



Your AI operating team for every back-office workflow

Made with Love by Intelsense AI and RunAgent

Operational Challenges

Common issues in manual back-office workflows

01

Slow communication

Information moves slowly across teams.

02

Limited scalability

Manual work becomes harder to scale.

03

Error risk

Repetitive tasks increase mistakes.

04

Unclear ownership

Process handoffs are not always clear.

05

Higher costs

More time and admin effort are required.

06

Compliance exposure

Delays and weak records create risk.

The Altercrew Ecosystem

A digital workforce, coordinated by the Head Agent

YOUR ORGANIZATION

Founders

Set direction and policy

Managers

Review and approve work

Employees

Request work, give feedback

one conversation

Head Agent
central coordinator

delegates & supervises

DIGITAL WORKFORCE

Marketing
agent

Payroll agent

Meeting
agent

Onboarding
agent

Finance
agent

+ any role

Sub-agents work like employees,
each with a defined role, permissions,
and boundaries.

Sub-Agents, Built in Minutes

No code, no consultants. Describe the role and Altercrew does the rest.

01

Describe the role

Plain language is enough: "a payroll agent that runs monthly payroll and flags anomalies."

02

Scope it

Pick its tools, set permissions, and name the human who approves its work.

03

Deploy

It joins the workforce immediately, governed by policy from its first action.

New Sub-Agent

2 min setup

ROLE

Payroll agent · runs monthly payroll

TOOLS

Gmail

Sheets

Slack

+ 500 more

PERMISSIONS

Read

Draft

APPROVER

Finance lead

Deploy agent

Payroll
agent · live

Meeting
agent · live

+ your next
role

Market Opportunity

Back-office work is the largest under-automated cost center in every company.

Companies already outsource this work at global scale. Altercrew replaces that spend with a governed digital workforce that never sleeps.

\$328B

Back-office outsourcing

Global BPO market value in 2025, work companies already pay others to do.

\$696B

BPO spend by 2033

Back-office spend keeps compounding at ~10% per year.

\$50B

AI-agent market by 2030

Growing at ~46% per year across the enterprise.

How Altercrew Works

Every task follows one governed loop



Approved work ships. Unapproved work waits. **Nothing is sent without sign-off.**

The Head Agent

The Head Agent manages the entire digital workforce, the way a chief of staff manages a team.

ONE COORDINATOR

Head Agent

The single point of coordination, and the single point of trust.

Manages

Delegates

Reports



Monitors every sub-agent's activity and performance



Delegates tasks to the right agent for the job



Communicates with founders, managers, and employees



Requests approval whenever human judgment is required




Controls access to every external tool and channel

Specialized Sub-Agents



A role library for the digital workforce: each agent with defined responsibilities, permissions, and boundaries



Every role operates through the same secured boundary: same policies, same approvals, same audit trail.




RunAgent Superbrowser

Specialized sub-agent

A real browser agent, not just web search

- **Acts like a person online.** It clicks, types, scrolls, fills forms, and navigates real websites, end to end.
- **Shared by the whole workforce.** Every sub-agent can hand it browser tasks; the Head Agent distributes the work and brings back the result.
- **Governed like everything else.** Runs behind the same boundary: policy-checked, approval-gated for sensitive actions, and every step logged.




Human-in-the-Loop Approvals

Every approval routes through your real organogram, straight to the right person

When Altercrew asks for input



Approve, reject, or comment in one tap. If the named approver is unavailable, the request falls back up the chain, so a decision is never lost.



Each cohort's approval requests surface to its own lead.
founders see strategy, not every LinkedIn post.

Communication Channels

Altercrew meets your team where it already works



The organization gives the Head Agent a dedicated enterprise account — every integration connects through it, and every message stays in your existing tools.

Secure Tool Access

Sub-agents never hold credentials. External access stays centralized in the Head Agent

500+

connectable tools & integrations



Sensitive credentials stay in one governed place, reducing exposure across the whole workforce.

