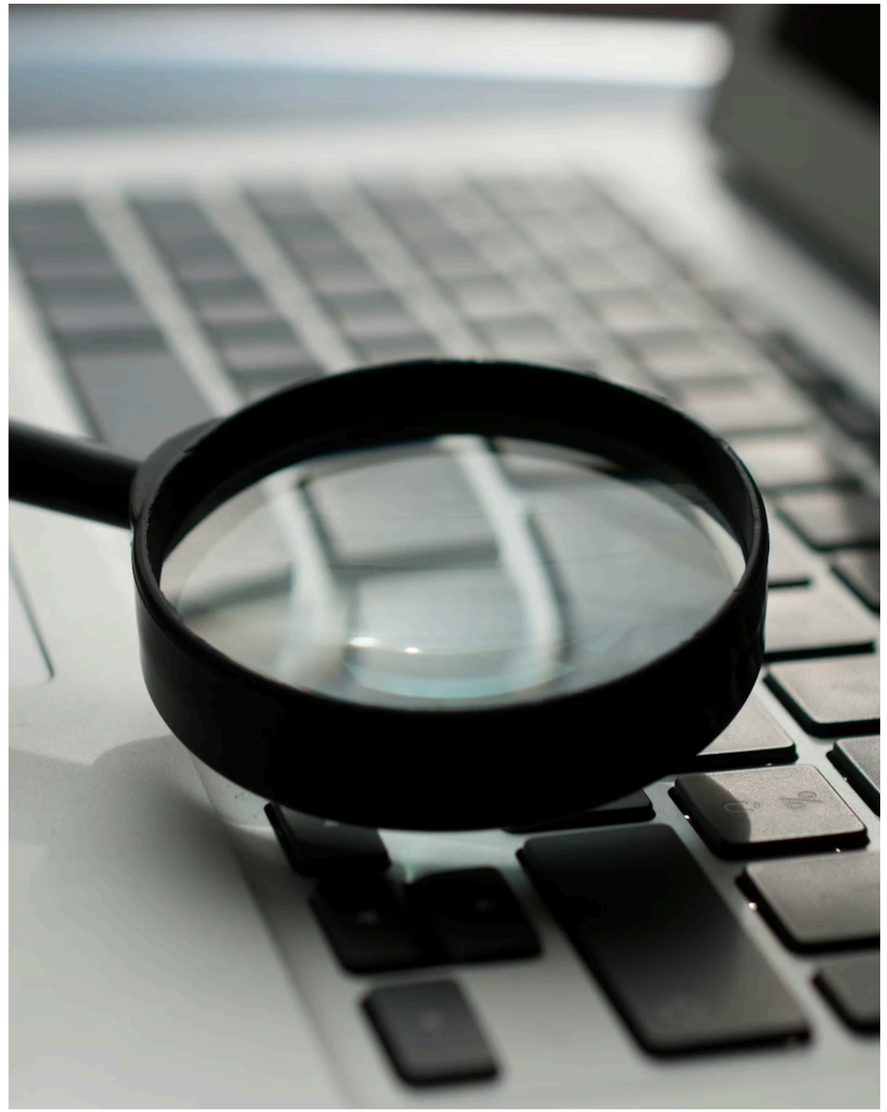


SRE Lens

Teams are **engineered systems.**

- **Load**
- **Capacity**
- **Margins**
- **Failure Modes**

Burnout is a **systemic failure mode.**

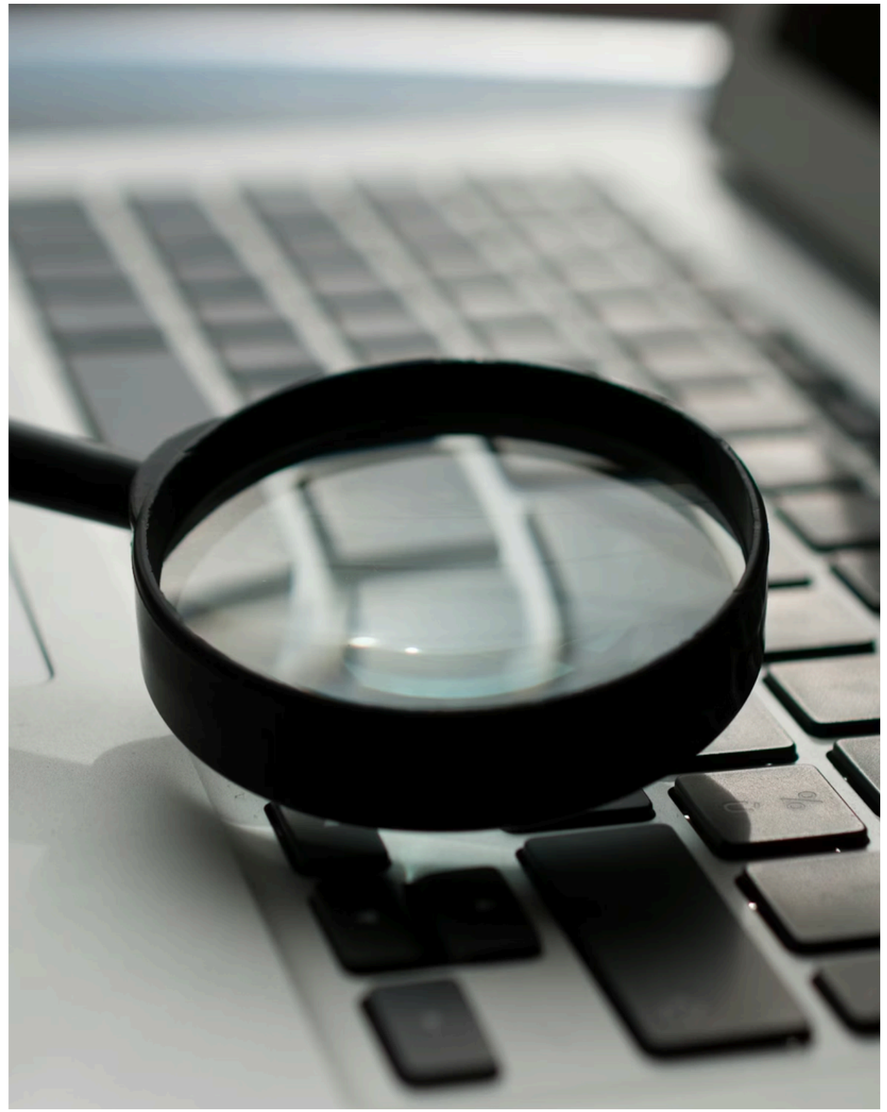


SRE Lens

Teams are **engineered systems.**

- **Load**
- **Capacity**
- **Margins**
- **Failure Modes**

Burnout is a **systemic failure mode.**



