CHAPTER 01 · EXECUTIVE SUMMARY — BRIDGING THE ENTERPRISE AI DIVIDE

Why Enterprise AI Pilots Stall—and How Bea Unifies All 3 Layers

Standard AI tools sit disconnected from core apps and leak data. **Bea (Business Execution Assistant)** unifies all three in one hub.

6 TOOLS → 1 AI CHIEF OF STAFF

THE 3 STRUCTURAL GAPS — Why Legacy Enterprise AI Pilots Stall

01 · Disconnected Systems of Record — SOLVED AT 8:30 AM →

Standalone chatbots summarize documents but cannot check live balances or execute write-back actions in **Workday**, **ServiceNow**, or **ERP‡**.

❌ Read-Only Chat | ❌ Manual Copy-Paste | ✓ Bea 1-Click Action

02 · The 95% Workforce Adoption Gap — SOLVED AT 11:00 AM →

Complex prompt engineering and technical workflow builders leave **95% of business staff** behind while bottlenecking the **5% technical builders**.

❌ Blank Prompt Box | ✓ Council 360° Sandbox | ✓ Instant Helpers

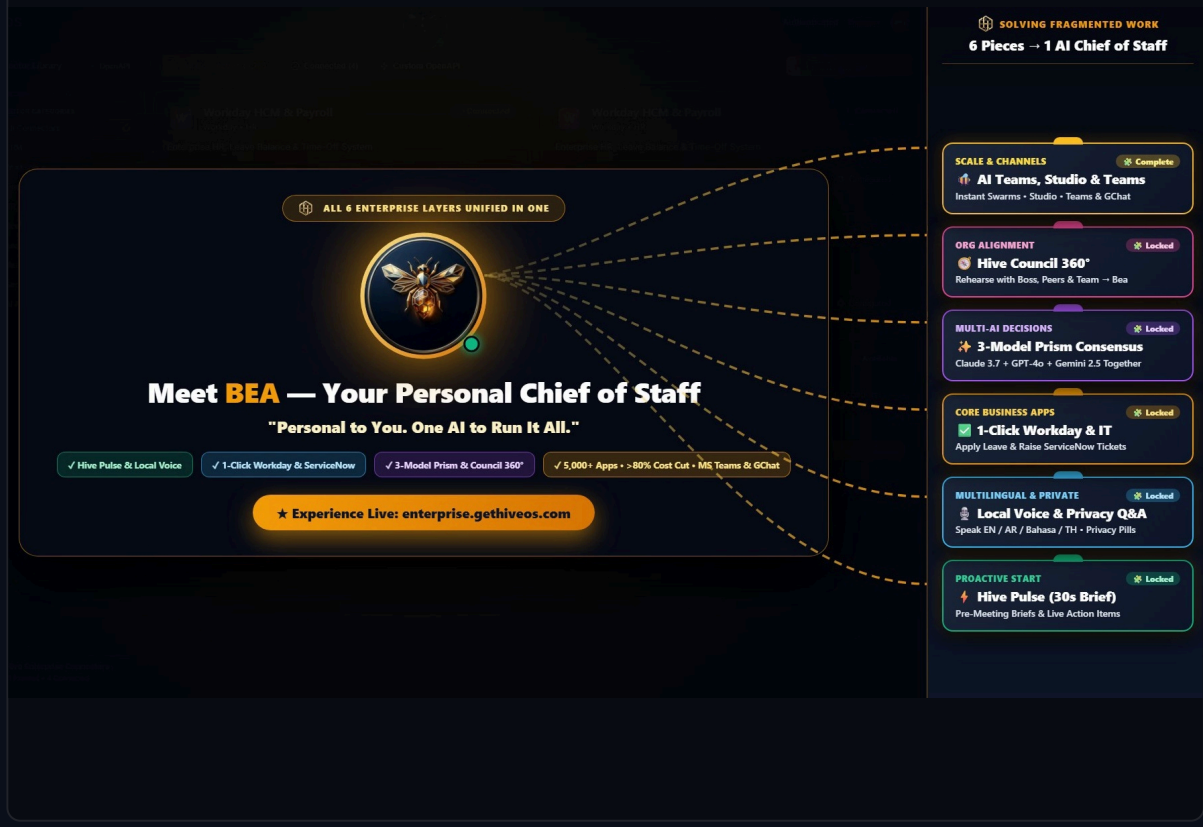
03 · Cloud PII Exposure & Token Burn — SOLVED AT 9:15 AM →

Sending raw HR/Finance records and routine queries to a single frontier cloud API exposes sensitive PII and inflates annual token spend by 5×*.

❌ Plaintext PII Egress | ✓ Privacy Pills† | ✓ ~82% Lower Spend*

Point tools force manual copy-pasting across tabs while exposing sensitive data.

