

PRODUCT MANUAL

Napster Agent Product Manual.

Build, deploy, and measure AI agents on Napster Agent — from first URL to live analytics.

VERSION

1.1

DATE

September 2026

FOR

People who build and deploy AI agents on the platform

SIGN IN

`spaces.napster.ai` (soon `agent.napster.ai`)

Napster Agent was formerly Napster Spaces — same product, renamed. © 2026 Napster. Product details reflect the current release and may change as Napster Agent is updated.

About this manual

For people building AI agents on Napster Agent — you sign in, you build an agent, you deploy it, you watch how it performs. No prior experience assumed.

End visitors — the people who talk to your agent — don't need this manual. Their experience is designed to need no training.

How this is organized

SECTION	WHAT IT COVERS
§2 How Napster Agent works	The mental model — build vs. run.
§3 Sign in and get oriented	Sign in, the two surfaces, the tabs.
§4 Build your first agent	Every step from creation to publish, tab by tab.
§5 What the agent does at runtime	Voice, text, screen, memory, summary, multi-language.
§6 Agent tools	Built-in tools + custom tools.
§7 Commerce	Shopify, WooCommerce, other catalogs, cart behaviour.
§8 Publish and share	Deployment surfaces, embed options, on-site controls.
§9 Analytics	Standard analytics + Shopping Analytics.
§10 Known limitations	What to be aware of.
§11 FAQ	Common questions.
§12 Glossary	Quick reference for terms.

Conventions

- Click paths use the format: **Builder** → **Section** → **Item**.
- Field names appear in *italics*; button labels appear in "quotes".

- Callouts flag tips, notes, and caveats.

NOTE

Napster Spaces was the earlier name for this product. The product itself is unchanged; some UI surfaces and URLs still say "Spaces." Where the UI shows one of these, this manual uses the same label.

Getting help. If something here doesn't match what you see, or you need an answer this manual doesn't cover — security and compliance reviews, pricing, integrations — contact your Napster representative or Napster Support.

How Napster Agent works

Two things to understand about any agent: how you **build** one, and how it **runs** for a visitor.

BUILD

You sign in, create an agent, give it a persona, connect its knowledge, configure its behavior, and publish it. Done in the Experience Builder.

RUN

A visitor loads the agent (via a URL or an embed). The runtime handles the live conversation — voice, text, tool calls, content cards, session recording. The visitor never touches the Builder.

An agent is composed of

- A **persona** — name, avatar, voice, personality.
- A **knowledge base** — the site's content, uploaded documents, product catalog, FAQs.
- **Tools** the agent can call at runtime — search, cart, custom.
- **Content cards** the agent surfaces when the conversation matches a trigger.
- A **deployment surface** — where the agent runs for end users.

Under the hood

- Real-time voice runs over **WebRTC**.
- The LLM is **Azure OpenAI**, using the Realtime model family for voice.
- All infrastructure runs on **Microsoft Azure**. **SOC 2 Type 2** audited annually.

03 · GET ORIENTED

Sign in and get oriented

Sign in

<https://spaces.napster.ai/> (soon becoming agent.napster.ai)

Sign in with Google, Microsoft, Apple, or email + password. One URL for everyone; your organization is determined by your account.

The two surfaces

Once signed in, you're in one of two places:

- **My Spaces** (soon renamed to **My Agents**) — the dashboard listing every agent you have access to. Search, toggle grid/list view, click into any agent to open its editor.
- **The agent editor** — the per-agent authoring surface where you configure everything. Opened by clicking into an agent.




Figure 1. The dashboard. Create an agent from a URL or prompt at the top; your agents are listed below with search, grid/list toggle, and filters.

Each agent card carries tags for its type (Site, Store), catalog (Shopify, WooCommerce), deployment (Cards), and status (Ready). The : menu on a card gives quick access to common actions without opening the editor.




Figure 2. Agent card menu — Edit, Copy / View Shareable URL, Share Settings, Copy Embed Code, Unpublish, Analytics, Delete.

The organization switcher

A single user account (email) can belong to multiple organizations. The switcher in the header lets you move between them. Data — agents, knowledge, sessions, analytics — is isolated per organization.

The agent editor tabs

Every agent shares the same tab set. Here's what each does; §4 covers each in depth.