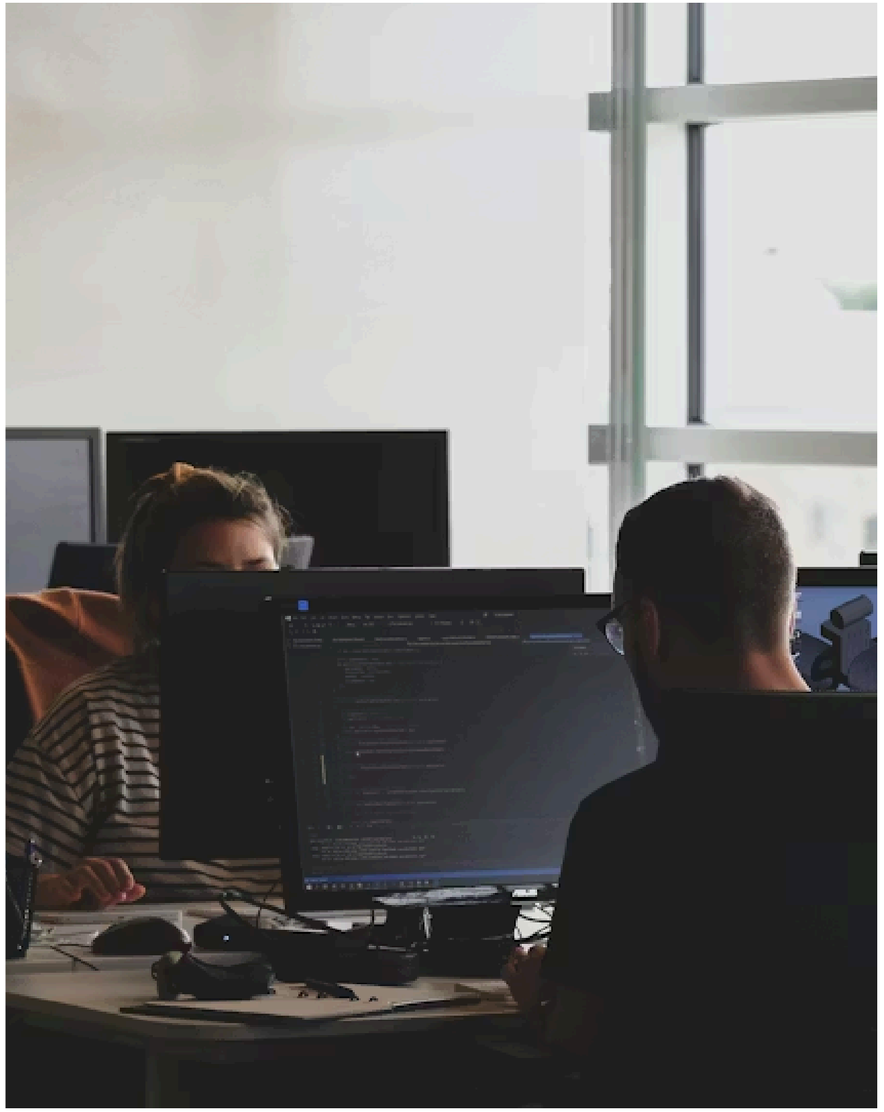
We CANNOT control what we DON’T observe.

We CANNOT control what we DON’T observe.

01

Make Load Visible

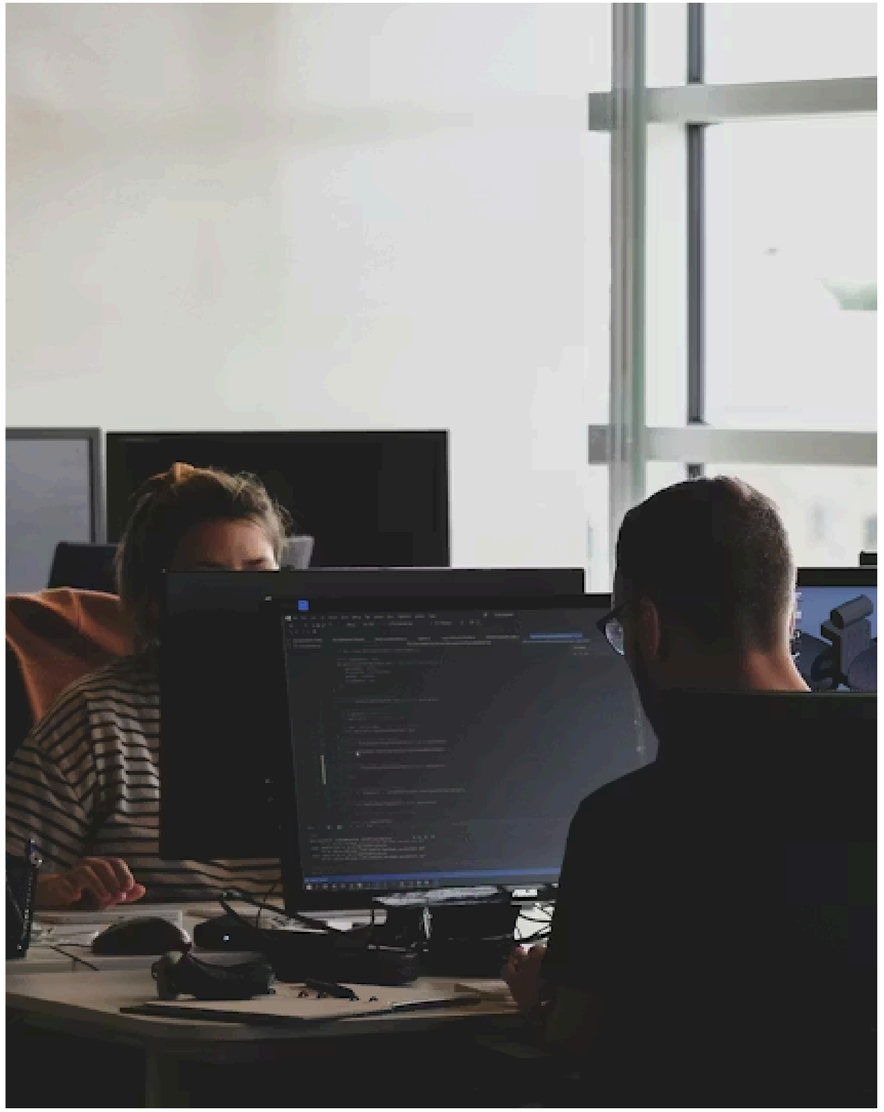
- COGNITIVE
- INTERRUPTIONS
- ON-CALL
- EMOTIONAL



01

Make Load Visible

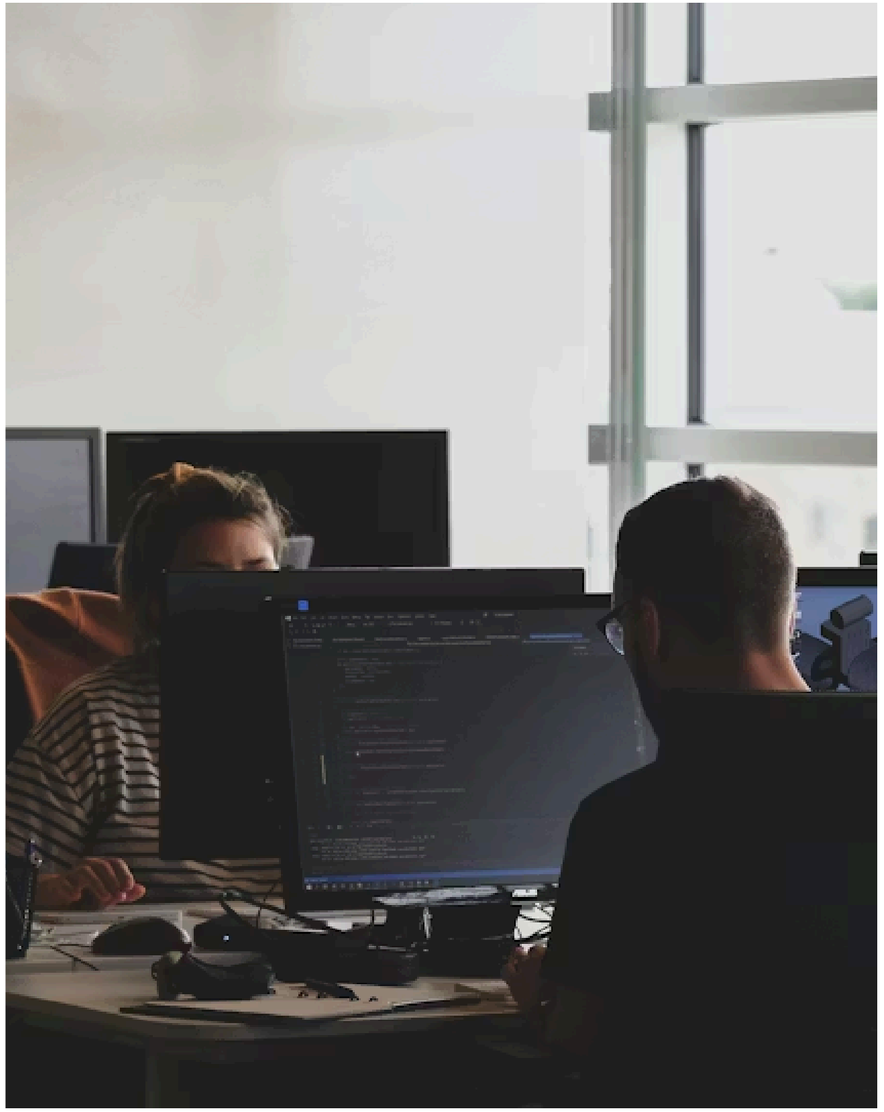
- COGNITIVE
- INTERRUPTIONS
- ON-CALL
- EMOTIONAL



