BLUEJAY

srmg
Saudi Research & Media Group

MEDIA & ENTERTAINMENT · GENERATIVE AI · AI PLATFORMS

Building a sovereign, AI-native Media OS for Arabic media

نظام تشغيل إعلامي سيادي قائم على الذكاء الاصطناعي

Bluejay is SRMG's in-house agentic platform for newsrooms. It runs editorial work as governed, repeatable pipelines, and nothing reaches readers without an editor's approval.

Jul 2025
Harness development started

700+
Person-hours, built by a two-person team

400+
Models across 10 providers, plus on-premise models

Nov 2026
Organisation-wide rollout

Newsrooms adopted AI faster than they could govern it

A multi-publication media group cannot let unreviewed text, unpublished sources or unpredictable spend run through consumer chat tools.

	BEFORE: OPEN-ENDED CHATBOTS	WITH BLUEJAY: GOVERNED PIPELINES
Workflow	Research, drafting, social, newsletters and multimedia in separate tools, with no handoff or shared memory	Multi-agent pipelines from research to every format, with persistent sessions and a shared library
House style	Arabic and English voice re-explained in every session, for every title	Each publication's voice and standards encoded as skills, applied by every agent
Data	Unpublished stories and embargoed material pasted into public tools	Editorial data kept in SRMG-controlled infrastructure; on-premise models for investigations
Accountability	No record of which sources or model produced a paragraph	Editor approval before publication and an exportable trace of every run
Cost	Agents search repeatedly; bills are opaque and unattributed	Context size and cost shown before a run; spend attributed to each desk
Ways of working	Free-form prompting that depends on each editor's skill	Form-first wizards, saved configurations and desk champions

One governed engine, built in-house from first principles

People
who steer the work
- Editors-in-chief
- Desk editors
- Producers
- Investigators
- Social & newsletter teams

Experience
form-first, not prompt-first
- **Second Brain**
- Podcast Creator
- Image Editor
- Social Video
- Newsletter Creator
- Content Repurpose
- **No-code agent builder**

Orchestration
loops with exit conditions
- **DAG multi-agent flows**
- Human-review nodes
- Scheduler & automations
- Refinement loops

Agent harness
Pydantic AI-native runtime
- Typed tools
- MCP
- Agent skills
- Context compaction
- Prefetch
- Session resume
- Hard stop
- Per-run cost accounting

Knowledge & data
SRMG-controlled
- **Vector RAG (Milvus, AWS)**
- Datasets (read/write)
- Feeds & web-as-feed
- Governed library
- Run ledger

Inference
model-agnostic, per step
- OpenAI
- Anthropic
- Google Gemini
- xAI
- Groq
- Amazon Bedrock
- DeepSeek
- Moonshot
- Xiaomi MiMo
- OpenRouter
- **On-premise models**

Governance plane
runtime primitives across every layer

Tool allow-lists
per agent and per edition

Credential isolation
agents never see raw secrets

Human-review gates
no ungoverned path to publication

Guardrails
Clarify on tone, audience, legal

Hard stop
cancel ends the run cleanly

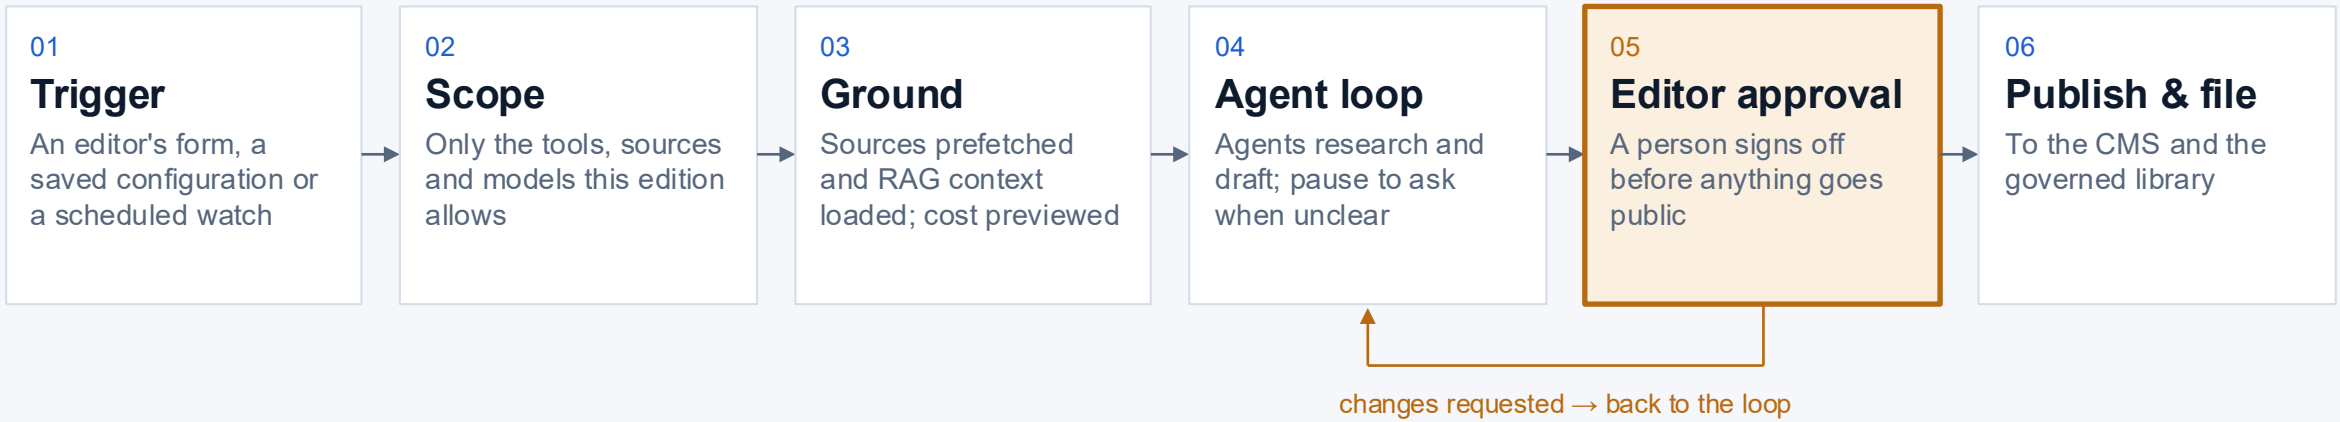
Explainability traces
exportable, auditable