TAB	WHAT YOU DO HERE
Avatar	Pick or upload avatars; pick voice.
Space	Background scene for the agent.
Knowledge	Configure knowledge sources — URL, uploaded documents, product catalog.
FAQs	Structured Q&A entries the agent will use.
Welcome	The agent's opening line.
Engagement	Toggles for summary email, lead capture, opt-in disclaimer.
Deflect	Topics the agent should refuse or redirect.
Prompt	The core persona and behaviour prompt.
Content	Content cards — video, image, YouTube, website link.
Preview & Update	Publish, embed code, on-site behaviour controls.

NOTE

The UI still uses "Space" and "Experience" in places. All three (Space / Experience / Agent) mean the same thing: the deployable unit made up of persona + knowledge + tools + deployment configuration.

04 · BUILD

Build your first agent

The fastest path takes 5 minutes. Everything else is refinement.

The 5-minute path

1. Open your dashboard.
2. Paste the **website URL** of the business the agent is for.
3. Pick **Auto** (or **Advanced** if you want to step through every setting yourself).
4. Click **Generate space** (*soon renamed to Generate agent*). Wait 1–5 minutes while the system analyzes the site, picks an avatar and voice, generates a background, builds the knowledge base, drafts the prompt, and publishes.
5. Speak with the agent — click **Preview** to try it. Preview mode doesn't consume minutes, use it freely.
6. Iterate in **Advanced** mode, tab by tab.

TIP · NO WEBSITE?

Skip the URL and use the **Prompt** tab. Describe the business in words — the agent is built from that description instead.

Tab by tab — the depth

Avatar. Pick one or more avatars from the pool, or upload your own (JPG or PNG, 5 MB max, landscape 16:9, face-camera, centre-framed). Multiple avatars per agent is supported — the system assigns one at random per session. Pick a voice from the Azure OpenAI voice library; the picker filters by the avatar's detected gender by default, expand to browse all voices.




Figure 3. The Avatar tab. Upload your own image or pick from the pool; each avatar gets its own voice. The editor tabs run along the top. Click "Update" to save.

Custom voice cloning is possible through custom work — contact your Napster representative.

Space. The background scene shown behind the agent. Auto-generated, but you can upload your own image. Toggle available to turn off the loading-screen morph/zoom animation if you'd rather show the image statically.

Knowledge. Where the agent gets its answers. Sources supported:

- **Website URL** — searched live via OpenAI Website Search. Relies on what's publicly indexed; `noindex` or auth-gated pages aren't visible.
- **Uploaded documents** — PDFs, FAQs, other files. Summarized on ingestion; retrieved by tool when the visitor's question matches.
- **Product catalog** — Shopify MCP, WooCommerce, or a custom feed. See §7.

Multiple sources per agent; the LLM decides which to consult per question. General LLM knowledge is always available as fallback and can't be disabled.

FAQs. Structured Q&A entries — question in, exact answer out. Use this when you want the agent to answer a specific question in a specific way — pricing, opening hours, refund policy, product-specific facts you don't want the LLM paraphrasing. The FAQ answer wins over general answers for those exact questions.

Welcome. The agent's opening line at the start of every session. Auto-generated, editable. You can also add **starter questions** here — example prompts to show the visitor what they can ask.

Engagement. Toggles for optional features:

- **Summary email.** The visitor asks for a recap; the agent captures their email and sends the summary after the session ends.
- **Lead capture.** For after-hours or referral scenarios — the agent fills out a contact form conversationally and forwards it to your email.
- **Opt-in disclaimer.** *Totally optional.* A gate screen with terms/conditions the visitor must accept before speaking with the agent.

Deflect. Topics the agent should refuse or redirect — sit on top of the built-in model safety guardrails. You don't need to add violence, politics, self-harm — those are handled by default. Add topics specific to the business (competitor comparisons, legal advice, medical claims).

Prompt. The generated system prompt — everything from the persona, knowledge, tone, and behavior. You don't have to edit it, but you can if you want tighter control. First outputs rarely nail everything; iterating on the prompt (or adding FAQs) is normal.

Language (*option inside the Prompt tab, not a separate tab*). Default language for the agent. Leave blank in most cases — the agent detects the visitor's language automatically. Force a specific language if you want strict control (e.g., a French client site where the agent should always respond in French).

Content. Cards the agent surfaces inline when the conversation matches a trigger. Each card has:

- *Name* — internal label, never shown to the visitor.
- *Trigger condition* — a natural-language rule ("When the user asks about pricing or costs").
- *Content type* — **Website URL**, **YouTube**, **Video** (upload), or **Image** (upload).
- *Description* — optional short text on the card.

At runtime the agent evaluates triggers against the live conversation and shows matching cards inline in the iframe context. Multiple cards per agent; if two triggers match at once, both cards can appear.