THE SOLUTION · HIVEOS ENTERPRISE — Meet Bea — Proactive AI Chief of Staff Unifying All 6 Layers



Bea connects your knowledge base, systems of record, and multi-agent studios in one governed hub.

Connected Foundations 5,266 APPS‡ | **Business Teams & Builders** BEA + AGENT STUDIO | **Pre-Inference DLP & FinOps** UP TO ~82% LOWER SPEND*

CHAPTER 02 · 8:30 AM WORKFLOW — SOLVING GAP #1 (DISCONNECTED SYSTEMS)

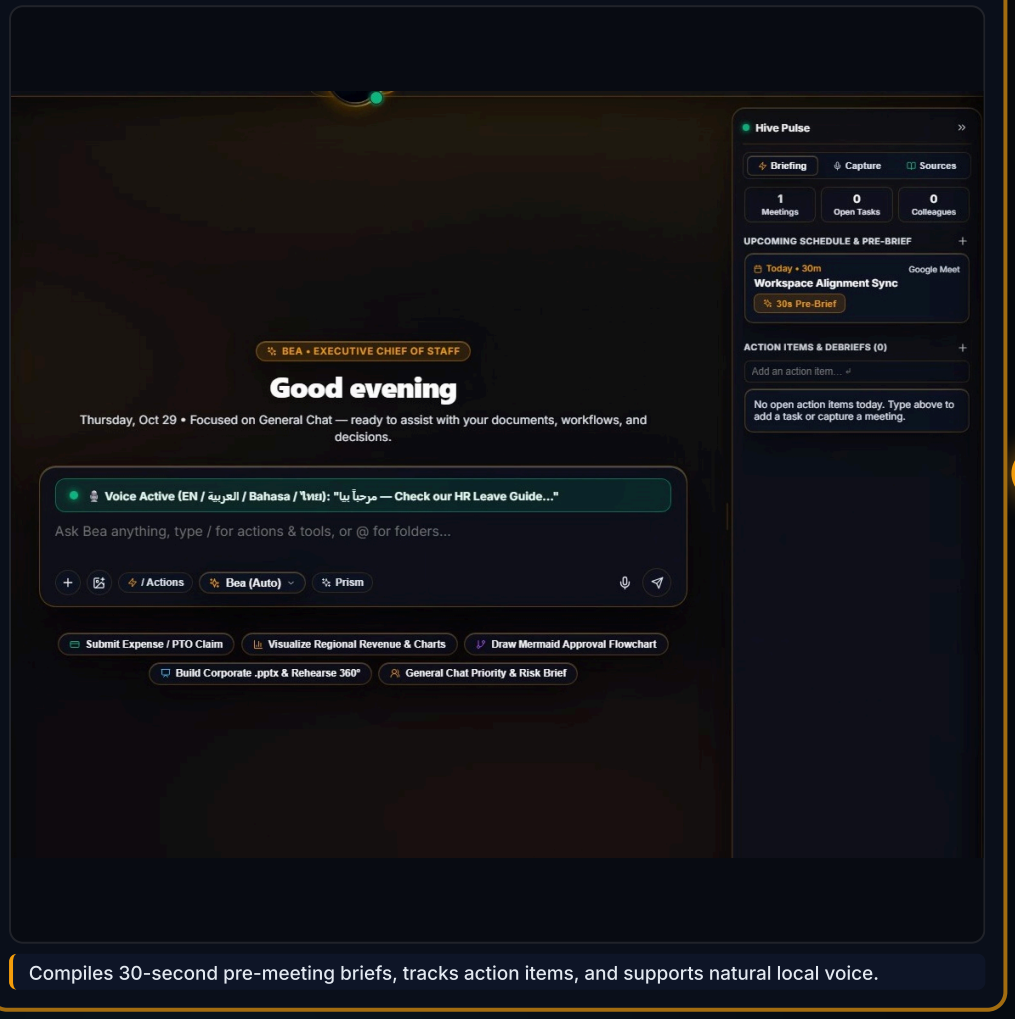
Start in "Today's Pulse." Execute in Workday in One Click.

ZERO PROMPTING · 1-CLICK WRITE-BACK

Before you type a single prompt, Bea prepares your 30-second pre-meeting brief* and drafts ready-to-submit HR and IT actions‡.