Cost & telemetry
per run, model and desk

SRMG enterprise systems CMS · analytics · marketing platforms

via REST API, webhooks and MCP

Every run is scoped, grounded, reviewed and recorded



Access, least privilege

- Each publication is an isolated configuration of agents, tools, models and UI
- Credentials held by the platform and referenced by name
- File access confined to designated workspace folders

Guardrails, human in the loop

- Clarify pauses on tone, audience, naming and redaction
- Draft → approve → publish; Bluejay never posts on its own
- Bounded loops and a hard stop on every run

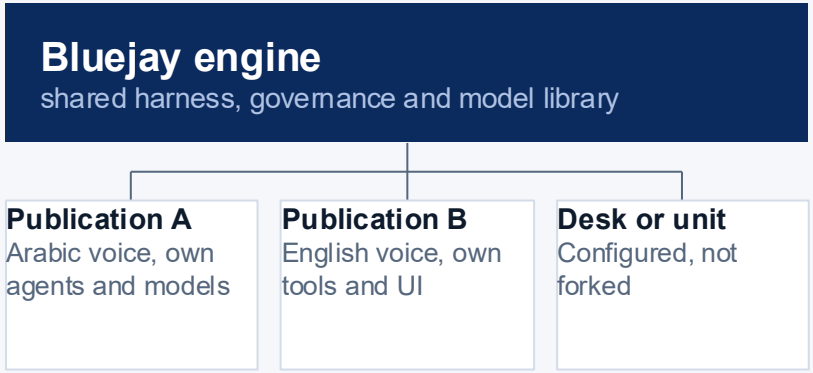
Telemetry and cost control

- Context size and expected cost shown before Run
- Spend attributed per workflow, model and desk
- Traces flag cost spikes and adversarial agent patterns

From reactive desks to proactive, governed production

Proactive breaking news	A scheduled watch finds breaking stories, qualifies them and returns the rationale plus first drafts in every format. **Hours become minutes.**
Deep research agents	Sources pulled into a governed library and built into long-form drafts with citation blocks; sensitive material runs on on-premise models.
Overnight intelligence	Automations monitor trends and topics so the daily brief is ready before the morning meeting.
Content repurpose	One approved article becomes social posts, a voiced podcast, a short video and a slide summary, each tied to the approved source.
Agentic knowledge base	Unstructured archives as RAG context, plus structured datasets agents can read and write.

