

PRIVACY-FIRST TIME INTELLIGENCE

# TimeCampus

See where your time actually goes — captured for you, explained honestly, and never through a screenshot.

AUTOMATIC CAPTURE

SELF-ASSEMBLING TIMELINE

FOCUS AND WELLBEING

WHOLE-OF-LIFE VIEW



---

## 01 The short version

TimeCampus works out where your hours actually went — across the tools you already use, with no timer to start and no screenshot ever taken.

The manual tools depend on a person remembering to press start and stop. The automatic ones buy their completeness with surveillance — and a picture of somebody's screen says nothing about whether real work happened.

TimeCampus records plain intervals from your calendar, inbox, editor and browser, and assembles them into one categorised timeline of the day.

### THE PROMISE

**Insight without surveillance. No screenshots, no keystroke logs, and nothing a machine guessed ever lands on an invoice.**

---

## 02 In numbers

---

### Zero

#### Timers to remember

Activity is captured as it happens.  
No start button to forget and  
no week to reconstruct on a  
Friday afternoon.

---

### None

#### Screenshots or keystrokes

Not a setting that ships switched  
off — a capability the product  
commits never to build, in any tier,  
at any price.

---

### 163

#### Tools for your AI agents

A published Model Context  
Protocol server exposes time,  
projects, meetings, focus, expenses  
and reporting.

---

### 6

#### Roles seeded per workspace

Owner, team lead, client-and-billing  
manager, capture administrator,  
member and read-only analyst.

---

## 03 Who it is for

---

### **Knowledge workers who lose the day to fragments**

The answer to where the hours went, without becoming a bookkeeper — focus ratio, context switches and the longest unbroken block.

---

### **Engineering and delivery managers**

Capacity planning grounded in what the team actually does. Team views are aggregate by design, so you plan without monitoring.

---

### **Freelancers and independent consultants**

Billing that rests on accurate, defensible data: the week is never reconstructed, and every attribution is a decision you made.

---

### **Privacy-conscious organisations and works councils**

The absence of screenshots and keystroke logging is architectural, and the individual keeps per-source consent.

---

### **Agencies and client-services teams**

Client portals, expense approval and project-level time make reporting self-serve instead of a chain of status emails.

---

### **IT and privacy administrators**

Capture configuration, categorisation rules and erasure requests belong to a role holding no approval or client-portal powers.

---

## 04 What breaks today

Every organisation already knows roughly what is wrong with its time data. The interesting question is why the existing fixes make it worse.

# The two failure modes

## The timer nobody starts

Manual tracking is abandoned the moment work gets intense, which is exactly when the data would have been most useful. What survives is a rough guess entered on Friday.

---

Activity is captured as it happens from the tools already open — calendar, inbox, editor, browser — with nothing to press.

## Completeness bought with surveillance

Staff hide work, trust collapses, and the resulting screenshots still cannot distinguish deep work from a stalled afternoon.

---

Only intervals are recorded — an application, a title, a link, a span of time. No pixels, no keystrokes, and no plan to add them.

# Four more things that break

---

## A day scattered across a dozen tools

Nothing knows about the day. Every source now feeds a single ingest path and lands on one chronological timeline.

---

## Burnout only visible after the damage

Work-pattern signals flag meeting overload, interruptions, low focus and too few breaks — each with the observation behind it.

And the two-thirds of your time no work tool touches — the commute, the queue, the evening — gets a separate, self-owned personal layer.

---

## Deep work crowded out by the calendar

Focus is scored daily and tracked as a trend, paired with a Pomodoro timer, so protecting it becomes a visible decision.

---

## Numbers you cannot interrogate

Every score reports the share of the window it could account for, and returns nothing rather than a fabricated figure.

---

## 05 What you get

One product with nine working parts. The first three make the record honest; the next three make it useful at work; the last three make it yours.

