



SALES & PARTNER PLAYBOOK

Botlit

Botlit lets a support, HR or IT team put governed AI agents on Slack, Teams and the web — and see exactly what every answer cost.

Agentic AI

Multi-provider models

Governed by default

Per-run cost audit

Slack & Teams



PREPARED BY

Burdenoff Consultancy Services Pvt. Ltd.

PRODUCT

www.botlit.ai

BOOK A CALL

calendly.com/vignesh-burdenoff/1-1

00 How to use this playbook

Read the opportunity signals and the one-question test first; they will tell you in a single conversation whether a prospect is worth pursuing. Keep the elevator and the persona hooks close to hand for live conversations, and use the objections section before you promise anything about compliance or runtime maturity. When a conversation gets serious, use the handoff checklist rather than trying to sell the technical detail yourself.

IN A HURRY

Read sections 01–03 and 18. That is enough to recognise an opportunity and open a conversation.

BEFORE A MEETING

Add 06–08, 10 and 16 — who buys, what they buy, what it costs, and what they will push back on.

HANDING OVER A LEAD

Sections 19–22 tell you exactly what to send, to whom, and what you earn for it.

CONTENTS

01 Product Overview

02 One-Line Pitch

03 30-Second Elevator Pitch

04 Problem We Solve

05 Value Proposition

06 Ideal Customer Profile

07 Buyer Personas

08 Top Use Cases

09 Product Capabilities

10 Pricing & Plans

11 Competitive Landscape

12 Why Us / Differentiators

13 Customer Proof & Credibility

14 Demo Information

15 Sales Talking Points

16 Objection Handling

17 Discovery / Qualification Questions

18 How to Identify a Good Opportunity

19 How to Introduce Us

20 Lead Handoff Process

21 Partner / Referral Commercials

22 Contact Information

23 Useful Links & Sales Assets

24 Product Mockups — Complete Visual Appendix

01 Product Overview

WWW.BOTLIT.AI

Botlit is a platform for building, governing and deploying AI agents. A team creates a bot app — the user-facing assistant, such as a Customer Support Bot — and attaches specialist agents behind it, each with its own instructions, its own model and its own knowledge. Those agents answer from knowledge bases built out of the organisation's own documents, and they can be published to web chat, Slack, Teams, Discord, WhatsApp and webhooks. Every run the platform executes is written down with its status, its inputs and outputs, its token counts and its cost. Because Botlit is a module of the Burdenoff Workspaces platform rather than a standalone tool, roles, permissions, audit, billing and file storage are already in place around it.

CATEGORY

Enterprise agentic AI platform, for mid-market and enterprise teams that need AI assistants under real governance rather than a bolt-on chatbot.

WHERE IT STANDS TODAY

Pre-launch. The site runs in waitlist mode, checkout is not enabled, and there are no reference customers. The management surface, the agent runtime, the execution ledger and the permission model are built and exercised by automated end-to-end tests; several runtime capabilities are deliberately labelled as partial, and compliance certification has not been achieved.

HOW CUSTOMERS GET IT

Cloud, multi-tenant, no infrastructure for the customer to run. Each customer gets a workspace that is the tenancy boundary for every agent, knowledge base and conversation, and agencies typically take one workspace per client. Model inference runs on the customer's own provider account, using keys they register and can revoke. Onboarding today is a guided early-access conversation rather than self-serve signup.

Web console for building, testing and operating agents

Web chat, Slack, Teams, Discord, WhatsApp and webhook channels

Shared team inbox where every bot conversation lands beside human messaging

GraphQL API for embedding agents inside the customer's own product

Embedded visual workflow builder for multi-step automations

Documentation site at docs.botlit.ai

Desktop, mobile, CLI, SDK and editor surfaces are specified but not yet released

02 One-Line Pitch

SAY EXACTLY THIS

"Botlit lets a support, HR or IT team put governed AI agents on Slack, Teams and the web — and see exactly what every answer cost."

SHORTER, FOR A CHAT MESSAGE

Governed AI agents for support and internal teams, with a receipt for every run.

THE CATEGORY SENTENCE

Botlit is an enterprise agentic AI platform for organisations that need AI assistants their security and finance people will sign off on.

03 30-Second Elevator Pitch

SPOKEN, ROUGHLY 30 SECONDS

"Most teams answering the same questions all day have tried a chatbot, and most of them switched it off. It either made things up, or nobody could tell what it cost. Botlit is an agentic AI platform. You build a bot app, ground it in your own documents, point it at whichever model provider you already pay for, and put it on Slack or Teams. Every single run is recorded with its tokens and its cost, so the answer to "what is this thing doing" is a query, not a guess."