01

Make Load Visible

- COGNITIVE
- INTERRUPTIONS
- ON-CALL
- EMOTIONAL



Human Observability Metrics

Live indicators mirroring your proposed SLIs.

WARNING

0.82 *ratio*

Verification Load

Reviewing AI output vs. creating value. High ratios indicate imminent cognitive burnout.

DANGER

12 *hrs left*

Error Budget

Remaining system margin for innovation and recovery before mandatory Safe Mode.

DANGER

18 */ day*

Context Swaps

Average task interruptions. Leads to high human MTTR (Mean Time To Flow).

OK

0.85 *mesh*

Team Redundancy

Knowledge distribution across nodes. 1.0 = No Single Points of Failure.

WARNING

High *latency*

Fatigue Signal

Leading indicator based on incident recovery duration trends.

WARNING

1 *Hero*

SPOF Nodes

Number of individuals holding siloed critical system knowledge.

Circuit Breaker Status: ALERT

Human Error Budget exhausted. System is approaching cascading failure. Enforce SAFE MODE: freeze all non-critical feature work and redirect capacity to documentation and rest.

TRIGGER SAFE MODE

Human Observability Metrics

Live indicators mirroring your proposed SLIs.

WARNING

0.82 *ratio*

Verification Load

Reviewing AI output vs. creating value. High ratios indicate imminent cognitive burnout.

DANGER

12 *hrs left*

Error Budget

Remaining system margin for innovation and recovery before mandatory Safe Mode.

DANGER

18 */ day*

Context Swaps

Average task interruptions. Leads to high human MTTR (Mean Time To Flow).

OK

0.85 *mesh*

Team Redundancy

Knowledge distribution across nodes. 1.0 = No Single Points of Failure.

WARNING

High *latency*

Fatigue Signal

Leading indicator based on incident recovery duration trends.

WARNING

1 *Hero*

SPOF Nodes

Number of individuals holding siloed critical system knowledge.

Circuit Breaker Status: ALERT

Human Error Budget exhausted. System is approaching cascading failure. Enforce SAFE MODE: freeze all non-critical feature work and redirect capacity to documentation and rest.

TRIGGER SAFE MODE

02

Define Leading Indicators

Human MTTR	→	Pulse surveys 24h, 72h, 1 week
Verification Load	→	AI Signal-to-Noise ratio
Human Error Budget Burn Rate	→	How fast the team is eating into its non-feature capacity

02

Define Leading Indicators

Human MTTR	→	Pulse surveys 24h, 72h, 1 week
Verification Load	→	AI Signal-to-Noise ratio
Human Error Budget Burn Rate	→	How fast the team is eating into its non-feature capacity