Workflow Example: LinkedIn Post

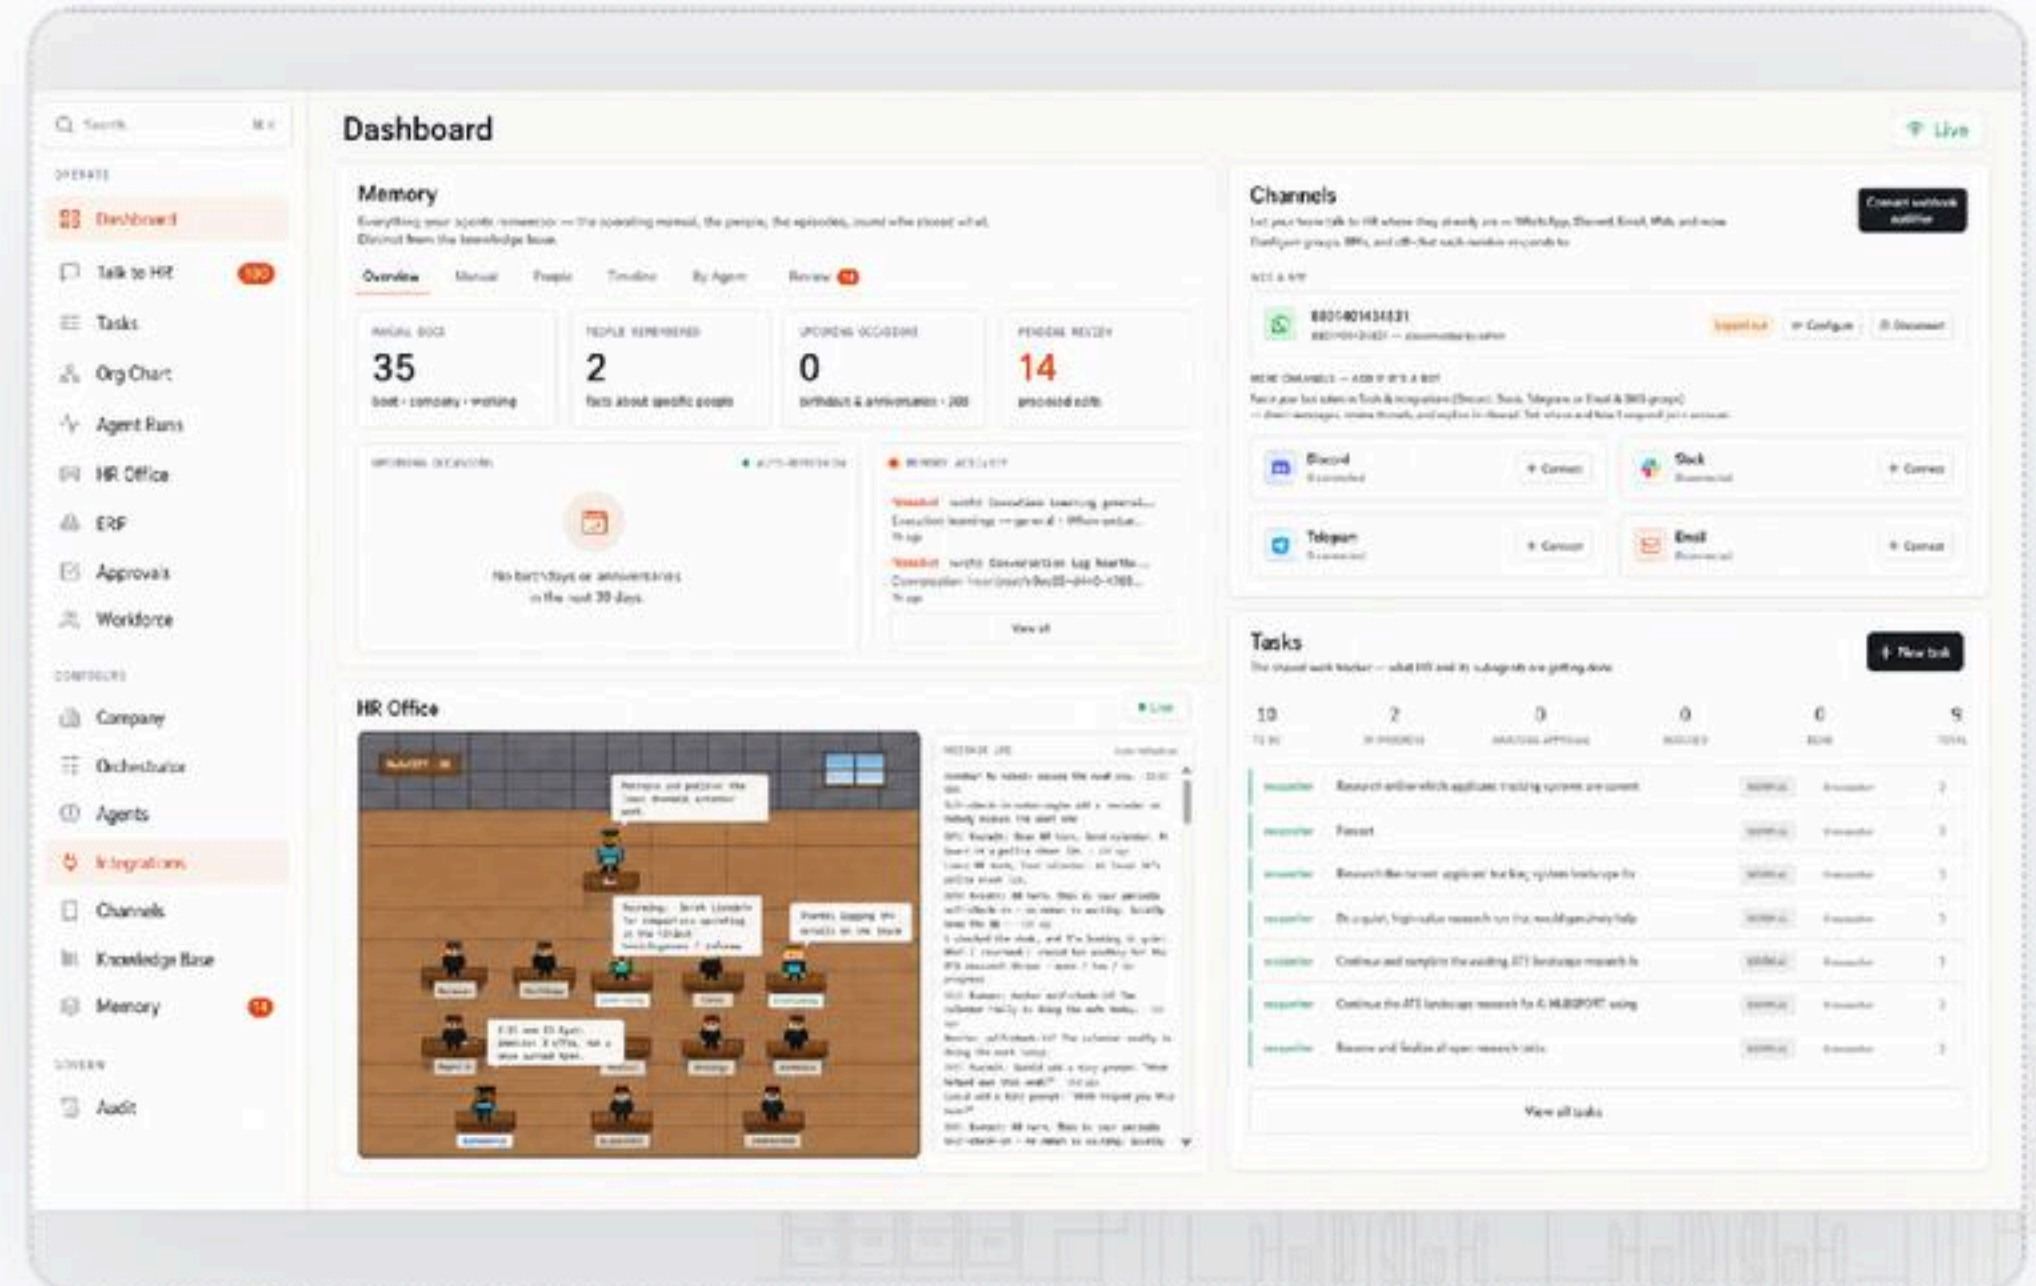
From request to publication, every step governed