ONE ENGINE, MANY GOVERNED EDITIONS



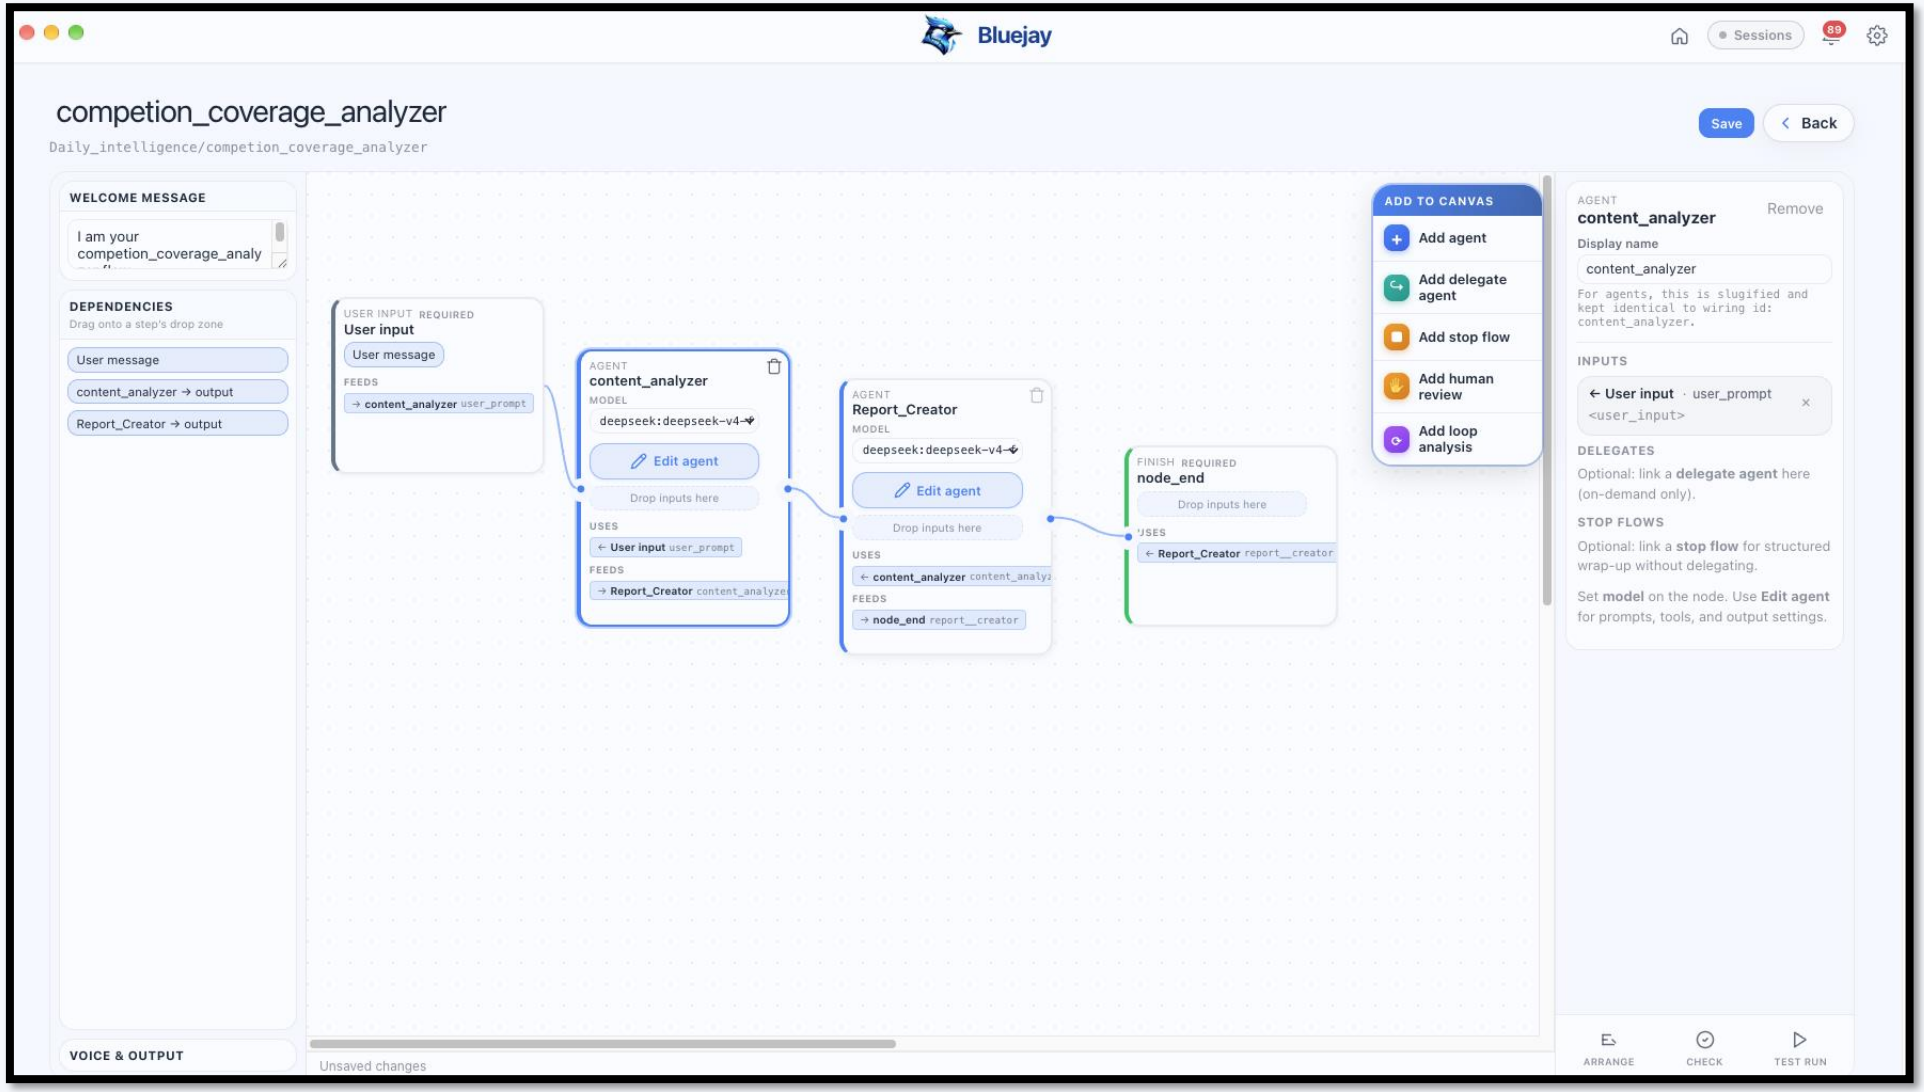
Each edition runs in isolation with its own agents, tools, models and interface, so new titles join by configuration rather than new code.