8:30 AM · MORNING BRIEFING

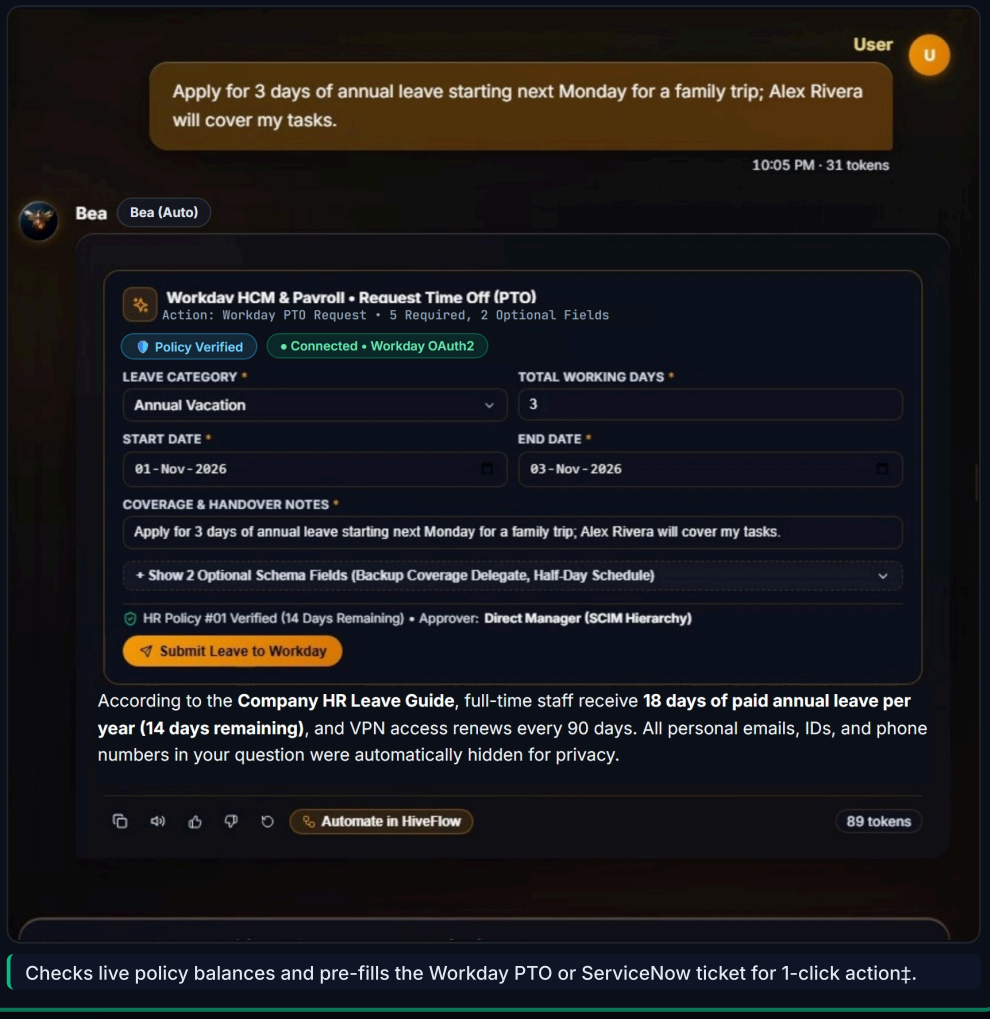
Today's Pulse (30s Pre-Brief*) & Local Voice



Compiles 30-second pre-meeting briefs, tracks action items, and supports natural local voice.

8:35 AM · 1-CLICK EXECUTION

1-Click Write-Back in Workday & ServiceNow‡



Checks live policy balances and pre-fills the Workday PTO or ServiceNow ticket for 1-click action‡.

30s Pre-Meeting Briefs* CALENDAR + DOCS

Natural Local Voice EN · AR · MS · TH · ZH

System-of-Record Actions‡ 1-CLICK WRITE-BACK

CHAPTER 03 · 9:15 AM WORKFLOW — SOLVING GAP #3A (DATA PRIVACY & VERIFICATION)