Admin Dashboard

The control plane for founders and admins


- Configure the Head Agent and create sub-agents
- Define roles, responsibilities, and permissions
- Connect external tools and channels
- Set org hierarchy and approval workflows
- Review activity, security, and operational logs



Employee Dashboard


The everyday workspace for employees, just like working with a real team

- Ask the Head Agent anything, like messaging a colleague
- It delegates to the right sub-agent behind the scenes
- The finished work comes back reviewed and approved
- Notified only when your input or approval is needed
- Task status and escalations at a glance, like a real team



The Execution Boundary

Agents propose. Systems enforce.



- Agents hold only descriptions of tools, never the tools themselves.
- A denied or approval-gated action makes **zero** external calls. A tested property, not a prompt.
- Agents never see a token or password. Secrets resolve inside the runtime, microseconds before use.

Policy Engine

Deterministic rules the agent can never edit or see around

FIVE POSSIBLE VERDICTS · MOST RESTRICTIVE WINS

1 Deny action never happens

2 Require approval waits for a named human

3 Redact sensitive data scrubbed first

4 More context "tell me why first"

5 Allow proceeds, and is logged

Default deny

If no rule matches an action, the answer is no. Safety is the starting position, not an add-on.

Rules live in configuration, not prompts

Policy loads from files the agent never touches. It cannot talk its way past a rule or rewrite its own permissions.

Approval binds to the exact content

Each approval is fingerprinted to the exact payload a human saw. Change one word after sign-off and the approval no longer applies.

Approver identity is re-checked

On resume, the system re-verifies who approved, their role, and the content fingerprint. Server-side, every time.

Always-On Autonomy

Not a chatbot waiting for prompts. A coordinator that notices work and starts it

A heartbeat, not a prompt

24/7

The Head Agent wakes on its own pulse, scans for birthdays, overdue tasks, new hires and waiting approvals, and assigns work without being asked.



Work that survives anything

Long tasks are durable runs: they sleep on a timer, park while waiting for a human, and resume days later, across restarts, on any machine.



Budgeted and pausable




A daily action budget, quiet hours at night, and an operator kill-switch that stops everything instantly.



Exactly once

An idempotency ledger guarantees an approved send happens exactly once. Crash, resume, or replay, it is never duplicated.

LIFE OF A DURABLE RUN



Security & Isolation

Enterprise-grade guarantees, engineered in from day one



Firecracker microVM isolation

Each agent workload runs in its own secure micro virtual machine, the same isolation technology behind AWS Lambda. Agents talk only through controlled internal channels.



Centralized credentials

Tokens and passwords live in a vault no agent can read. The system resolves a secret inside the isolated runtime, microseconds before the call, and never hands it to a model.



Tamper-evident audit chain

Every proposal, decision, approval, and effect is hash-chained: alter or drop one record and the chain breaks. Who did what, why, and who approved it is always answerable.



Permission-based everything

Roles, scopes, and operational boundaries per agent; disabled tools simply do not exist in an agent's world, so there is nothing to bypass.

Trust is structural, not promised: the agent **cannot** do what it shouldn't, because the code path doesn't exist.

Modular ERP Architecture

BUILT & RUNNING

A headless ERP is already implemented, each module operated by its own cohort of sub-agents

Head Agent

central coordination layer: hierarchy, delegation, approvals, visibility across every module

Marketing

Content · social ·
campaigns · research ·
reporting

Finance

Invoicing · payroll ·
reconciliation · close

Human Resources

Recruiting · onboarding ·
records · policy

Sales

Pipeline · outreach ·
proposals · CRM hygiene

Customer Support

Tickets · knowledge base ·
escalations

Operations

Process tracking · vendor
coordination

Procurement

Sourcing · purchase orders
· renewals

Administration

Scheduling · documents ·
compliance prep



Roadmap

From digital workforce to a total business ecosystem

01

Digital Workforce

LIVE

Head Agent + core sub-agents
Approval workflows and notifications
Gmail, Slack, WhatsApp, LinkedIn integrations

02

Department Modules

IN BUILD

Agent cohorts for every business function
Org-aware routing and permissions at scale
Firecracker microVM isolation in production

03

Total Business Ecosystem

NEXT

One platform from HRM to CRM to full ERP
Enterprise deployments and compliance
Marketplace of specialized agent roles



The back office that runs itself

A governed digital workforce, coordinated by the Head Agent, supervised by your people, auditable end to end.

Made with Love by Intelsense AI and RunAgent

founders@intelsense.ai