Preview & Update. Where you publish. See §8.

Access control

Set per agent, in Preview & Update under *Share settings*. Most agents are **Open** — that's the default, and it's required if you're embedding the agent on a website (visitors can't be asked to log in to talk to a widget).

- **Open** — public URL, anyone with the link can talk. Use for agents on your website.
- **Invite-only** — login-gated with email allowlisting; invited visitors authenticate before entering. Only relevant for closed-audience use cases (e.g., internal tools, gated pilots).

Crisis and safety

Every new agent automatically includes a **crisis and safety protocol** for handling sensitive interactions (self-harm, medical emergencies, abuse). On by default; you don't configure it per agent.

05 • RUNTIME

What the agent does at runtime

Platform-level capabilities. You don't configure most of these per agent — they just work.



Figure 4. An agent in Preview Mode (Private). The control bar at the bottom holds mic, file upload, text chat, screen share, stop, and hang-up.

Voice + text. Real-time voice-first conversation with optional text chat. Visitors can switch modes within a session across all deployment surfaces. Mic, speaker, and hang-up controls sit at the bottom of the conversation surface.

File uploads. Visitors can upload files (PDFs, images) during a session for the agent to review and respond to in context. Useful for support (share an error screenshot) or guided workflows (upload a document you want feedback on).

Screen awareness. Visitors can share their screen. The agent sees the active screen and can reference it during the conversation — real-time visual assistance, guided workflows, screen-based support.

KNOWN LIMITATION

Stopping and re-enabling screen share mid-session may cause the agent to hallucinate. Recommend a new session instead.

Conversation memory.

- *In-session* — the agent remembers prior turns of the current conversation.
- *Cross-session (authenticated visitors)* — logged-in visitors on invite-only agents get past-session context.
- *Anonymous visitors* on open agents have no memory across sessions — every session starts fresh.

Multi-language. Auto-detects the visitor's language and switches. Same agent serves a multilingual audience — no separate agent per language.

Pre-session loading states. Between the visitor clicking Start and the agent becoming ready, the UI walks through: initial prompt → mic permission check → agent setup → almost ready → long-wait recovery if init takes longer than expected. Phase labels are localized.

Idle detection. If the visitor stops interacting, the agent asks "are you still there?" and shuts down. No minutes wasted.

Conversation summary. Optional per agent (toggle in Preview & Update). When on, visitors see a Summary button, enter their email, and receive a recap after the session. Works in preview and live.

Agent tools

The specific actions an agent can take at runtime. The LLM decides which tool to invoke based on the visitor's question and the prompt.

Built-in tools

These tools are available out of the box in the product. Which ones actually get invoked depends on the agent's type and which features you've enabled — the LLM picks the right tool for each visitor question.

TOOL	WHEN IT'S INVOKED
Knowledge base search	Whenever the visitor's question matches content in your uploaded documents.
Web search (URL)	Whenever the visitor asks something covered on the configured website. Uses OpenAI Website Search.
Product search	Only if the agent runs against a Shopify or WooCommerce store, and the visitor asks about products.
Cart operations	Same as product search — only if the agent runs against a Shopify or WooCommerce store, and the visitor wants to add / update / remove items.
Content card trigger	Whenever a content card's trigger condition matches the live conversation. Surfaces the matching card inline.
Lead capture	When lead capture is enabled and the visitor asks to be contacted (or the agent decides a lead should be captured).

You influence tool selection through the prompt — you can't force a specific tool at build time.

Custom tools

Custom tools are not configured through the Builder UI. If you need one — a specific data lookup, an escalation to a human on your side, an action into your own system — contact your Napster representative.

Commerce

Native commerce lets the agent search a merchant's product catalog, manage a shopping cart during the conversation, and hand off to the merchant's native checkout. **Napster Agent never touches payments** — cart is agent-internal, checkout is on the merchant's side.

Shopify

Connects via Shopify's public storefront protocols. No setup, no authentication needed — the endpoints are built into every Shopify store.

- **Catalog search** — via the Storefront MCP endpoint (`https://{shop}.myshopify.com/api/mcp`). Tool: `search_catalog` .
- **Cart operations** — via the newer **UCP Cart MCP** endpoint (`https://{shop}.myshopify.com/api/ucp/mcp`). Tools: `create_cart` , `get_cart` , `update_cart` , `cancel_cart` . Shopify moved cart tools to UCP in June 2026; the older Storefront MCP cart tools are deprecated and will stop working after August 31, 2026.