02

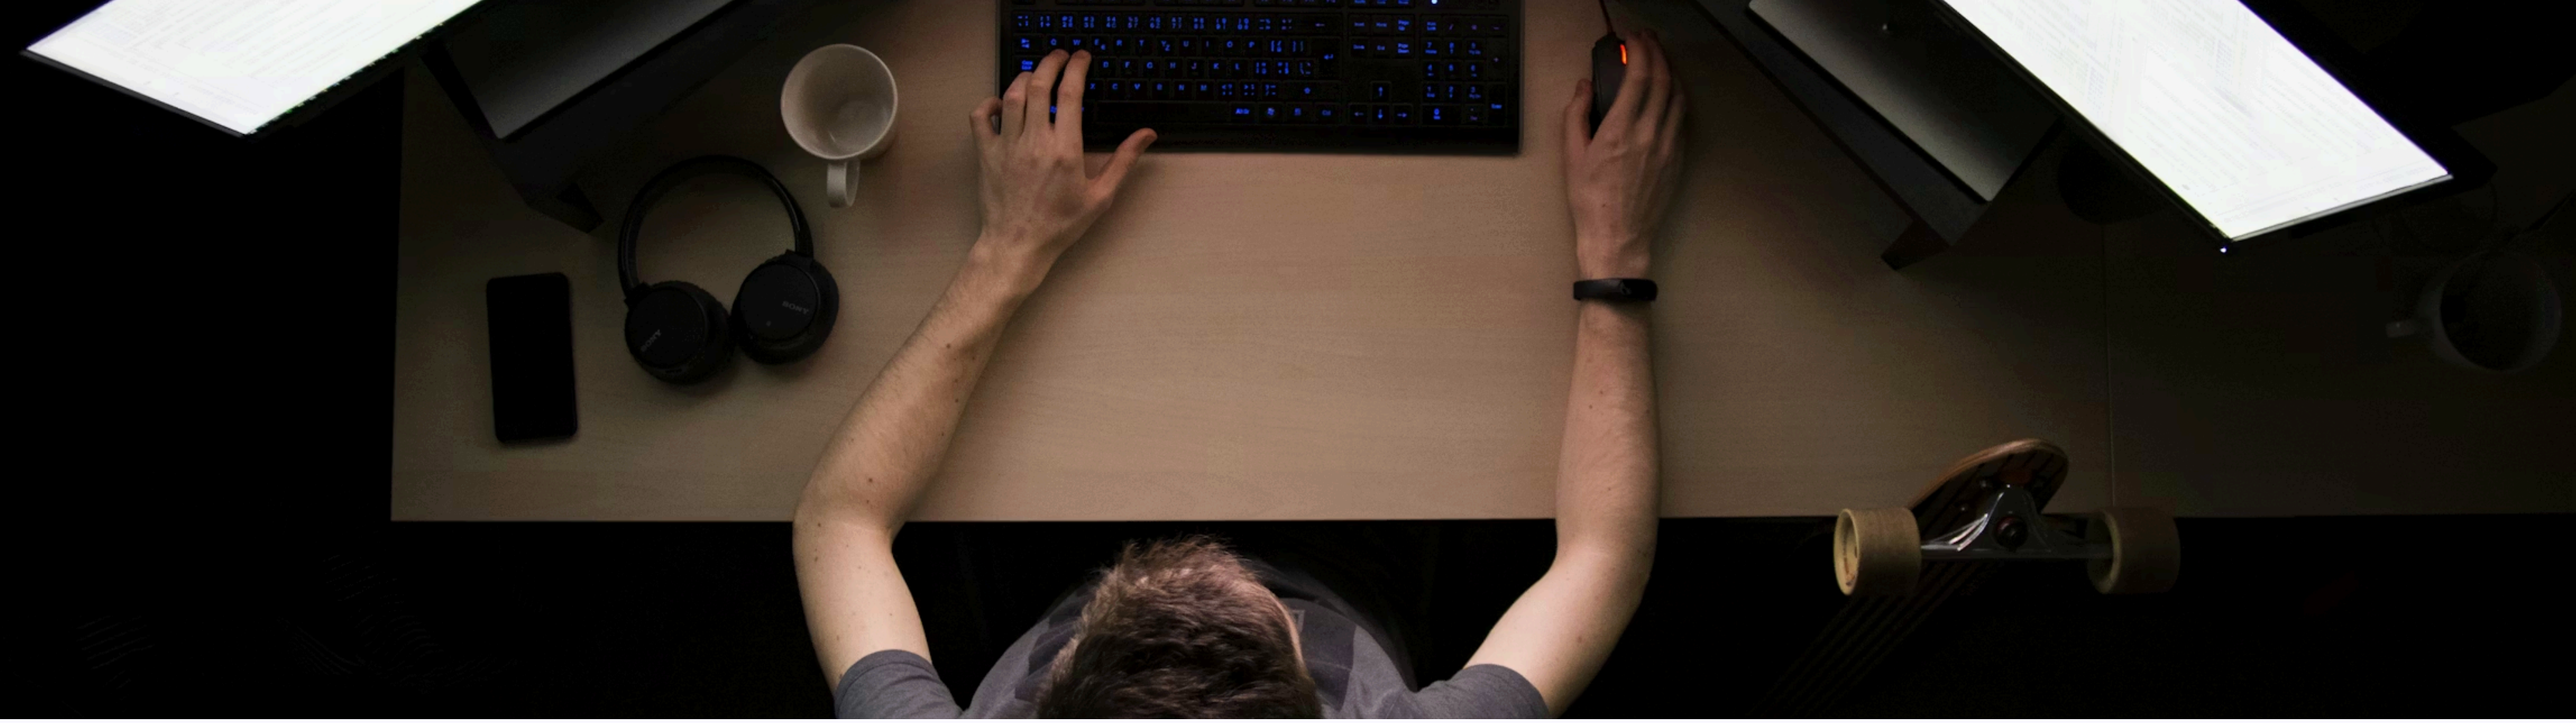
Define Leading Indicators

Human MTTR	→	Pulse surveys 24h, 72h, 1 week
Verification Load	→	AI Signal-to-Noise ratio
Human Error Budget Burn Rate	→	How fast the team is eating into its non-feature capacity

02

Define Leading Indicators

Human MTTR	→	Pulse surveys 24h, 72h, 1 week
Verification Load	→	AI Signal-to-Noise ratio
Human Error Budget Burn Rate	→	How fast the team is eating into its non-feature capacity