---

## Parts 1-3 **Making the record honest**

### CAPTURE

#### **One ingest path for every source**

Browser, command line, editor, calendar and email feed normalised batches through a single endpoint. Repeat submissions are ignored, so a client re-sending yesterday cannot inflate the week.

### TIMELINE

#### **A day that assembles itself**

Events are categorised into meetings, deep work, communication and breaks, then merged with time entries, meetings, focus sessions and email activity into one chronological view.

### ATTRIBUTION

#### **Naming the project behind each hour**

Your project names and rules score first; a model may refine. Whichever wins is a suggestion carrying its confidence and method — never written onto the billing hour until you accept it.

---

## Parts 4–6 Useful at work

### FOCUS

#### **Focus scoring and Pomodoro**

A daily focus score with a trend over time, plus a built-in Pomodoro timer with sessions, goals and settings. Adaptive distraction blocking is specified but not yet built.

### WELLBEING

#### **Honest overload signals**

Work-pattern insights and weekly recommendation cards read sustained workload early. Individual signals stay private; team signals are aggregate. Predictive burnout modelling is roadmap.

### DELIVERY

#### **Projects, approvals, expenses, portals**

Projects and tasks with budgets, a timesheet approval queue, an expense workflow, and read-only client portals whose link can be switched off. Native invoicing is planned.

---

Parts 7-9

# Making it yours

## PLANNING

### Team capacity and resource plans

Resource plans, bulk allocations across people and a workspace-wide capacity overview — planning capability most single-purpose trackers do not attempt.

## LIFE

### The whole of your finite time

Life profile, values, goals, balance targets, life timeline, life-in-weeks, a time-well-spent alignment score and a dead-time radar. No management role in the product can read it.

## EXTEND

### Channels, applets, widgets, workflows


Channels bring activity in, applets add focused functionality, widgets compose the dashboard each role needs, and workflows automate routine handling by rule.

---

## 06 Put to work

Situations the product was built around, told from the point of view of the person who has the problem.

# Getting the day back




KNOWLEDGE WORKER

## See where the day actually went, without a single timer

Activity from the tools you already had open is organised into a categorised timeline. You open it next morning and read the day back.

---

Awareness that costs no administration, and a record you did not have to maintain.



FREELANCER

## Plan a realistic day and then defend it

Reserve focus blocks for the work that matters, and see where reality is drifting from the plan while there is still time to act.

---

A busy day that still ends with the important things finished.

# Focus, and the overload behind it




SOFTWARE ENGINEER

## Make deep work something you can actually see

Context switches and the longest unbroken block are measured explicitly, and the pattern insights name interruption frequency as a warning.

---

An argument for uninterrupted time that rests on your own data.



TEAM LEAD


## Notice overload while it is still reversible

Meeting overload, too few breaks, a drop in focus. Individual signals stay private to the individual; what a lead sees is aggregate.

---

A chance to rebalance before exhaustion, without a performance file.

# Client work that reports itself




AGENCY OWNER

## Replace status emails with a client portal

Each client gets a read-only portal for their own projects, reached by a share link you can switch off or regenerate at any moment.

---

Transparency the client can serve themselves, and hours back for delivery.



CONSULTANT

## Produce a trustworthy time report in seconds

Because activity was captured as it happened and approvals already ran, summaries are generated on demand and grouped by project or client.

---

Reports that survive a client asking where a figure came from.

---

## 07 **AI that reads your day, not your screen**

A captured hour arrives as an application, a window title, a link and a span of time, and something has to decide what it was. Naming the project behind each hour is the entire difference between a log and an answer.

# The order of operations

Deliberate, and unusual. Each step can only ever narrow what the next one is allowed to do.

## 01

### Your rules score first

Your own project names and category rules are matched and scored. Below a set bar the product offers nothing rather than guessing.

## 02

### One model, one attempt, twelve seconds

Only if you configured a provider. A timeout, a garbled reply or an invented project name is discarded and the rule-based answer stands.

