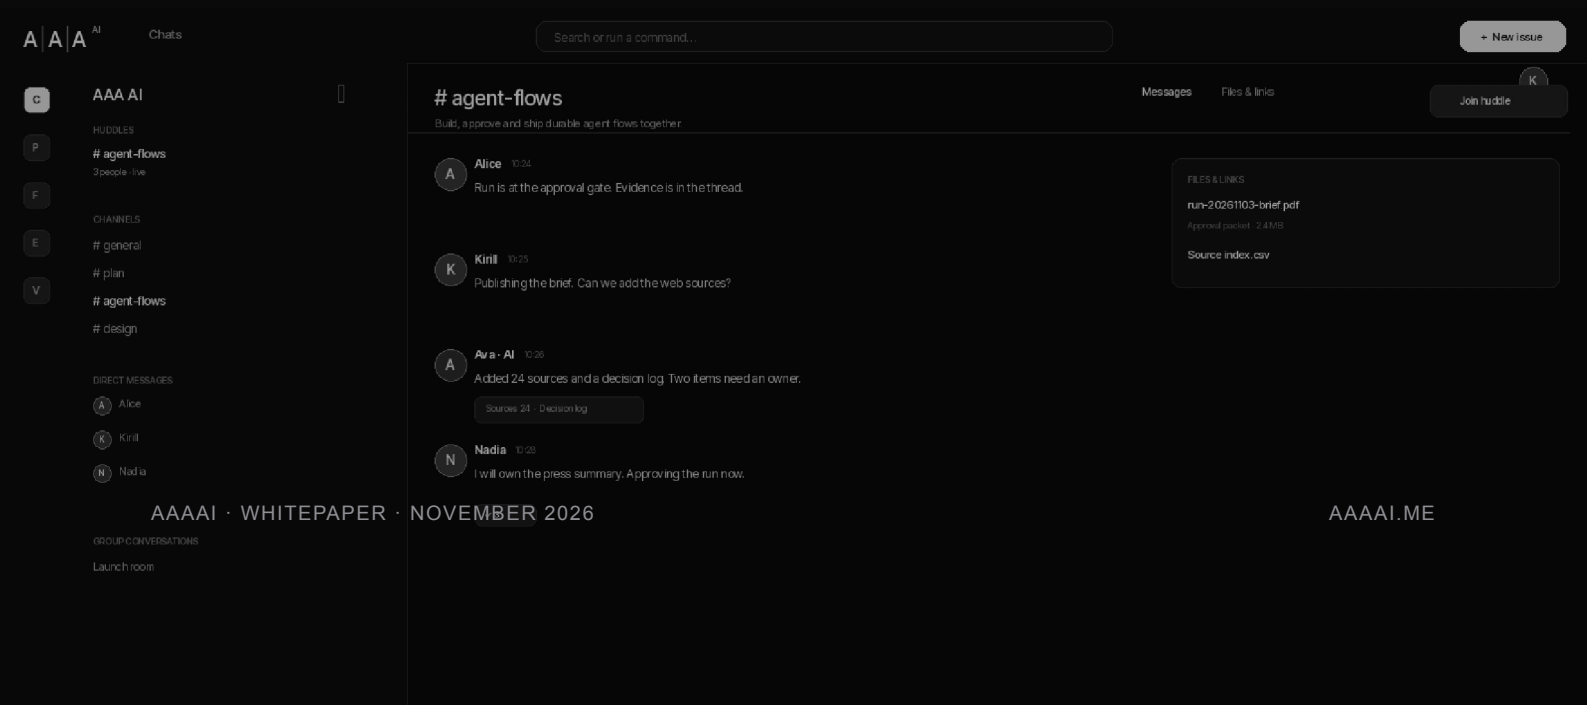
A|A|A AI

PRODUCT WHITEPAPER · CHAT

Chat

Multi-expert conversations, a Slack-style workspace and live calls with AI notes — one continuous thread from question to action.



WHITEPAPER · 02

The conversation is the workspace.

Chat is the front door to AAA AI. It is where a question becomes a decision, a decision becomes a task, and a task becomes coordinated work across people, agents and systems.

Most AI chat products optimize a single answer. AAA AI optimizes the loop around it: multiple experts reason and verify, tools and retrieval ground the answer, memory preserves context, and the same thread can move into a Plan issue, an Agent Flow or a live call.

10
expert modules with transparent reasoning

Multi
surface chat across web, iOS, Android and macOS

What the product is

- **Multi-expert reasoning.** Specialized experts debate, critique and verify; the planner synthesizes an answer instead of relying on a single model pass.
- **Context that persists.** Documents, web evidence, memory and workspace history remain attached to the conversation.
- **Workspace chat.** Channels, direct messages, group conversations, huddles, reactions, files and unread counts in a familiar Slack-style room.
- **Calls and AI notes.** Voice, video, screen share and recordings with live summary, plan, decisions and action items.

> A conversation should not lose its context when it becomes work.

WHITEPAPER · 03

One loop: ask, verify, act.

Chat is not a separate product from the rest of the operating system. It is the surface where the operator chooses what to do next.

How a conversation moves

STAGE	WHAT HAPPENS
Ask	The request is routed through the chosen model, available tools and relevant experts.
Verify	Experts compare assumptions and evidence; transparent mode exposes disagreement, confidence and the final synthesis.
Act	The outcome can become a Plan issue, an Agent Flow run, an approval-gated email, a note, a task or a live call.

Modes for different work

- **Chat** for general multi-expert reasoning and document work.
- **Task** and **Team** for structured or collaborative execution.
- **Auto-scale**, dynamic experts and cognitive scaling when the work needs a larger agent pool.

Workspace chat

Every account starts in **#general**. Public and private channels, DMs, group conversations and huddles share the same call stack. A call in progress shows a Join action directly in the room; unread counts and attachment metadata are persisted, with attachment bytes in S3-backed object storage.

WHITEPAPER · 04

Trust, identity and deployment.

Identity and permissions

Chat uses the same authenticated account as every other surface. Workspace membership, channel visibility and direct-message access are enforced by the shared application, not by the model. Role and policy controls apply to tools, agent spend and enterprise deployment.

Data handling

- Chat history, channel membership and attachment metadata are persisted in the application database; object storage keeps files and links.
- Model choice remains yours: local Ollama, private endpoints, cloud providers or a mixed model layer.
- Memory, documents and workspace context are scoped to the authenticated user and deployment.
- Enterprise deployments can run in cloud, VPC, on-prem or fully offline configurations.

What to standardize

- Define which conversations are shared, private or retained, and which tools agents may call.
- Choose a default model and a verification policy per team.
- Connect calls, email and Plan so the thread can move without copy/paste.

> The best chat product is not the one with the longest answer. It is the one that finishes the work.