Configuration: Shopify store URL per agent. No API keys.

Caveat: the store owner can disable the public storefront API — if they do, the integration breaks.

WooCommerce

Connects via WooCommerce's REST API. Agent queries the catalog, surfaces results as product cards, manages an internal cart, hands off to the store's native WooCommerce checkout.

Configuration: WooCommerce store URL per agent. Requires the store to permit API access for catalog and cart operations.

KNOWN LIMITATIONS

- **HTML entities** in product text sometimes leak through in long-form descriptions.
- **If the store has restricted the cart or checkout API**, hand-off fails.

Other catalog formats

For stores not on Shopify or WooCommerce — all paths require contacting us; there's no self-serve upload in the Builder for these today.

- **Product JSON** or **Google Shopping XML** — we can ingest these into the vector store, but there's no UI to upload them yourself. Send us the file and we'll get it in.
- **Custom format** — a code update is needed to support a new format. Contact your Napster representative.
- **Client's own product API** — API integration is custom work: we build a translation layer per client (an OpenAI call takes the conversation + your API filters and translates the visitor's natural-language query into accepted API parameter values). Talk to your Napster representative to scope this.

Cart behaviour (universal)

The agent generates its own internal cart during the conversation. When the visitor is ready to purchase, the agent redirects to the merchant's native checkout. **There is no write action to the merchant's cart, order, or inventory systems** — the cart is entirely agent-internal.

Publish and share

Publish at **Agent editor** → **Preview & Update**.




Figure 5. The publish panel. Custom URL and deployment mode on the left, live preview on the right. "Copy URL" copies the standalone link.

Deployment surfaces

SURFACE	DESCRIPTION
Standalone URL	Every agent has a dedicated URL. Shareable, bookmarkable. Best for demos.
Embedded (Full Experience)	Iframe embed of the complete experience — agent, content, and UI.
Embedded (Agent Only)	Iframe embed of just the agent, no surrounding UI. Supports content cards.
Cards (floating)	<i>Coming soon.</i> Floating agent that answers with cards for products and pages.

The Preview & Update panel

- **Custom URL** — the standalone agent URL (editable).
- **How it appears** — pick a deployment mode (see above).
- **Embed Code** — the `<script>` tag to paste into a website. Copy button.
- **Settings** — Share settings (Make Experience Private), Cards Display.

IMPORTANT

Any change made in the editor needs an **"Update"** click at the bottom to push to the live experience. Nothing propagates automatically.

More options are coming soon (agent activity, installation help, on-site behaviour controls). This section will grow as those land.

Analytics

Agent editor → Analytics (per-agent) or top-right icon → Analytics (org-wide).

A shared date filter (7d / 30d / 60d / 90d / custom / all) and Export Report sit at the top and apply to every block.

Org-wide analytics — all agents




Figure 6. Org-wide analytics. Totals across the top; a per-agent performance table below, filterable by session or creation date.

Opened from the top-right icon → Analytics. One view across every agent in your organization.

- **Top stat cards** — Agents created, total Interaction Minutes, and total Sessions for the selected period.
- **Performance table** — one row per agent with Number of Sessions, Total Interaction Minutes, Average Session Length (min), and Creation Date. Sortable; paginated.
- **Session date / Creation date** — toggle whether the date range filters by when sessions happened or when agents were created.

- **Date range + Export Report** — pick the period; export the table.

Click into any agent from here, or use **Analytics** in its card menu, to see the per-agent view below.

Standard analytics — every agent

- **Top stat cards** — Sessions, Average Session duration, Total Interaction time.
- **Sessions list** — date, time, length, AI-generated summary, transcript viewer per row.
- **Active sessions filter** — toggle labelled *"Active sessions only — At least one user input captured"*. Scopes every block to sessions where the visitor typed or spoke. Default on.
- **Export to CSV.**

Shopping Analytics — commerce-enabled agents

Adds a full commerce funnel:

- **Revenue captured** — top KPI + funnel + per-product. Currency follows the store. "Captured" is measured up to checkout hand-off; not settled revenue.
- **Funnel** — Products Discussed → Products in Cart → Products Checked Out. Each stage counts unique products. Inter-stage and overall conversion %.
- **Sessions table** — one row per session with per-session product counts. Four filter chips (All / Discussed / In cart / Checked out) scope the table. Search box scans summaries; Export All Transcripts downloads conversation history for the current scope.
- **Per-session Detail panel** — opens a modal titled *"Products in this session"* with session ID, outcome, and a per-product table showing In-cart and Checked-out quantities plus a **Checked out / Dropped** status badge.
- **Top Products table** — per-product breakdown with sessions, units, revenue captured, and conversion %. Export CSV.