## 03

### Then it stops, and you decide

The label sits beside the hour with its confidence and model name. Accepting copies it across. Setting the project yourself marks the record as your decision.

An AI system whose most important design property is the thing it is not allowed to do.

# Four AI design rules

---

## **A log of hours is not an answer**

Working out the project and the life area behind each interval is the only output worth reading, so inference is the product rather than a garnish on it.

---

## **Scored on balance, not on output**

The time-well-spent score compares the shape of your week against the shape you asked for, so working longer cannot raise it. A calm week can beat a frantic one.

## **Every label says who decided it**

Rule, model or you — plus which model and how confident. A machine guess is never mistakeable for a human decision, in either direction.

---

## **Switch a source off and nothing is kept**

Consent is checked in the capture path before anything is stored. Withdraw a source and the batch is refused outright, and you are told how much was refused.

# What the AI can actually do

---

## Plain-rule project labelling

Your project names and rules label each hour with no AI provider configured at all.

The baseline behaviour is silence rather than noise.

---

## Time-well-spent alignment score

How your week actually split against how you said you wanted it to split.

With too little coverage it is flagged as an estimate, not presented as fact.

---

## Provenance on every label

Who decided — rule, model or you — plus the model name and the confidence.

You can audit your own timesheet rather than trusting it.

---

## Dead-time radar

Waiting, commuting, unlabelled activity and untracked gaps are separated and totalled.

Sleep is never a bucket, so rest is never presented back as waste.

---

## An assistant that knows your screen

Ask about this project or this meeting without naming it, on every signed-in screen.

It receives a short fixed list, and nothing about your device or browsing.

---

## The never-work-more guard

Every candidate suggestion declares whether its net effect is more work.

Those are dropped before they are saved — the suggestion never exists.

# And what keeps it in bounds

## **Consent is a gate, not a settings page**

The capture path checks per-source consent before a single event is stored, and the refusal is counted and reported back to you.

## **Permissions are checked before the AI, not by it**

An employer or workspace account cannot reach your personal analysis, and cannot even confirm that it exists.

## **No guess ever lands on an invoice**

Only your acceptance or your own edit writes the project that drives billing, and your edit is marked as yours for good.

## **The model may only choose from your own list**

If it returns an invention, a malformed reply or nothing at all, the answer is discarded and the rule-based label stands.

## **Nothing runs unattended**

Generating suggestions is a button you press. A shared record of model calls is committed before anything runs on its own.

## **Model usage is billed to you, not bundled**

AI categorisation runs on your own provider key, which keeps your staff data inside your own provider relationship.

STATED PLAINLY

## What is not built yet

The accept and reject controls are not on the timeline, so that review loop currently runs through the API. Corrections fix the hour you corrected but do not yet teach the next label.

There is no shared record of model calls and no fixed evaluation set, which means a drop in answer quality would be noticed by you rather than caught by us.

AI usage is counted but not capped. Switching a source off stops new capture but does not retroactively remove what was already stored. Natural-language questions over your time data are roadmap.

---

## 08 Who does what

---

### Workspace owner

Owns the time estate end to end — capture governance, approvals, expenses, portals, capacity planning — and holds the deletion authority behind retention and erasure requests. Reads aggregate team summaries only.

---

### Team lead

Approves timesheets and expenses, owns resource plans, reviews aggregate analytics. Raw capture, timeline, application usage and focus scores are withheld from this role on purpose.

---

### Client and billing manager

Runs client-facing delivery: portals and share links, billable projects and time, expense reimbursement. Reads approvals and capacity but cannot approve or plan.

---

### Capture administrator

The IT or privacy specialist running ingestion: capture channels, application categorisation, email hygiene, geofences and erasure. Holds no approval or expense powers.

---

### Member

Their own timeline, time entries, focus sessions, submitted timesheets and expenses — and the only role in the product that holds any personal life-time data at all.

---

### Analyst

Produces utilisation, billable-ratio, meeting-cost and capacity reporting and changes nothing. Every permission is read-only, with no raw capture and no client-portal access.