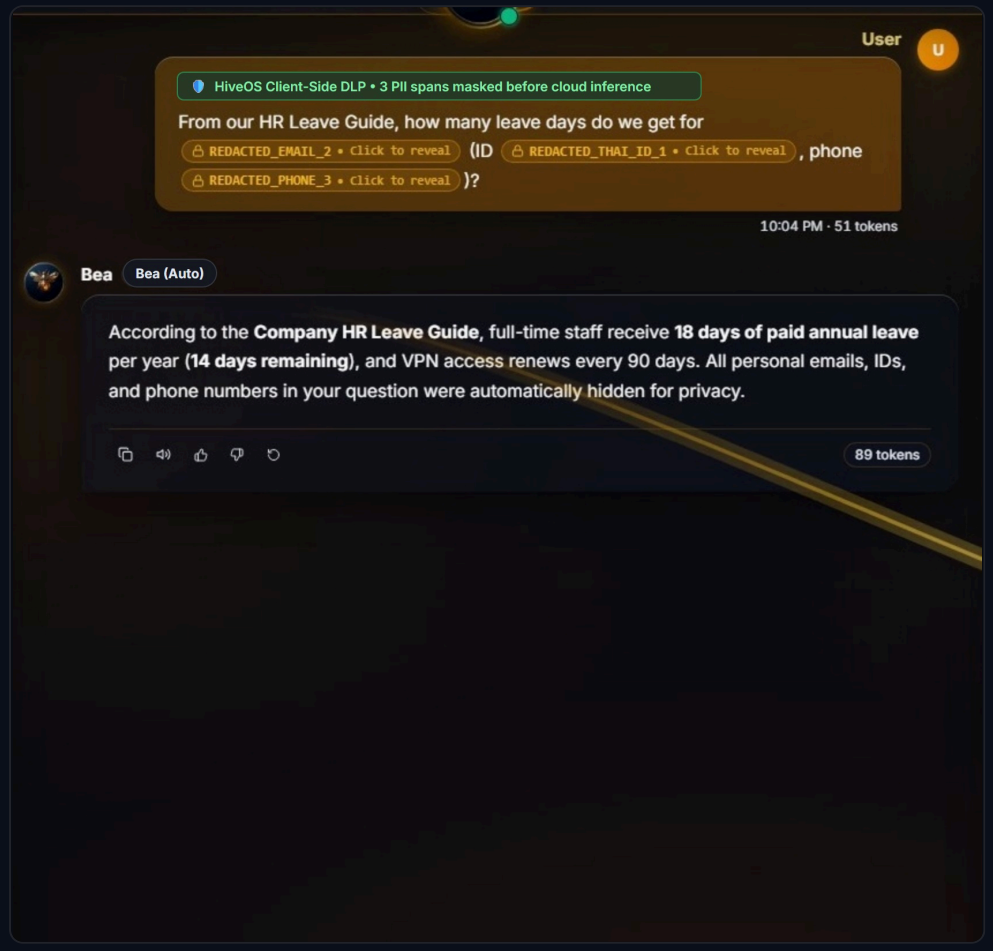
Mask Sensitive PII Locally. Verify Answers Across 3 Models.

Client-side Privacy Pills tokenize personal data before cloud inference†, while Hive Prism cross-checks Claude, GPT-4o, and Gemini‡.

PRE-INFERENCE PII REDACTION†

9:15 AM · CLIENT-SIDE DLP†

Clickable PII Privacy Pills (Tokenized Locally)



Masks emails, national IDs, and phone numbers locally into RBAC-gated pills before cloud calls†.

9:20 AM · HIVE PRISM

Claude 3.7 + GPT-4o + Gemini 2.5 Consensus‡



Queries three frontier models in parallel to flag policy conflicts and synthesize one verified answer‡.

RBAC-Gated Privacy Pills† CLICK TO REVEAL

3-Model Cross-Check‡ CLAUDE · GPT · GEMINI

Tenant Policy Guardrails† PDPA & GDPR READY†

CHAPTER 04 · 11:00 AM WORKFLOW — SOLVING GAP #2A (EXECUTIVE ALIGNMENT)