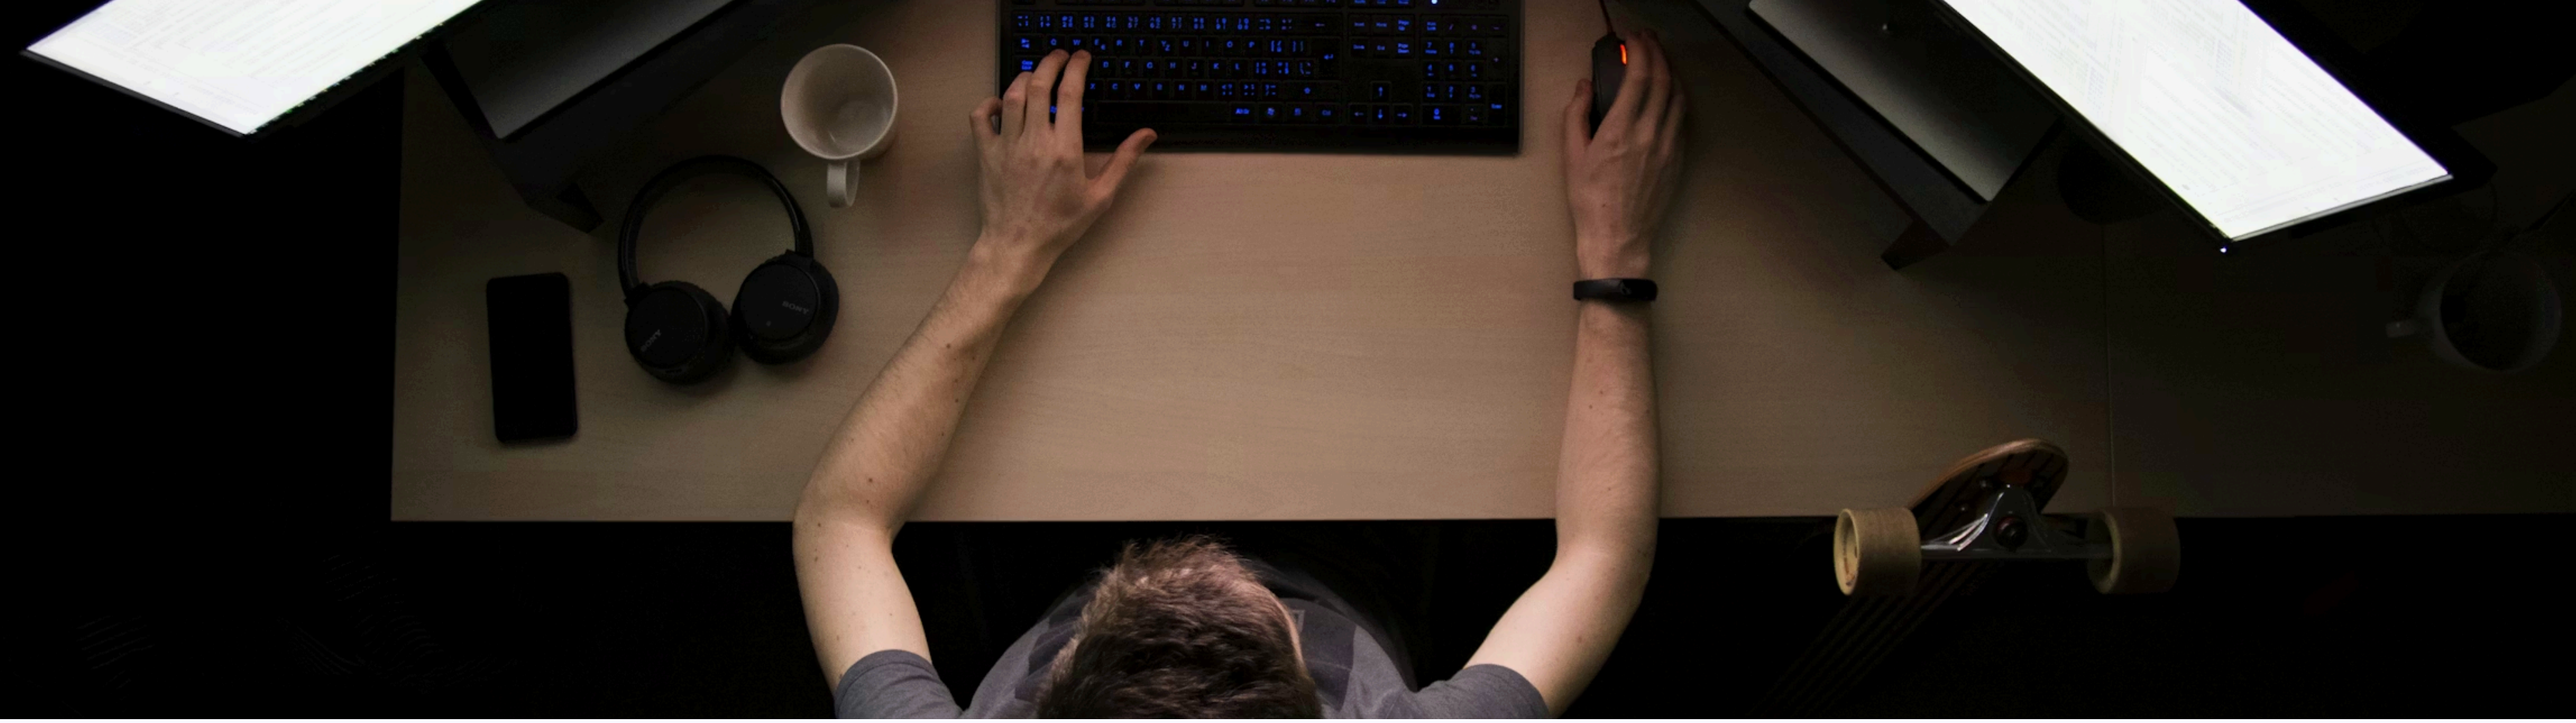
03

Redundancy Mechanisms

Heroes are unacknowledged single points of failure.

Advocate and implement:

- knowledge sharing
- pairing
- distributed ownership



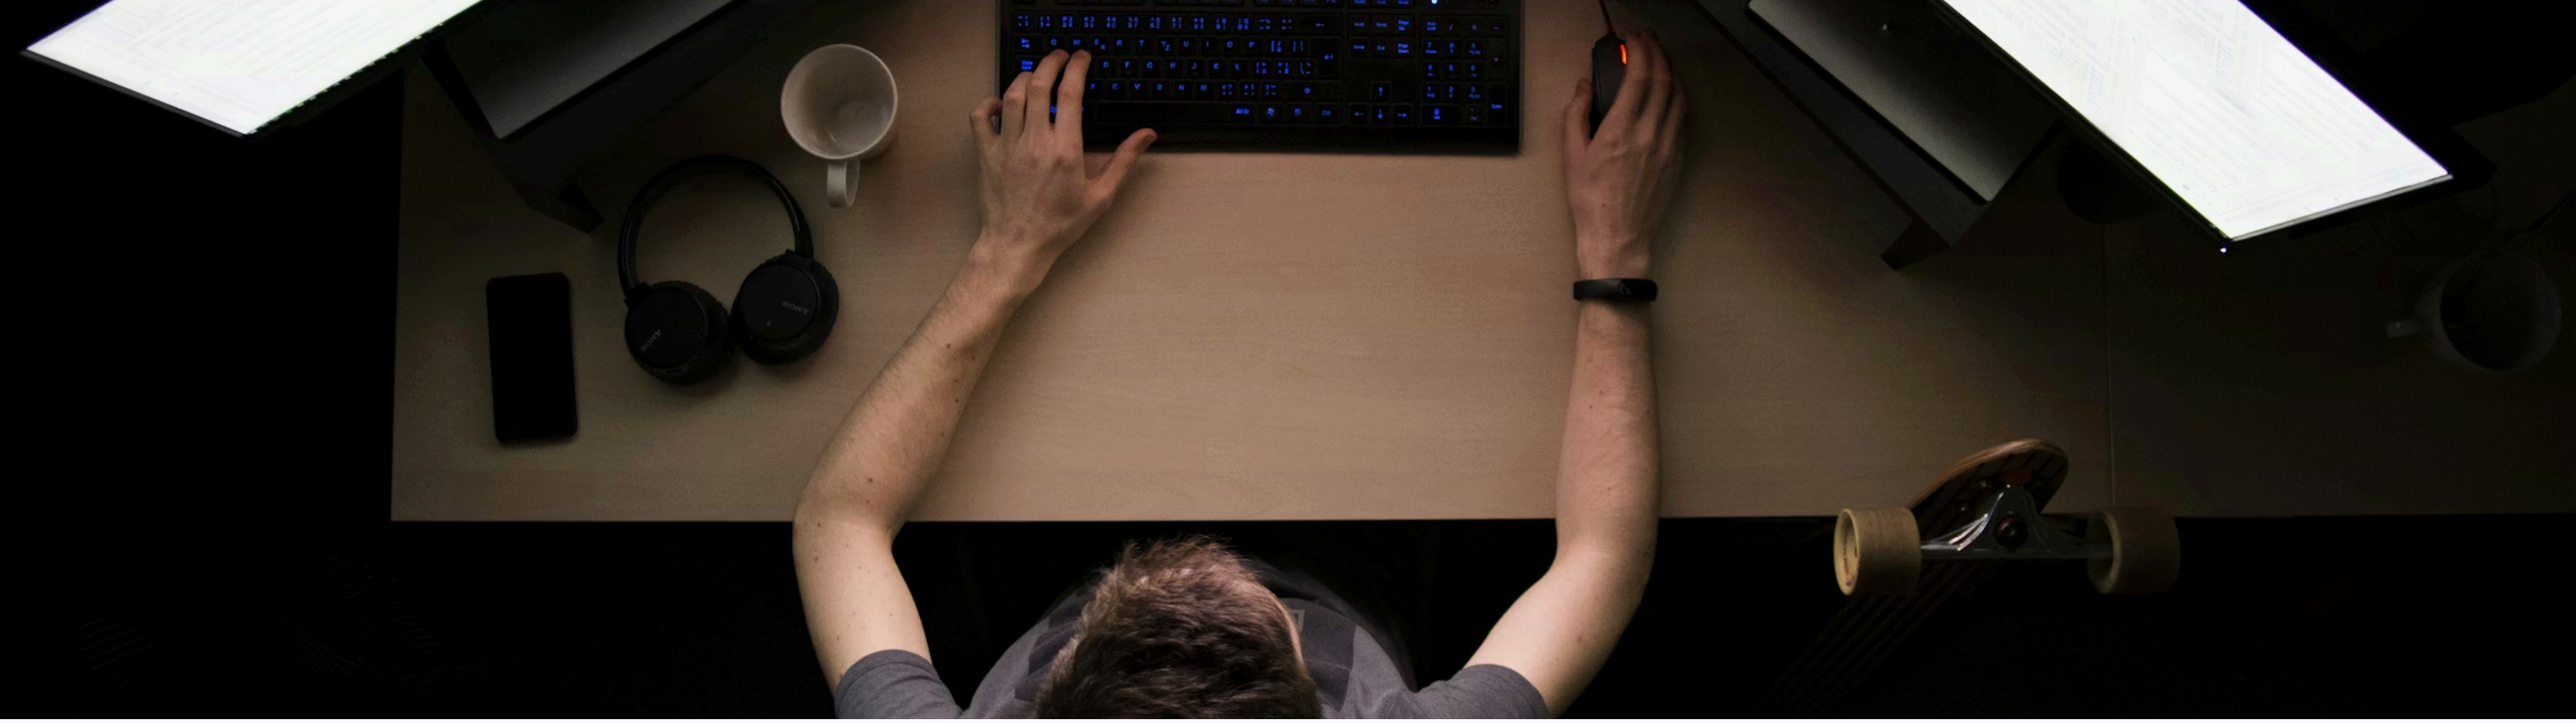
03

Redundancy Mechanisms

Heroes are unacknowledged single points of failure.