---

---

## 09 Under the hood

TimeCampus is API-first. Every client — web, mobile, desktop, extensions, the command-line tool, the agent server and the SDKs — talks to the same federated GraphQL API. There is no privileged path.

Authorisation is decided at the API edge and never inside the product. Every read and write declares the permission it requires, and the gateway enforces it before a request reaches any product logic.

The product is a domain service on a shared multi-tenant platform, so identity, roles, billing, audit, notifications and files are inherited rather than rebuilt. Captured activity flows over an event stream.

### THE SPLIT THAT MATTERS

The individual's whole-life estate is enforced as self-scoped only. The workspace owner's permission set omits every one of those resources, and that omission is the access control. In the other direction, raw work patterns and location rows are owner-filtered, and management reaches only aggregate summaries carrying no coordinates.

# Capability by area

---

## Capture ingest

A single endpoint takes normalised batches. Consent is checked before storage, refusals are counted, and repeat submissions are discarded.

---

## Analytics and scoring

Focus scores with trend, work-pattern insights, an alignment score reported with its coverage, and reclaimable-time analysis that excludes sleep.

---

## Timeline assembly

Auto-categorised events merged with entries, meetings, focus and email. Overlapping sources are reconciled rather than added together.

---

## Professional-services layer

Projects with budgets, meetings with action items, approval requests, expenses, revocable client portals and resource plans with allocations.

---

## Attribution pipeline

A deterministic keyword classifier with a confidence floor, and optional model refinement bounded to one attempt and twelve seconds.

---

## Inherited platform services

Single sign-on, role-based access control, multi-tenancy, quotas, an exportable audit timeline, notifications, file storage, tagging and accessibility.

# Standards it speaks, and what it plugs into

## OPEN STANDARDS

- OpenID Connect single sign-on across every client surface
- OAuth 2.0 device-code flow for terminal and agent clients
- SAML and SCIM for enterprise identity and provisioning
- GraphQL as the single product API, one schema behind every client
- Model Context Protocol, so any compatible agent operates under your permissions
- Webhooks, CSV export, TLS in transit and encryption at rest

## PLUGS INTO

- Calendars — Google, Outlook and Apple, for meeting and focus-block context
- Email — measured from message metadata rather than message content
- Editors — VS Code and JetBrains, without reading file contents
- Browsers — a Chromium extension that tags usage and holds a focus session
- Chat and project tools — native connectors are roadmap; the open capture endpoint covers them today
- Anything else — normalised activity batches from your own scripts and jobs

# How it deploys

## Multi-tenant cloud

The default shape. Your workspace is isolated by tenant boundary, with identity, roles and quota enforced at the API edge.

## Dedicated tenancy

An enterprise option for organisations needing separation from the shared estate, alongside multi-entity structures.

## Self-hosting

An enterprise deployment choice for organisations whose policy requires the product to run inside their own estate.

## Regional data residency

An enterprise control over where your time data is stored, sold as an option rather than assumed.

## Bring your own model provider

AI categorisation runs on your own key or subscription. With no provider configured, rule-based labelling still works.

## Offline-tolerant capture

Clients collect while disconnected and deliver when connectivity returns, and the ingest path discards repeats.

---

## 10 Security & trust

---

### No screenshots, no keystroke logging

Architectural, not a default setting. There is no capability to enable and no plan to build one, in any tier. Idle-time policing is likewise absent.

---

### Per-source consent checked before storage

Revoke a source and new activity from it is refused outright, with the volume of refusals reported back to you.

---

### Aggregate-by-default team views

Raw work-pattern and location rows are owner-filtered, so a single-row lookup returns nothing rather than a colleague's data.

---

### Client-portal secrets treated as credentials

A share link is a bearer credential, so reading a portal is deliberately withheld from the read-only analyst role.