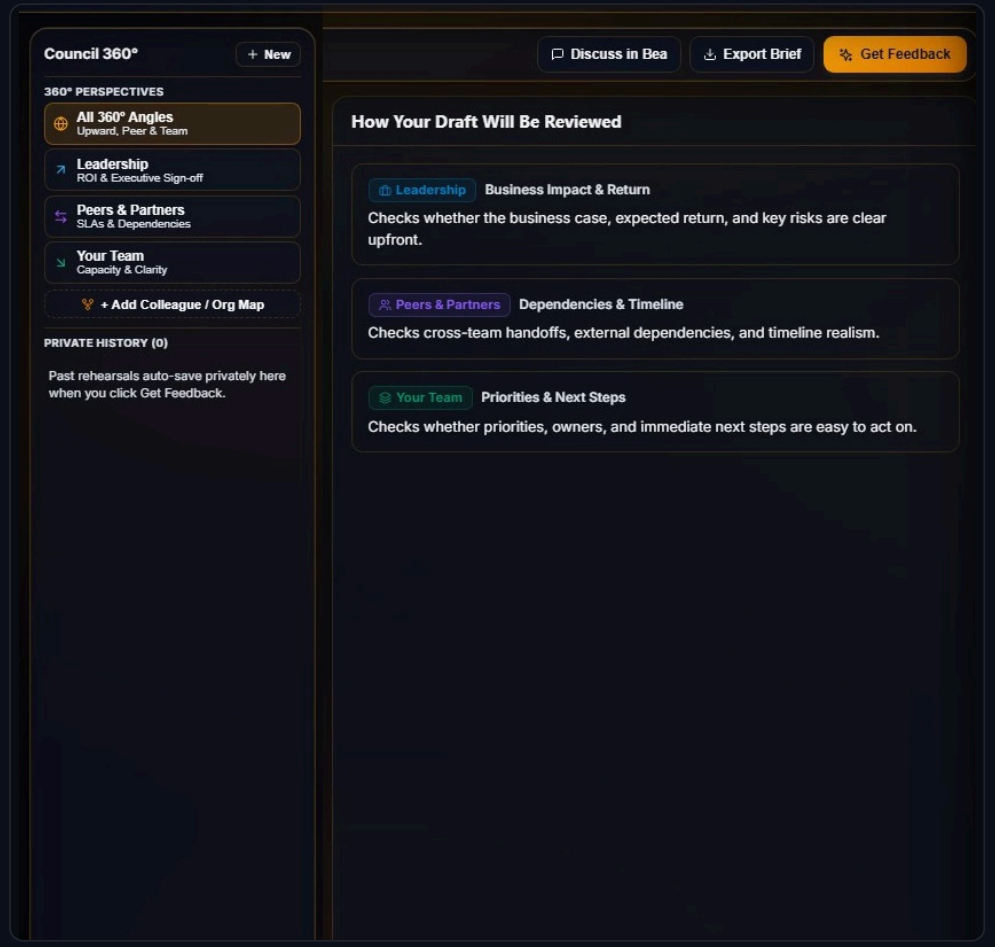
Stress-Test Strategic Proposals in a Private Executive Sandbox

Simulate objections from leadership, peers, and Finance/Risk/Ops personas in Hive Council 360°—then apply revisions in one click.

UP TO 96% ALIGNMENT SCORE*

11:00 AM · PRIVATE SANDBOX

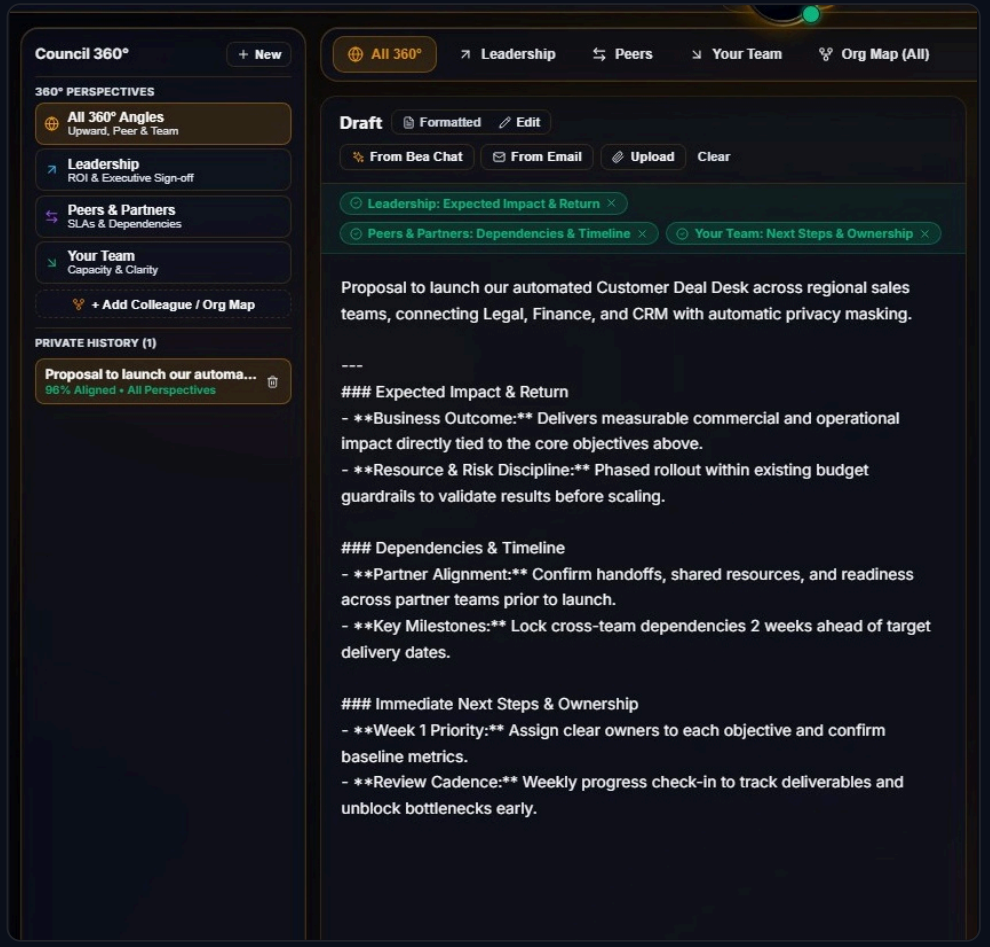
Simulate Review Across Leadership, Peers & Team



Select stakeholder perspectives from your org map in a private sandbox visible only to you.