Advocate and implement:

- knowledge sharing
- pairing
- distributed ownership



03

Redundancy Mechanisms

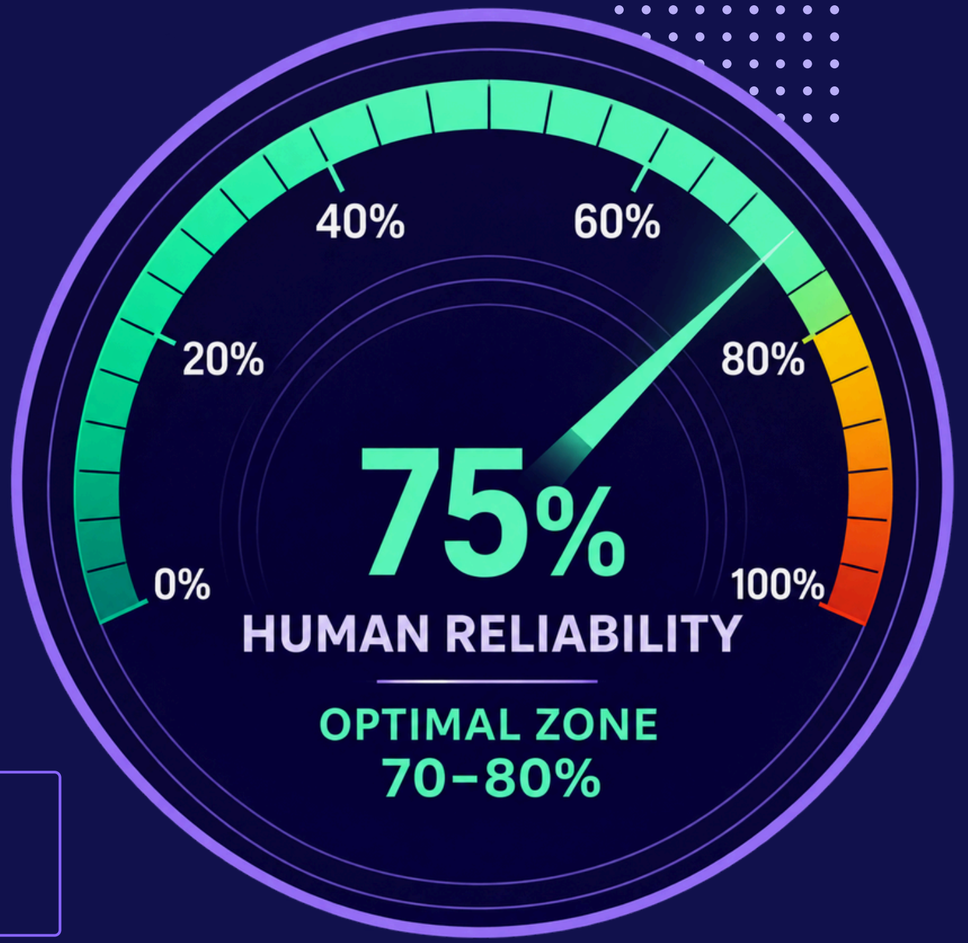
Heroes are unacknowledged single points of failure.

Advocate and implement:

- knowledge sharing
- pairing
- distributed ownership

04

Establish Limits



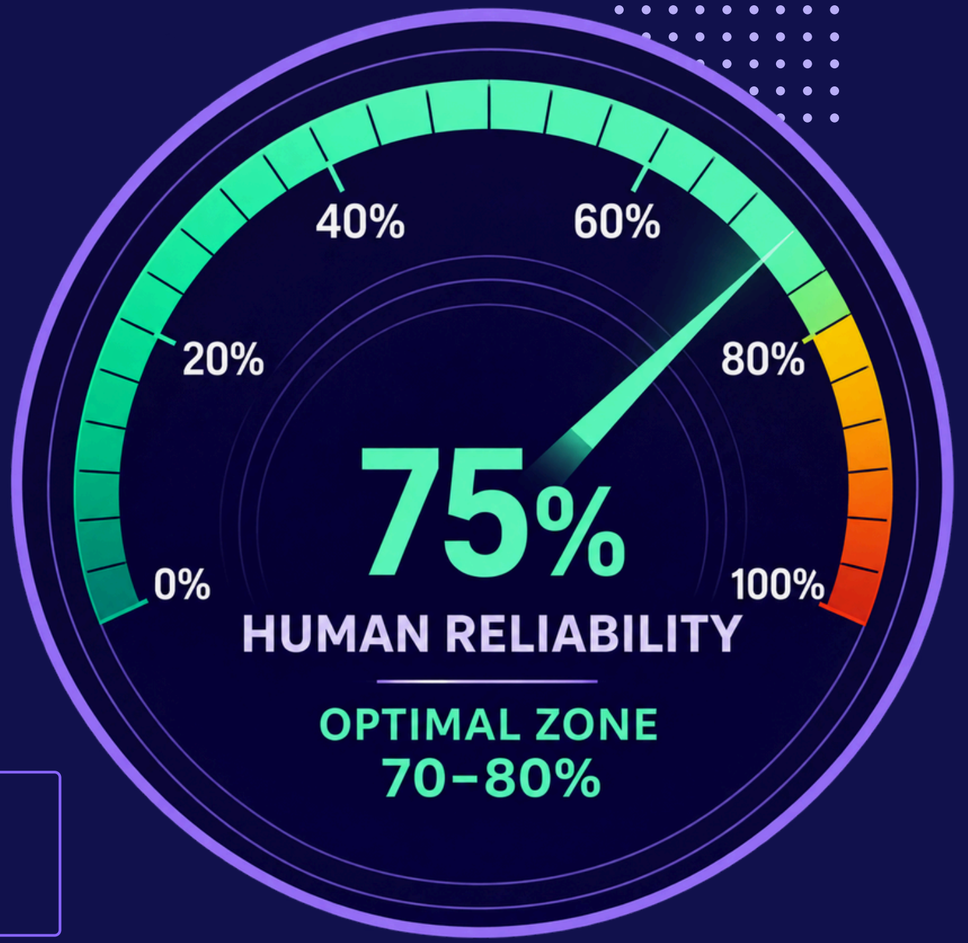
- **LOAD-SHEDDING**
- **COGNITIVE RECOVERY**
- **CIRCUIT BREAKERS**
- **AI THROTTLING**

But what if I’m an IC? … **Start with yourself!**

[4] https://www.riseworkplacewellbeing.co.uk/news/burnout-series-understanding-the-42-rule/
[5] https://github.com/dnanuti/cognitive-observability/blob/main/.claude/skills/observe/SKILL.md