---

### Authorisation decided at the API edge

The product never validates tokens and never re-implements access control. The gateway decides allow or deny before any product logic runs.

---

### The personal estate is self-scoped only

Fourteen life-time resources are enforced as belonging to the individual. No management role holds any of them.

---

### Separation of duties by construction

Capture governance cannot approve timesheets. Client portals cannot plan capacity. Submitting and approving are different permission levels.

---

### Exportable audit and activity trail

Every action is recorded with actor, target and result, searchable and exportable. Audit streaming is an enterprise roadmap option.

STATED PLAINLY

## Where the boundaries are still coarse

TimeCampus is pre-launch and holds no third-party security certifications today. GDPR alignment, SOC 2 principles and ISO 27001 practices describe what the product is being engineered toward — they are not attestations.

Withdrawing consent for a source stops new capture but does not retroactively remove what was already stored.

Because a person must be able to record their own activity, that grant also permits reading a colleague's raw capture within the same workspace. Narrowing it is committed work.

---

## 11 One product, many front doors

### Web application

The full experience, and the surface every other client is measured against.

### Embedded product surface

The same experience composed into the wider workspace shell rather than a separate tab.

### Agent server

163 published tools, each acting as you and passing the same permission checks your clicks do.

### Mobile applications

In active development. Automatic capture is opt-in on Android and foreground-only on iOS.

### Desktop, browser and editor

A background agent, a Chromium extension and a VS Code extension. In active development.

### CLI and SDKs

Terminal sessions and reports, plus Node and Python SDKs. In active development.

### GraphQL API

The same federated API the web application uses, and where the accept and reject loop runs today.

### Documentation and community

Product guides, API reference and tutorials, with a community hub planned alongside the launch.

### Status, honestly

The web experience leads. Much of the surrounding ecosystem is in active build ahead of the 2026 launch.

## 12 Plans

Priced by the size of the team you manage — employees plus contractors — not by how many administrators log in. Every plan includes the whole platform; what changes is capacity, history, support and governance.

|                                                                                                                                                                         |                                                                                                                                                                                          |                                                                                                                                                                                                                |                                                                                                                                                                                                     |                                                                                                                                                                                                                                             |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>FREE</p> <p><b>\$0</b></p> <p>Trying TimeCampus. No card, no expiry.</p> <hr/> <p>15 managed profiles</p> <p>50 time entries per month</p> <p>30 days of history</p> | <p>STARTER</p> <p><b>\$99/mo</b></p> <p>Small businesses up to around 50 people.</p> <hr/> <p>50 profiles, 5 admin seats</p> <p>500 time entries per month</p> <p>90 days of history</p> | <p>GROWTH</p> <p><b>\$399/mo</b></p> <p>Scaling companies up to around 250 people. 14-day trial.</p> <hr/> <p>250 profiles, 3 workspaces</p> <p>5,000 time entries per month</p> <p>365 days, 99.5% uptime</p> | <p>BUSINESS</p> <p><b>\$999/mo</b></p> <p>Up to around 1,000 people, with an SLA.</p> <hr/> <p>1,000 profiles, 10 workspaces</p> <p>50,000 time entries per month</p> <p>730 days, 99.9% uptime</p> | <p>ENTERPRISE</p> <p><b>From \$12/member/mo</b></p> <p>Larger or regulated organisations. Annual contract.</p> <hr/> <p>Production SSO and provisioning</p> <p>Audit streaming, data residency</p> <p>Dedicated tenancy or self-hosting</p> |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Planned launch prices, before tax. Annual billing reduces each tier; four regional bands set the price by billing country. Web billing is the canonical price.

# What the price does and does not include

## **Add-ons raise one limit, not the whole plan**

Capacity packs cover profiles, projects, focus seats, storage, API volume and integrations. Feature add-ons cover deeper AI throughput, extended analytics, longer audit history and portal branding.

## **Model usage is not included in the plan price**

AI categorisation, summarisation and report drafting run on your own provider key. TimeCampus governs the workflow and the limits; your provider bills you for the model.