11:02 AM · 1-CLICK "APPLY ALL"

Board-Ready Executive Synthesis in Seconds*



Adopts cross-functional feedback in one click and sends the polished plan back to Bea for action.

Org-Chart Stakeholder Personas LEADERSHIP · PEERS · TEAM

Private Rehearsal Sandbox NO PREMATURE ALERTS

1-Click "Apply All" & Act READY TO SHARE

CHAPTER 05 · 2:00 PM WORKFLOW — BRIDGING BUSINESS USERS (95%) & BUILDERS (5%)

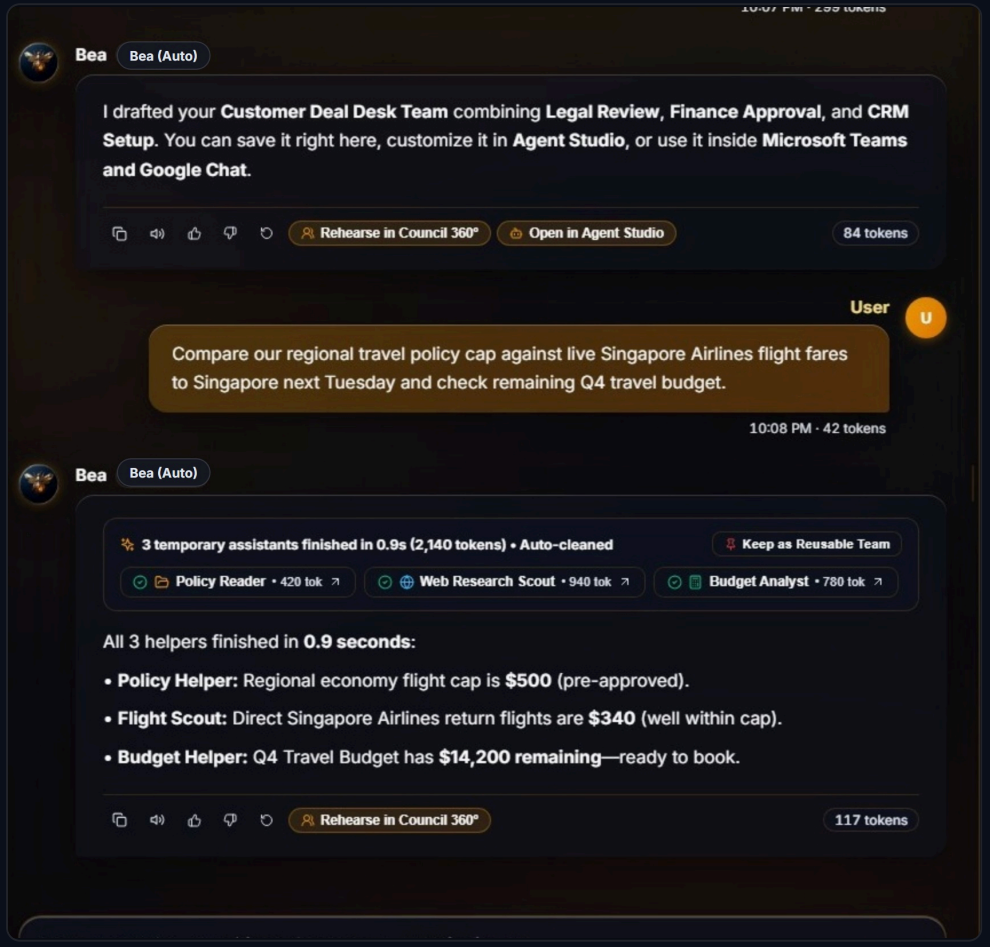
Instant Task Swarms for Business Teams. Visual Studio for Builders.

Everyday staff spin up parallel helper agents right inside chat, while technical builders design governed multi-agent workflows.