04

Establish Limits



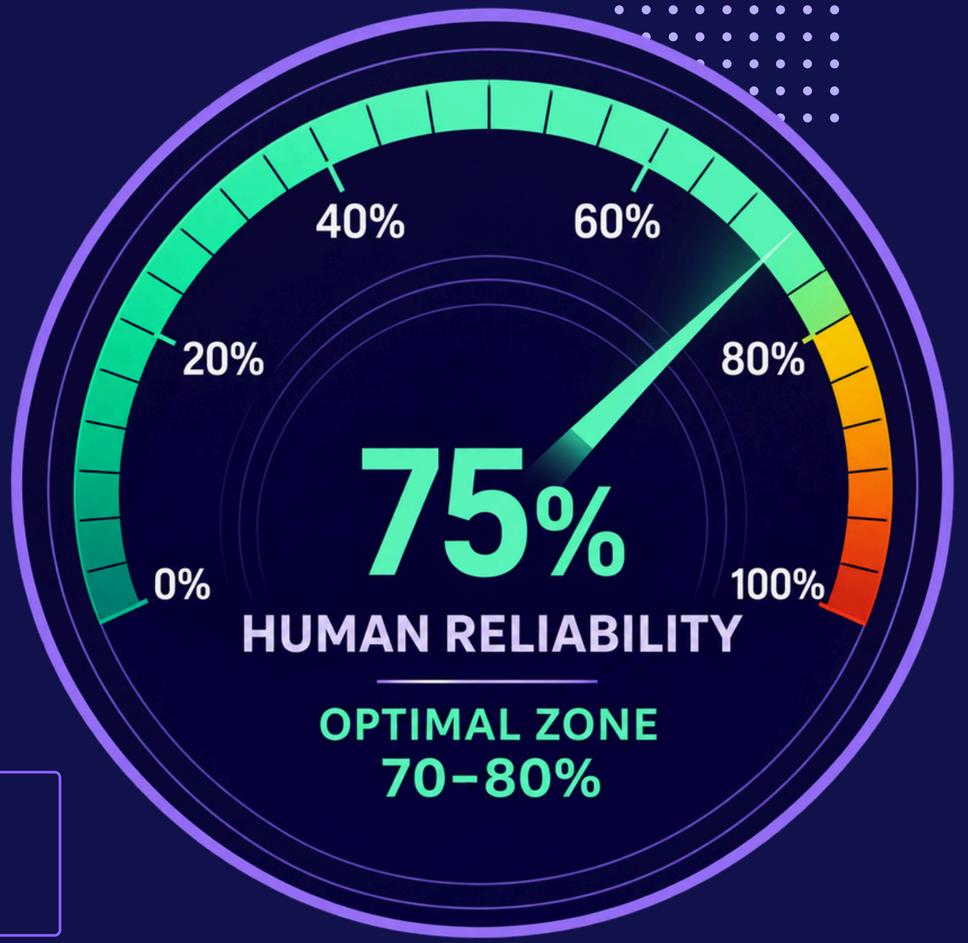
- **LOAD-SHEDDING**
- **COGNITIVE RECOVERY**
- **CIRCUIT BREAKERS**
- **AI THROTTLING**

But what if I’m an IC? … **Start with yourself!**

[4] https://www.riseworkplacewellbeing.co.uk/news/burnout-series-understanding-the-42-rule/
[5] https://github.com/dnanuti/cognitive-observability/blob/main/.claude/skills/observe/SKILL.md

04

Establish Limits



- **LOAD-SHEDDING**
- **COGNITIVE RECOVERY**
- **CIRCUIT BREAKERS**
- **AI THROTTLING**

But what if I’m an IC? … **Start with yourself!**

[4] https://www.riseworkplacewellbeing.co.uk/news/burnout-series-understanding-the-42-rule/
[5] https://github.com/dnanuti/cognitive-observability/blob/main/.claude/skills/observe/SKILL.md

Check-ins (`/battery`)

Command	What happens
`/battery`	Asks how you feel right now on the Samn-Perelli scale (1 fully alert … 7 completely exhausted) and adjusts: 1–4 carry on, 5 suggests a break, 6–7 suggests stopping and writes a "where we are" note
`/battery 4`	Same, with the rating given up front
A number 1–7 in reply to a nudge's "battery check"	Recorded the same way, with source `nudge`
After a rating of 5+	Optionally asks two NASA-TLX items (mental demand, frustration, 0–10) to tell what kind of help fits
`/battery month`	Monthly check: 4 exhaustion questions (last 4 weeks) and 7 optional ones on workload, control, support, recognition, fairness, meaning and work–life. At most once per 28 days
"skip"	Recorded as a skip, never treated as a low score

Every answer is one line in `checkins.md` in Claude's memory folder for the project.

Reflection (`/observe`)

Command	What you get
`/observe`	The current session: retry loops, yak shaving, frustration, stretch length, context switching. No git, no check-in questions
`/observe 24h`, `/observe Nh`	Git history for that window (commit times, churn, reverts, scope scatter, unmerged branches), compared with your earlier snapshots
`/observe 7d`, `/observe 1w`, `/observe Nd` (N ≥ 7)	The above, plus **recovery** against your own 4-week baseline (days without Claude sessions, late and long days, work streak, recovery debt), your check-ins, and at the end either the weekly burnout question or, every 28 days, the monthly check
`/observe 30d` (or longer)	The above, plus **does it match how you feel**: which recovery numbers moved together with your self-reports (shown after ~8 paired weeks)

Output length follows the evidence: a sentence when things look normal, observations plus a concrete next step (split the task, change approach, note where you are) when they don't. Each run saves a short snapshot (no prompt text) so later runs can tell a bad day from a pattern. Outside a git repo, a timeframe run stops and suggests plain `/observe` instead.

Session patterns

Last 84 days (since Aug 17) · generated 2026-11-08 22:02+00:00

Generated locally from Claude Code session logs. Prompt text is never shown or stored here, only counts. Every number is a rough proxy: it can't see breaks away from the keyboard or work outside Claude.

Active hours	Days with no activity	Outside working hours	Longest stretch	Interruptions	Replayed nudges
316.9	**18**	**15%**	**192 min**	**185**	**162**
66 active days	see 7. Recovery	of prompts	median 87 min	1,006 failed tool calls	2.45 per active day