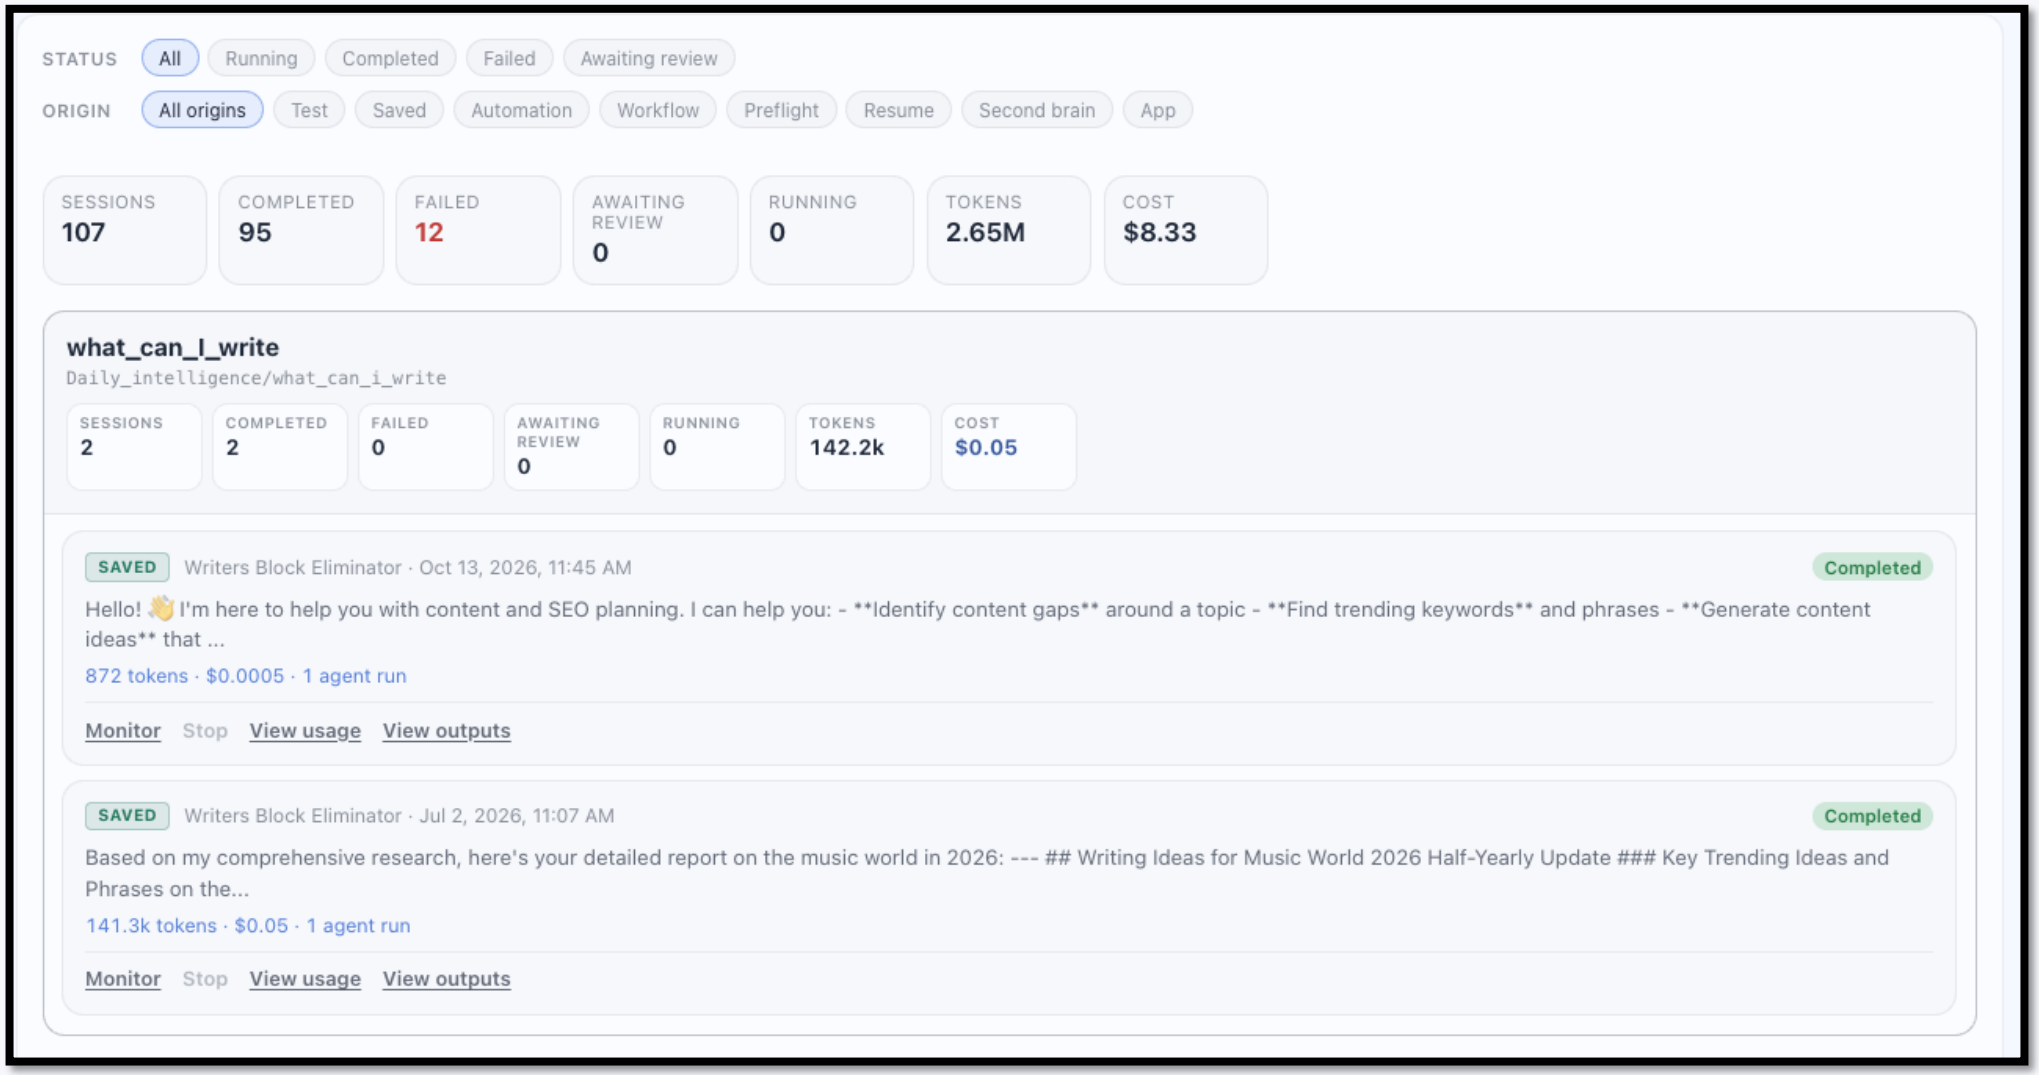


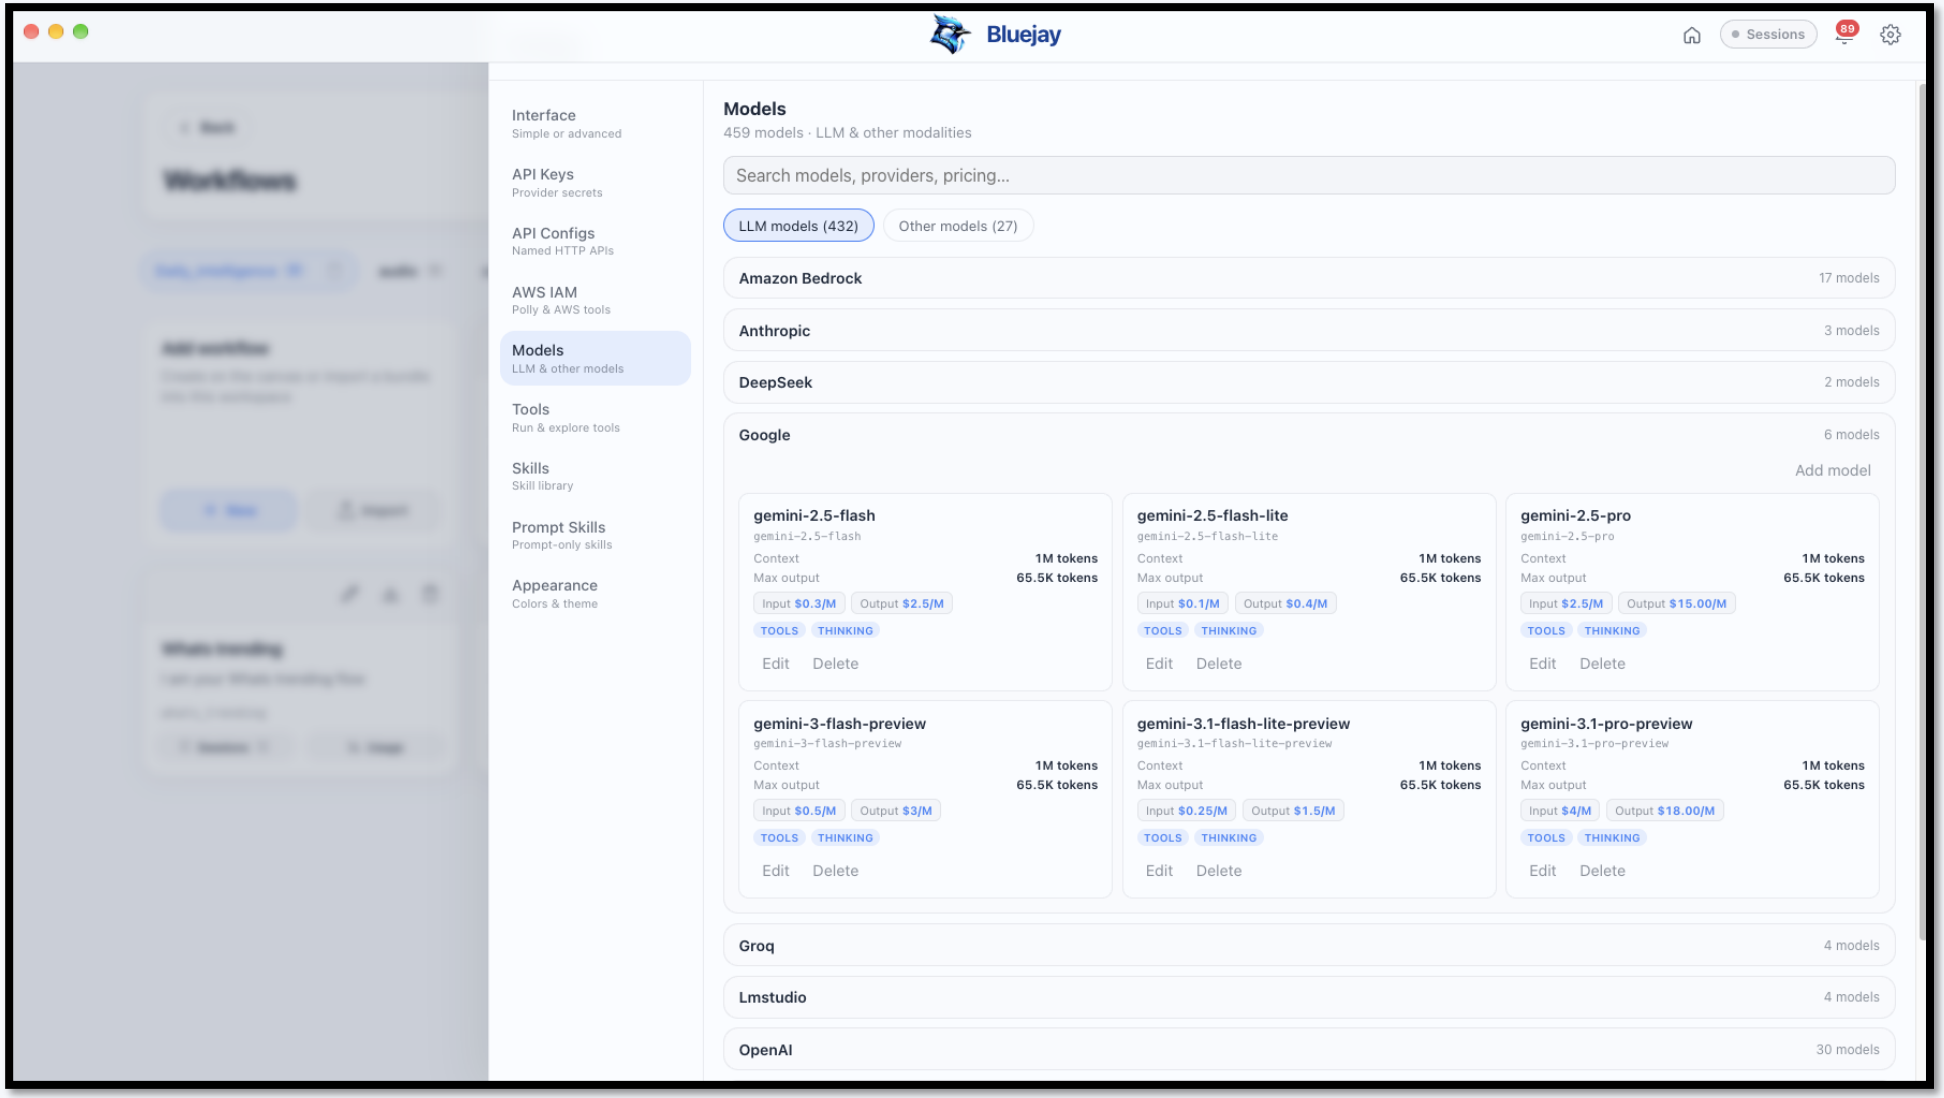
JUL 2025
Harness build
Pydantic AI runtime