PARALLEL MULTI-AGENT ROUTING*

FOR BUSINESS TEAMS (95%)

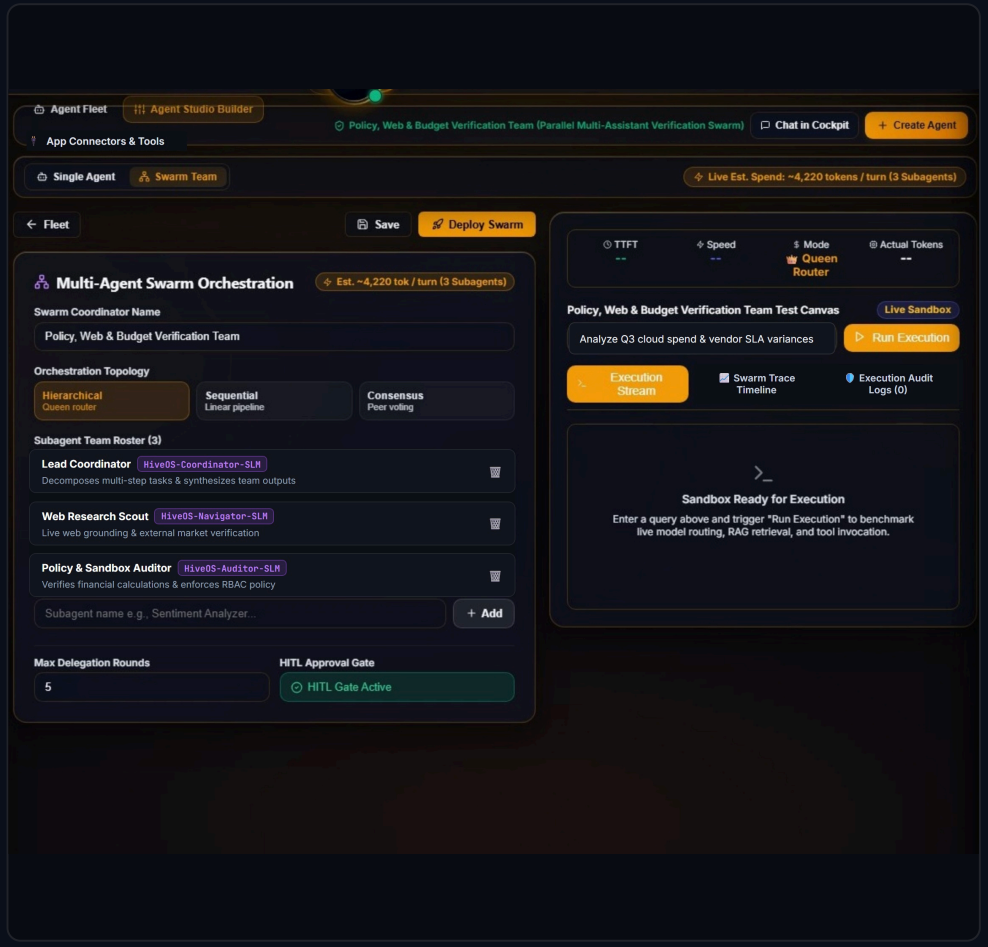
3 Parallel Helpers Spin Up Automatically in Chat*



Checks travel policy, live flights, and Q4 budgets in parallel—auto-cleaning or saving in 1 click*.

FOR TECHNICAL BUILDERS (5%)

Agent Studio: Swarm Topology & HITL Gates



Configure hierarchical swarms, sovereign SLM tiers, HITL gates, and 1-click MS Teams/GChat publish‡.