## **Regional pricing is the normal price, not a discount**

Four bands set the price by billing country. Growth is \$399/mo on the standard band and around \$279, \$199 and \$139 on the others.

## **What happens at a limit, and on the way down**

At the limit new usage may be blocked until you upgrade or reduce usage; existing data stays accessible. On downgrade, data above the new limit becomes read-only rather than being deleted.

---

## 13 Why teams choose it

---

### Whole-of-life time intelligence, already built

No incumbent in the time-tracking category offers whole-of-life intelligence at all, which makes this the most defensible difference and the least imitable one.

---

### The machine proposes and the human applies

A suggestion sits beside the hour, never on it — enforced in the data model rather than in a warning dialogue.

---

### Privacy is structural, not a stance

Team views are aggregate by construction, so the boundary survives a change of management, a change of policy and a bad day.

---

### Numbers that state their own coverage

Scores report the share of the window they could account for. Interrogability is a feature, not a debugging aid.

---

### Capture without a daily approval ritual

A normalised ingest path fed by browser, command line, editor, calendar and email. No start button and no daily confirm step.

---

### Programmable to an unusual depth

163 agent tools, Node and Python SDKs, a CLI, extensions and the same GraphQL API the product itself uses.

# Honest answers to hard questions

---

**You are pre-launch. Why would we bet on that?**

---

You should weight it accordingly. What is real is documented capability rather than promise, and where something is roadmap this document says so on the same page — which is the only thing that makes the built claims worth anything.

---

**Automatic capture sounds like monitoring with better manners.**

---

The distinction is testable. There is no screenshot capability, no keystroke logging and no idle-time policing to switch on. If somebody wants to watch a specific person's screen, this product is genuinely unable to help them.

---

**Can an AI misattribute our billable time?**

---

It can produce a wrong suggestion; it cannot produce a wrong invoice. Only your acceptance or your own edit writes the project that drives billing, and the model can only choose from your existing project names.

---

**The integration catalogue is thin next to the incumbents.**

---

It is, and that is currently the honest weakness. Chat, documents and project-tool connectors are not built; those sources arrive through the open capture endpoint today. If you depend on a specific connector, check it before committing.

---

# Where it sits against the alternatives

## VS MANUAL TIMER TRACKERS

If you want a dead-simple manual timer, or mature profitability reporting today, they win. TimeCampus is for the case where the timer itself is the problem.

## VS AUTOMATIC TIMESHEET TOOLS

The closest category, and their drag-and-drop daily timeline is better than ours today. We add focus scoring, overload signals, a programmable platform and the life-time layer.

## VS DESKTOP FOCUS TRACKERS

On pure focus coaching they are ahead — our adaptive blocking is planned. TimeCampus wins on breadth: every device, billing-grade project data and capacity planning.

## VS EMPLOYEE MONITORING SOFTWARE

No overlap. If your requirement is screenshots, keystroke logs, GPS or idle policing, TimeCampus will never satisfy it. What you get instead is data staff have no incentive to game.

And against spreadsheets and calendar archaeology: the record is reconstructed rather than observed, so it is least accurate exactly when it matters most.

START TODAY

# Join the waitlist before the 2026 launch.

Early access and founding-member pricing go to the list first, and the fastest way to judge the product is to read what it admits it cannot do yet.

---

JOIN THE WAITLIST

[timecampus.com/waitlist](https://timecampus.com/waitlist)

---

HOW THE AI ACTUALLY WORKS

[timecampus.com/ai](https://timecampus.com/ai)

---

PRICING

[timecampus.com/pricing](https://timecampus.com/pricing)

---

TALK TO US

[contact@timecampus.com](mailto:contact@timecampus.com)

---



A product of Burdenoff Consultancy Services Pvt. Ltd. Pre-launch: capabilities described as planned or roadmap are not yet available. Prices are the planned launch configuration, before tax.