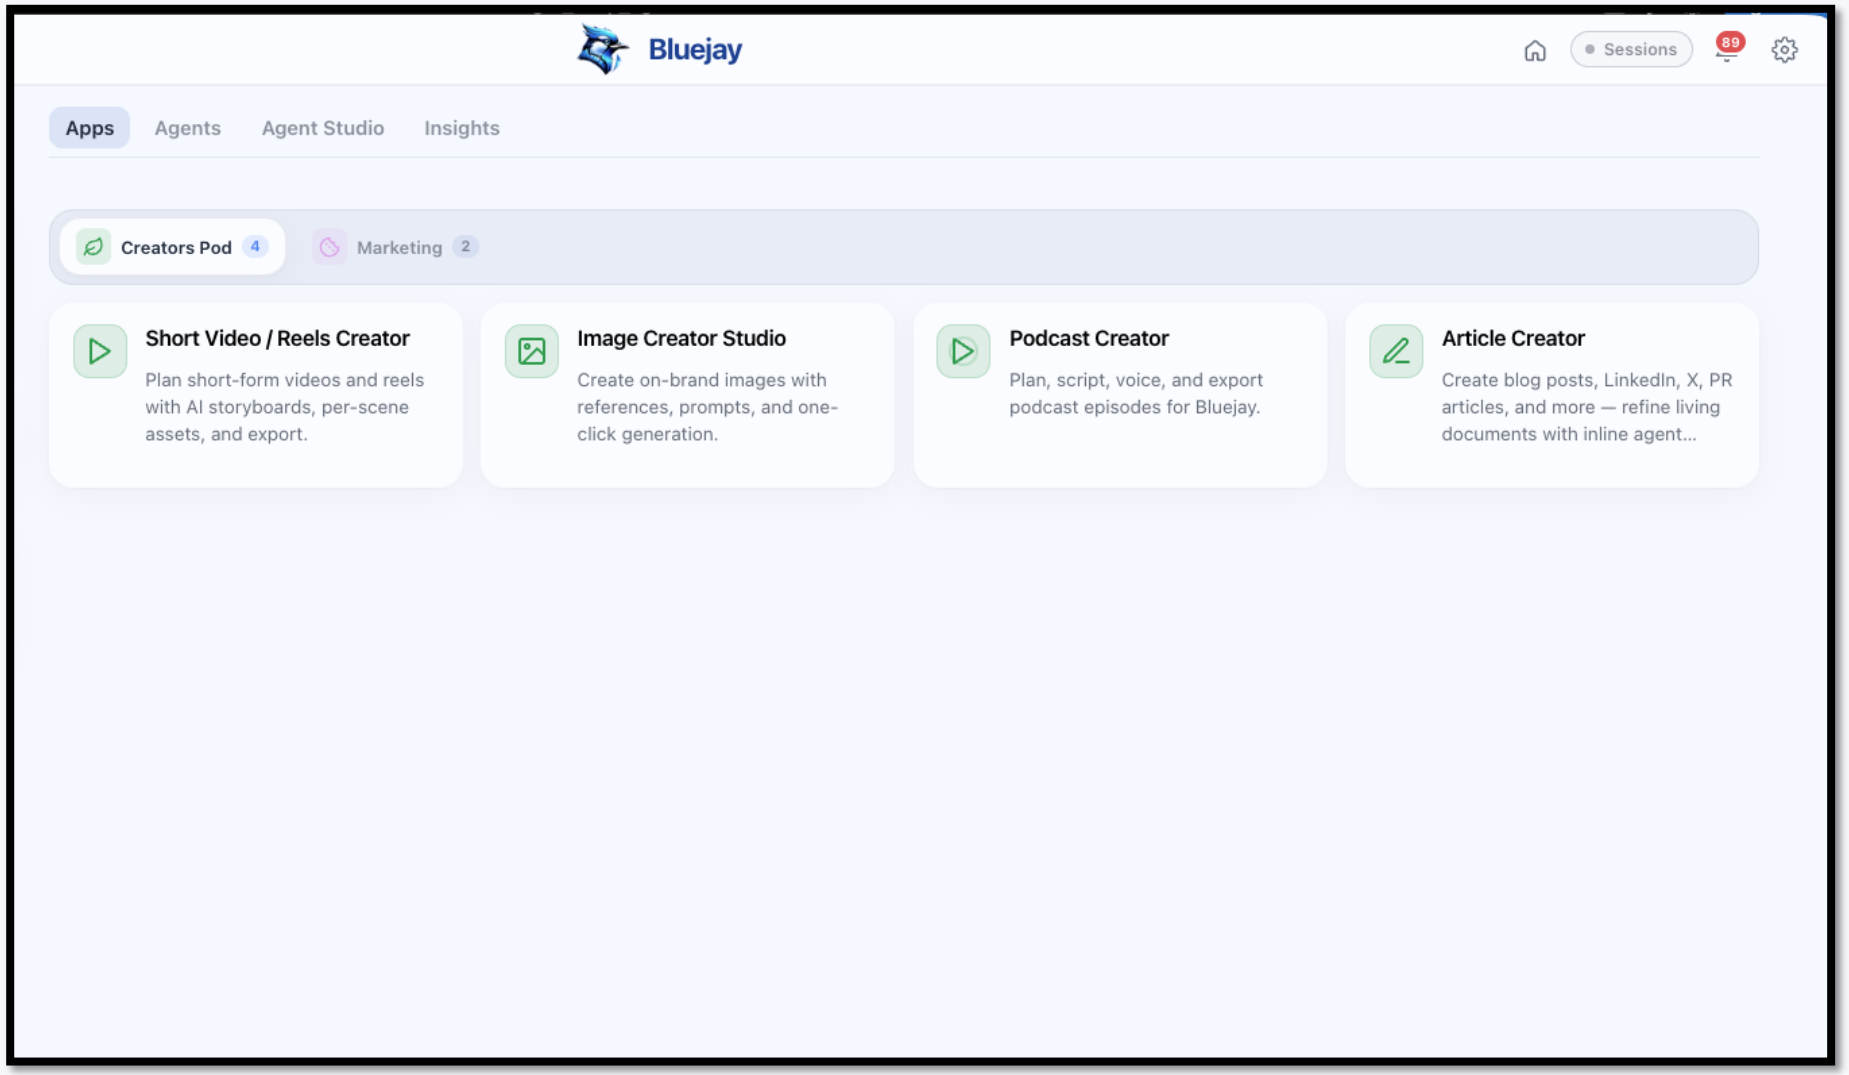
JUN 2026
Public beta
SRMG editorial teams