Single-Use Helper Swarms — AUTO-CLEANED MEMORY

Governed Agent Studio — HITL GATES ACTIVE

Omnichannel Deployment‡ — MS TEAMS & GOOGLE CHAT

CHAPTER 06 · 5:00 PM WORKFLOW — ENTERPRISE SCALE, FINOPS MODEL & ROLLOUT

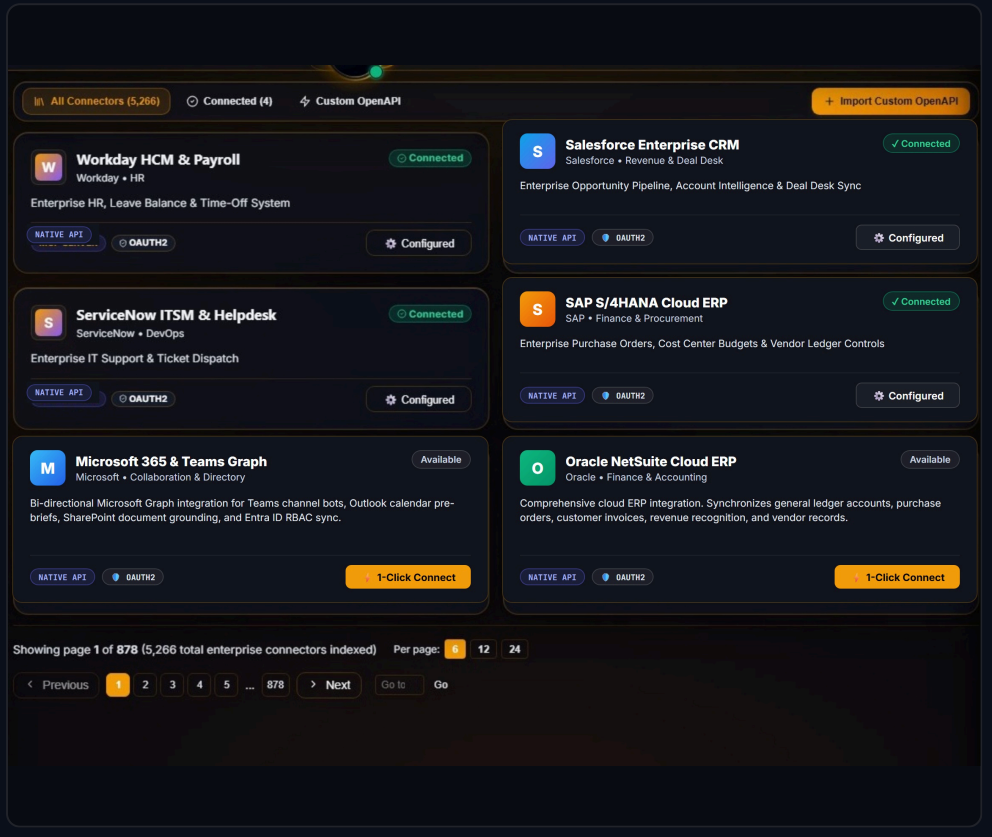
Connect 5,266 Enterprise Apps‡. Cut Annual AI Token Spend by ~82%.*

MODEL: $420K → $74K / YR (-82.4%)*

Route routine operational tasks to sovereign local/VPC Small Language Models (SLMs), reserving frontier APIs for complex synthesis*.

5:00 PM · CONNECTOR HUB‡

5,266 Pre-Indexed Connectors (Scoped OAuth2)



Pre-indexed integrations across ERP, CRM, HCM, ITSM, and storage with RBAC and HITL approval gates‡.

TOKEN FINOPS & ROLLOUT*

Illustrative 1,000-Seat Annual Token Model*

Unrouted 100% Frontier API Baseline* — $420,000 / YR*

HiveOS Tiered Routing (Sovereign SLM + Prism Consensus)* — $74,000 / YR*

EST. SAVINGS: $346,000 / YR (82.4%)*

WEEK 01*	WEEK 02*	WEEK 03*	WEEK 04*
Pulse & SSO	DLP & Prism	HR & IT Apps	Sovereign SLM

MANAGED CLOUD	PRIVATE VPC / BYOC	SOVEREIGN ON-PREM
Day-1 Trial Sandbox	AWS · Azure · GCP‡	Air-Gapped K8s

WATCH DEMO & CLAIM PRIORITY SANDBOX (+2 SEATS)

https://gethiveos.com

welcome@gethiveos.com →

*FINANCIAL & PERFORMANCE ESTIMATES: Cost comparison figures ($420,000/yr baseline vs. $74,000/yr hybrid routing; $346,000/yr or 82.4% estimated reduction) represent an **illustrative financial model** for a 1,000-seat enterprise executing ~10,000 daily prompts across 250 business days (comparing 100% frontier API token inference against tiered routing where routine operational queries execute on sovereign Small Language Models and complex synthesis escalates to multi-model consensus). Figures measure **estimated LLM token inference expenditure only** and exclude cloud hosting, hardware, implementation, and HiveOS licensing fees. Actual savings, latencies (e.g., 0.9s, 30s), alignment scores (96%), and deployment timelines (Weeks 01–04) vary by customer workload, infrastructure, and tenant configuration, and do not constitute a contractual guarantee, warranty, or binding SLA. All company metrics, names, national IDs, fares, and leave balances shown in UI screenshots are **fictitious demonstration data**. †SECURITY & COMPLIANCE: Client-side PII Privacy Pill redaction applies to supported entity patterns configured in tenant DLP policies; regulatory compliance (including Malaysia PDPA, Singapore PDPA, and EU GDPR) is a shared responsibility between Logic42 and the customer (Data Controller). ‡THIRD-PARTY TRADEMARKS & IP NOTICE: Workday, ServiceNow, Microsoft 365, Microsoft Teams, Google Workspace, Google Chat, Salesforce, SAP, Oracle NetSuite, Anthropic Claude, OpenAI GPT-4o, Google Gemini, AWS, Azure, and GCP are trademarks of their respective owners; Logic42 is not affiliated with, endorsed by, or sponsored by them. Connector integration requires valid customer licenses and API credentials. Architecture and materials protected under **MyIPO Copyright Act 1987 (Act 332) Voluntary Notification CRLY2026E**. © 2026 Logic42. All rights reserved. Provided "AS IS" for informational evaluation only.