KNOWN LIMITATIONS

Tracks up to checkout hand-off only (no post-hand-off revenue, refunds, or completion); no UTM / marketing attribution; cohort segmentation limited to date range and active-sessions toggle.

Known limitations

What to be aware of. Sorted roughly by impact.

Platform / model

- **Double answers on tool-augmented queries.** When a tool takes time to return, the model produces an initial response from memory, then a second response once the tool result arrives. Inherent to how the models are trained; improvement depends on the model provider.

Avatar

- **Lip-sync quality.** Not yet at target quality. A new visual model is in active development. **Tip:** avoid avatars framed with big close-up faces — lip movement is much more noticeable at that scale. A slightly wider frame (waist-up or bust) hides the issue better.
- **Screen share re-enable.** Stopping and re-enabling screen share mid-session may cause the agent to hallucinate — recommend a new session.

Commerce

- **Cart-add reliability.** The agent occasionally reports a successful cart add when items aren't actually added.
- **No customer account access.** Visitors are anonymous during conversations. No order history, no OAuth pass-through.

Deployment

- **Cross-page persistence.** True persistence across a traditional multi-page website is not deliverable today. An intercom-style MVP is in progress.
- **Client-specific UI customizations can disappear after updates.** Manual maintenance required for clients with heavy custom UI.

Analytics

- No post-hand-off revenue tracking.

Frequently asked questions

Can the agent speak languages other than English?

Yes. Multi-language with in-session switching. Same agent serves a multilingual audience.

Can we use a custom avatar or voice?

Custom avatar images can be uploaded (5 MB max, landscape 16:9). Custom voices are possible through custom work. Standard voices come from the Azure OpenAI voice library.

Does the agent scrape the client's website?

No. URL search uses OpenAI Website Search — relies on what's publicly indexed. `noindex` or authenticated pages aren't visible.

Can the agent authenticate the visitor?

Not for open agents. For invite-only agents, the visitor authenticates before entering. No OAuth pass-through to the client's site.

Can the agent work across multiple pages of the client's site?

Yes for SPA sites — persistence works well. For traditional multi-page sites, true persistence isn't deliverable today (see §10); an intercom-style MVP is in progress.

What LLM does the agent run on?

Azure OpenAI, using the Realtime model family for voice.

Is client data used to train models?

No. Napster Agent uses Azure OpenAI in a mode where data is not used for training. SOC 2 Type 2 audited annually.

Can we export session transcripts?

Yes. Per-row transcript viewer in the Sessions list; Export All Transcripts for the current Sessions scope.

Can the agent do product cards for non-Shopify / non-WooCommerce stores?

Yes — via Product JSON, Google Shopping XML, or a custom API translation layer. See [§7](#).

How do we handle escalation to a human?

Via a custom tool per client. Not out-of-the-box today.

Glossary

TERM	DEFINITION
Agent / Experience / Space	The core deployable unit. One persona + one set of knowledge + one deployment configuration.
Azure OpenAI	Our LLM provider.
Content card	A video, image, YouTube link, or website link the agent surfaces inline when a conversation trigger matches.
Deployment surface	Where the agent runs for end users (standalone, embed, partner).
Embed code	The <code><script></code> tag pasted into a client's website to deploy the agent.
Experience Builder	The authoring UI where creators configure agents.
Knowledge base	The sources the agent can consult: prompt, URL search, uploaded docs, FAQs, product catalog, general LLM knowledge.
Napster API	The developer API layer exposing the agent runtime for custom integrations.
Napster Spaces	The earlier name for Napster Agent.
Organization / org	A tenant. Users, agents, subscriptions, and analytics scoped per org.
Persona	The agent's identity: name, avatar, voice, personality, tone.
Preview mode	The internal test view of an agent. Sessions don't consume billable minutes.
Runtime	The live conversational engine handling agent-visitor interactions.
OpenAI Website Search	The search service used for URL search over publicly indexed content on the configured website.
Shopify MCP	Shopify's public Storefront Model Context Protocol endpoint for catalog and cart operations.
Tool	A specific action the agent can call at runtime (KB search, URL search, product search, cart, custom).

TERM	DEFINITION
Vector store	Storage used for ingested product catalog and knowledge base content.
WebRTC	The real-time communication protocol used for low-latency voice.