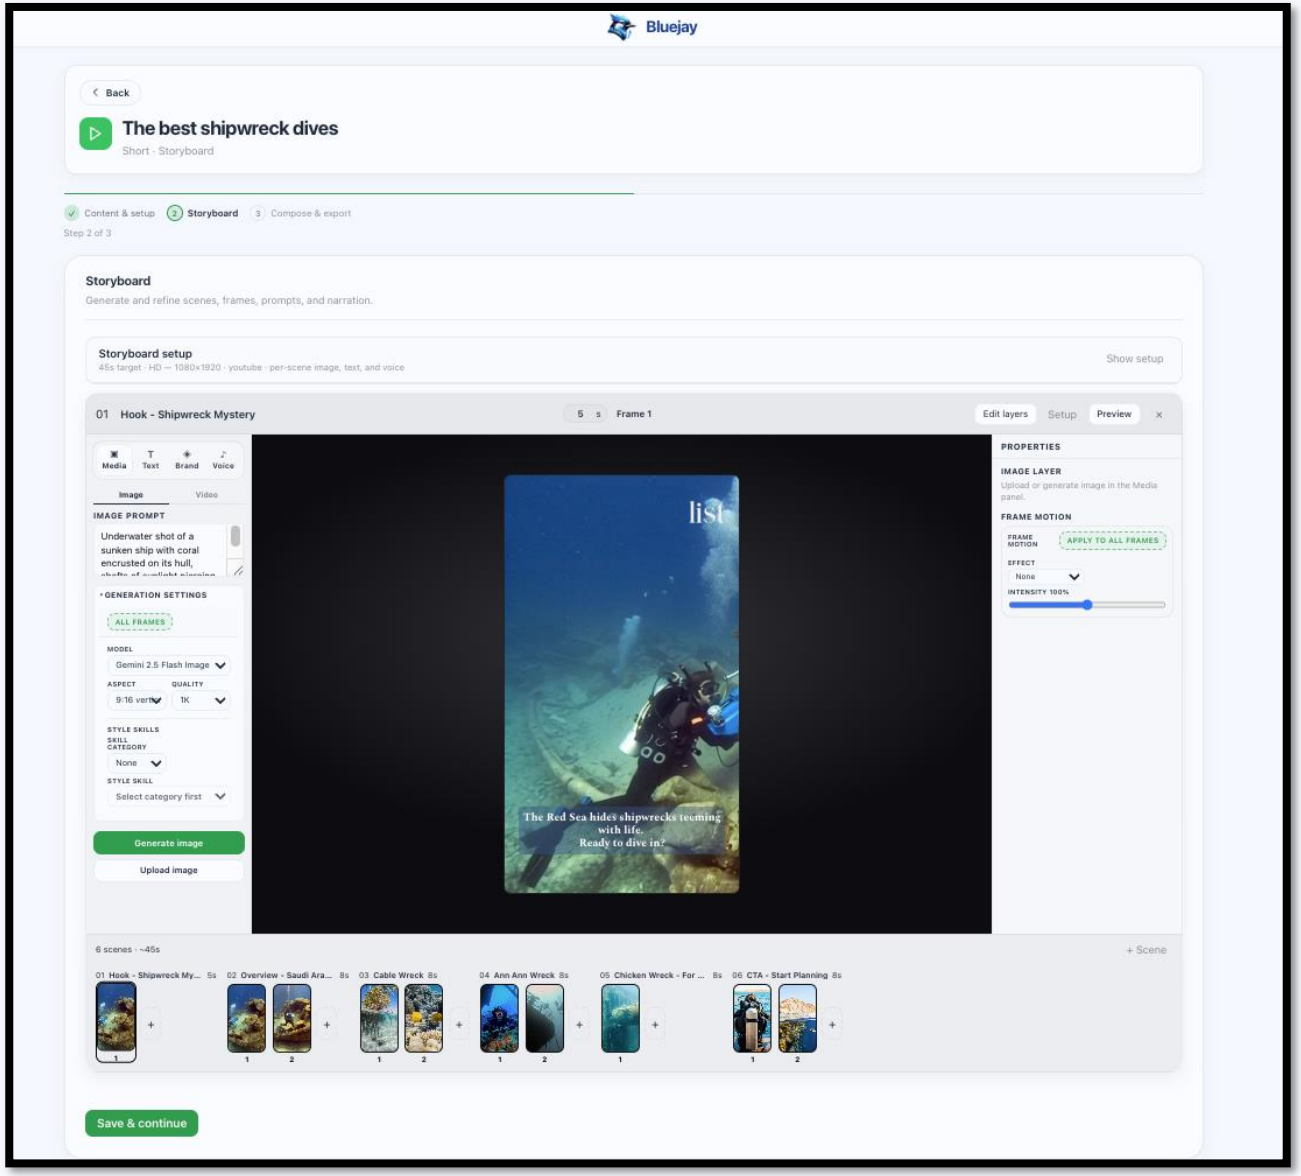
NOV 2026
Org-wide rollout
Projected, mid-month