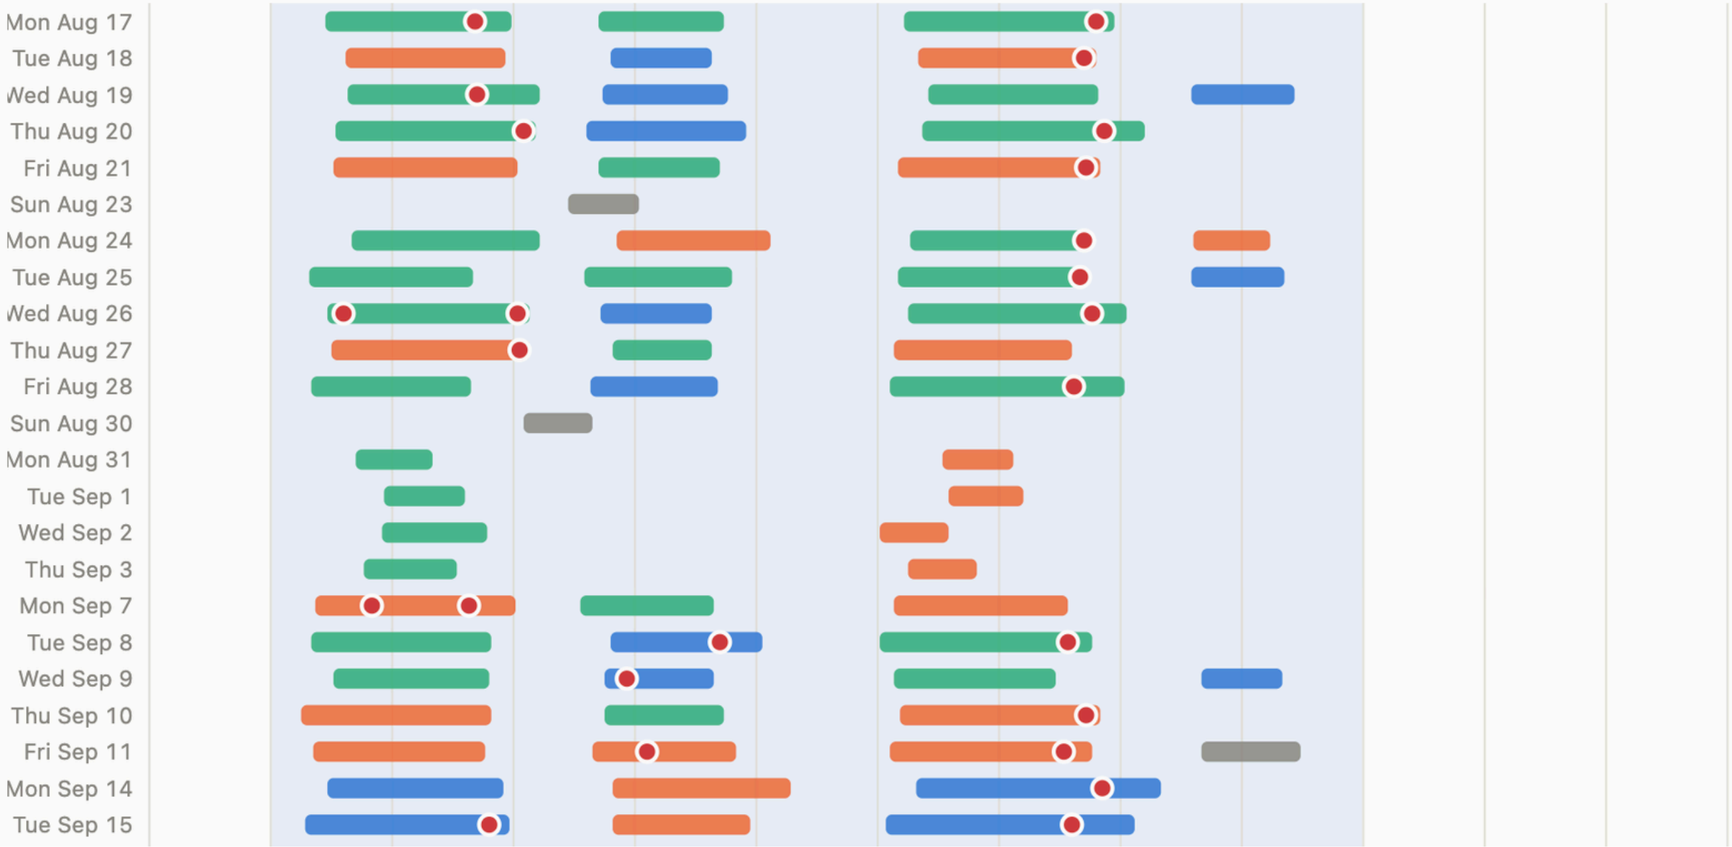
- Busiest day: Wed Oct 7 with 10.0 active hours (08:25–22:49).
- 98 stretches passed 90 min, 12 passed 150 min.
- Counting only gaps between your prompts, the longest stretch would be 190 min instead of 192: Claude working for 10+ minutes looks like a break. The watcher's activity hooks now count both sides.
- 163 prompts were sent on weekends.

2. Breaks and stretches

A stretch is continuous work with no gap of 10+ minutes. These charts test whether the 90 and 150 minute nudge thresholds fit how you actually work.

Day timeline

Each bar is a session's stretch, one row per day; markers are where the watcher would have nudged

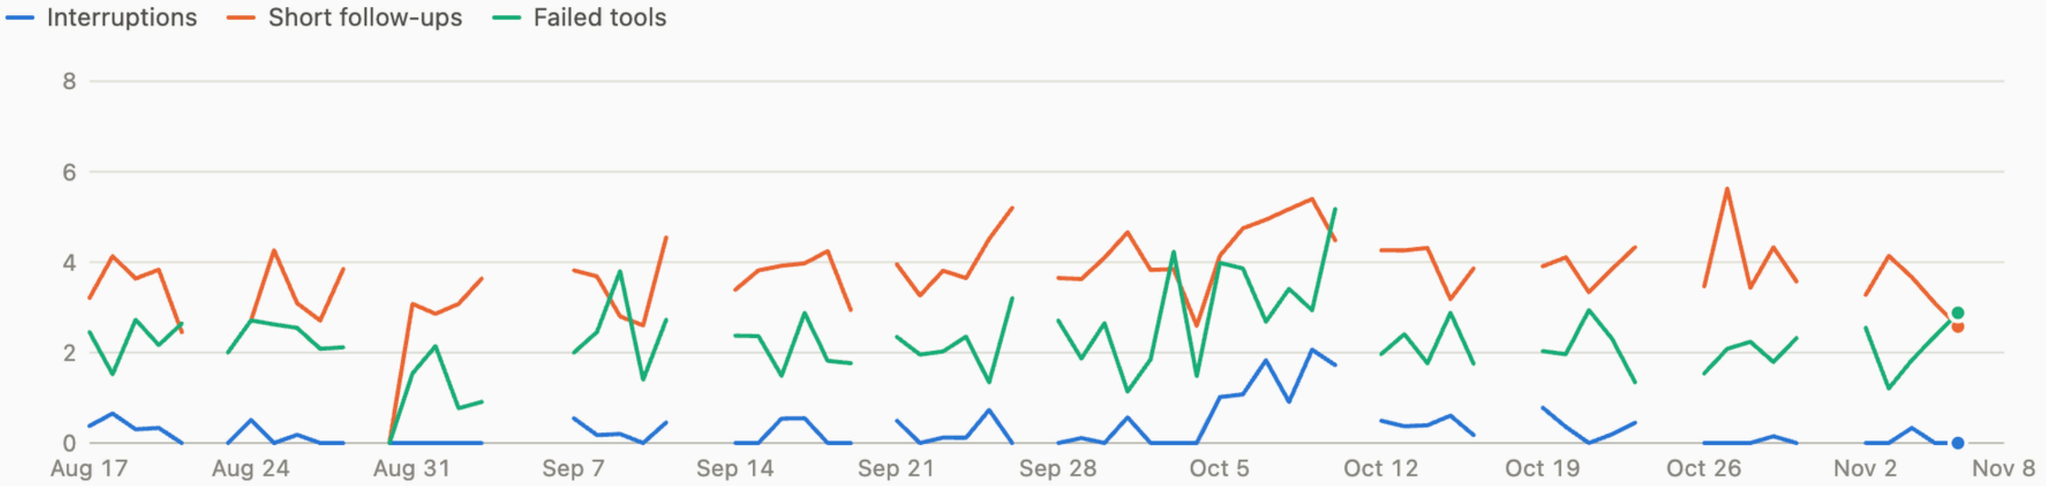


3. Friction

Interruptions (you stopping Claude), failed tool calls and short quick follow-ups are proxies for "that's not what I meant" loops.

Friction per day

Per 10 prompts, so busy days don't look worse just for being busy



Monthly check

Exhaustion 0–100 (mean of 4 items, higher = more often exhausted) · latest Nov 6

Exhaustion

38

previous 63 on Oct 9

Cause area (latest)	Type	0–100
Not enough time for tasks	demand	50
Influence over amount of work	resource	50
Support from colleagues	resource	75
Recognition	resource	50
Work distributed fairly	resource	50
Work feels meaningful	resource	75
Work drains private life	demand	50

05

Takeaway

01.

Burnout is a systems failure, not a people problem.

02.

High performance requires slack, not heroics.

03.

Boundaries work best when the system enforces them.

Thank You

For your time and attention

DESIGN TEAMS LIKE YOU DESIGN SYSTEMS.

LinkedIn

GitHub