Bluejay
Sessions
89
competion_coverage_analyzer
Daily_intelligence/competion_coverage_analyzer
Save
Back
WELCOME MESSAGE
I am your competion_coverage_analy
DEPENDENCIES
Drag onto a step's drop zone
User message
content_analyzer → output
Report_Creator → output
VOICE & OUTPUT
USER INPUT REQUIRED
User input
User message
FEEDS
→ content_analyzer user_prompt
AGENT
content_analyzer
MODEL
deepseek:deepseek-v4-
Edit agent
Drop inputs here
USES
← User input user_prompt
FEEDS
→ Report_Creator content_analyze
AGENT
Report_Creator
MODEL
deepseek:deepseek-v4-
Edit agent
Drop inputs here
USES
← content_analyzer content_analy
FEEDS
→ node_end report__creator
FINISH REQUIRED
node_end
Drop inputs here
USES
← Report_Creator report__creator
ADD TO CANVAS
Add agent
Add delegate agent
Add stop flow
Add human review
Add loop analysis
AGENT
content_analyzer
Remove
Display name
content_analyzer
For agents, this is slugified and kept identical to wiring id: content_analyzer.
INPUTS
← User input · user_prompt
<user_input>
DELEGATES
Optional: link a delegate agent here (on-demand only).
STOP FLOWS
Optional: link a stop flow for structured wrap-up without delegating.
Set model on the node. Use Edit agent for prompts, tools, and output settings.
Unsaved changes
ARRANGE
CHECK
TEST RUN


STATUS
All
Running
Completed
Failed
Awaiting review
ORIGIN
All origins
Test
Saved
Automation
Workflow
Preflight
Resume
Second brain
App
SESSIONS
107
COMPLETED
95
FAILED
12
AWAITING REVIEW
0
RUNNING
0
TOKENS
2.65M
COST
$8.33
what_can_I_write
Daily_intelligence/what_can_i_write
SESSIONS
2
COMPLETED
2
FAILED
0
AWAITING REVIEW
0
RUNNING
0
TOKENS
142.2k
COST
$0.05
SAVED
Writers Block Eliminator · Oct 13, 2026, 11:45 AM
Completed
Hello! 👋 I'm here to help you with content and SEO planning. I can help you: - **Identify content gaps** around a topic - **Find trending keywords** and phrases - **Generate content ideas** that ...
872 tokens · $0.0005 · 1 agent run
Monitor
Stop
View usage
View outputs
SAVED
Writers Block Eliminator · Jul 2, 2026, 11:07 AM
Completed
Based on my comprehensive research, here's your detailed report on the music world in 2026: --- ## Writing Ideas for Music World 2026 Half-Yearly Update ### Key Trending Ideas and Phrases on the...
141.3k tokens · $0.05 · 1 agent run
Monitor
Stop
View usage
View outputs


Bluejay
Sessions
89
Interface
Simple or advanced
API Keys
Provider secrets
API Configs
Named HTTP APIs
AWS IAM
Polly & AWS tools
Models
LLM & other models
Tools
Run & explore tools
Skills
Skill library
Prompt Skills
Prompt-only skills
Appearance
Colors & theme
Models
459 models · LLM & other modalities
Search models, providers, pricing...
LLM models (432)
Other models (27)
Amazon Bedrock
17 models
Anthropic
3 models
DeepSeek
2 models
Google
6 models
Add model
gemini-2.5-flash
gemini-2.5-flash
Context 1M tokens
Max output 65.5K tokens
Input $0.3/M
Output $2.5/M
TOOLS
THINKING
Edit
Delete
gemini-2.5-flash-lite
gemini-2.5-flash-lite
Context 1M tokens
Max output 65.5K tokens
Input $0.1/M
Output $0.4/M
TOOLS
THINKING
Edit
Delete
gemini-2.5-pro
gemini-2.5-pro
Context 1M tokens
Max output 65.5K tokens
Input $2.5/M
Output $15.00/M
TOOLS
THINKING
Edit
Delete
gemini-3-flash-preview
gemini-3-flash-preview
Context 1M tokens
Max output 65.5K tokens
Input $0.5/M
Output $3/M
TOOLS
THINKING
Edit
Delete
gemini-3.1-flash-lite-preview
gemini-3.1-flash-lite-preview
Context 1M tokens
Max output 65.5K tokens
Input $0.25/M
Output $1.5/M
TOOLS
THINKING
Edit
Delete
gemini-3.1-pro-preview
gemini-3.1-pro-preview
Context 1M tokens
Max output 65.5K tokens
Input $4/M
Output $18.00/M
TOOLS
THINKING
Edit
Delete
Groq
4 models
Lmstudio
4 models
OpenAI
30 models


Bluejay
Sessions
89
Apps
Agents
Agent Studio
Insights
Creators Pod 4
Marketing 2
Short Video / Reels Creator
Plan short-form videos and reels with AI storyboards, per-scene assets, and export.
Image Creator Studio
Create on-brand images with references, prompts, and one-click generation.
Podcast Creator
Plan, script, voice, and export podcast episodes for Bluejay.
Article Creator
Create blog posts, LinkedIn, X, PR articles, and more — refine living documents with inline agent...


Bluejay
Back
The best shipwreck dives
Short · Storyboard
Content & setup
Storyboard
Compose & export
Step 2 of 3
Storyboard
Generate and refine scenes, frames, prompts, and narration.
Storyboard setup
Show setup
01 Hook - Shipwreck Mystery
Frame 1
Edit layers
Setup
Preview
Media
Text
Brand
Voice
Image
Video
IMAGE PROMPT
Underwater shot of a sunken ship with coral encrusted on its hull,
GENERATION SETTINGS
ALL FRAMES
MODEL
Gemini 2.5 Flash Image
ASPECT
QUALITY
1K
STYLE SKILLS
SKILL CATEGORY
None
STYLE SKILL
Select category first
Generate image
Upload image
list
The Red Sea hides shipwrecks teeming with life. Ready to dive in?
PROPERTIES
IMAGE LAYER
Upload or generate image in the Media panel.
FRAME MOTION
APPLY TO ALL FRAMES
EFFECT
None
INTENSITY 100%
6 scenes · ~45s
+ Scene
01 Hook - Shipwreck My... 5s
02 Overview - Saudi Ara... 8s
03 Cable Wreck 8s
04 Ann Ann Wreck 8s
05 Chicken Wreck - For ... 8s
06 CTA - Start Planning 8s
Save & continue