IF THEY ARE TECHNICAL

Botlit separates the user-facing bot app from the specialist agents behind it, and routes between them at run time using keyword and pattern scoring with a deterministic fallback. Agents bind to models from eleven provider families whose credentials are vaulted and resolved at execution rather than stored in the agent definition, retrieval is injected into the turn from your own knowledge bases, and every operation on the GraphQL API carries a permission rule enforced at the gateway before a resolver ever runs.

IF THEY ARE A BUSINESS BUYER

Repetitive enquiries are a headcount problem that nobody wants to solve by hiring. Botlit lets an existing team stand up assistants that answer from company documents, deploy them on the channels people already use, and price the result honestly, because every run carries its own token and cost record. The governance a security review asks for — roles, permissions, audit, tenant isolation — is inherited from the platform rather than bolted on afterwards.

04 Problem We Solve

These are the five or six conversations that reliably turn into a Botlit opportunity. If a prospect describes any two of them unprompted, keep going.

The chatbot invents answers

Who feels it: Support leaders and the people who own the brand

What it costs them: A confident, fluent, wrong answer about a refund policy costs more goodwill than no answer at all, so the bot gets quietly switched off and the deflection savings never arrive.

How Botlit answers it: Botlit retrieves from the agent's own knowledge bases inside the turn, prepends the retrieved passages to the agent's instructions, and tells the model in the same breath to say it does not know rather than fabricate. Only knowledge bases whose index has finished building are consulted, and a retrieval failure degrades to an ungrounded answer with a logged reason instead of an error in front of a customer.

Nobody can say what the AI costs

Who feels it: Finance, and the operations lead who has to defend the line item

What it costs them: Spend is discovered on the invoice. A longer prompt, a model swap or a retrying integration can triple a bill silently, and there is no way to attribute cost to a team, a client or a use case.

How Botlit answers it: Botlit writes an execution record for every run — successful, failed, timed out or cancelled — carrying the agent, the model, the duration, the input and output tokens, the resolved cost, the number of tool calls and the number of retrieved passages. The same data aggregates per agent and per bot app, so cost per client or per use case is a query rather than a reconstruction from logs.

Prompt changes with no way back

Who feels it: Whoever owns the assistant's behaviour, usually a builder or ops lead

What it costs them: Agent behaviour lives in a system prompt that three people edit. When answers start hedging, nobody can say what changed, and the previous version exists only in someone's chat history.

How Botlit answers it: Agents and bot apps are versioned. You iterate in draft, publish a version with a changelog, and roll back to the previous published version when a change regresses. Versions move through draft, published and historical states, so the lineage of every assistant is inspectable rather than folklore.

The AI stack is locked to one vendor

Who feels it: CTOs, VPs of Engineering and procurement

What it costs them: Model pricing, model quality and data-residency choices all belong to a single supplier. Negotiating leverage is nil, and moving cheap tasks to a cheaper model means a rewrite.

How Botlit answers it: Providers are registered per workspace across eleven families — OpenAI, Anthropic, Google, Cohere, Mistral, Groq, Together, HuggingFace, Azure OpenAI, AWS Bedrock and custom endpoints, which is also how a self-hosted model server is connected. Credentials are vaulted and resolved at execution, never written into an agent definition, so swapping a model binding does not mean rewriting the agent.

Governance is a promise, not a control

Who feels it: Security reviewers, compliance officers and IT admins

What it costs them: The pilot passes because it is a pilot. The rollout stalls in security review because nobody can demonstrate who may do what, which tools an agent may reach, or what it actually did.

How Botlit answers it: Every query and mutation on the Botlit API carries a permission rule that the gateway enforces from the composed schema before a request reaches any resolver, so authorisation is not reimplemented feature by feature. Workspace scoping isolates tenants in the data model, five seeded roles cover the platform admin, builder, knowledge curator, operator and auditor personas, and guardrail policies for rate limits, budgets and content are evaluated in the request path.

Every client engagement starts from scratch

Who feels it: Digital agencies, consultancies and internal platform teams

What it costs them: The same support bot is rebuilt for the fifth client with the same prompt archaeology, so margin never improves with experience and good work is not an asset.

How Botlit answers it: Botlit distributes ten kinds of package — apps, agents, tool packs, prompt templates, skills, retrieval profiles, policy bundles, channel templates, agent-to-agent definitions and knowledge-base templates. Each installs linked, so it keeps tracking the author's improvements, or forked, so the client owns the copy outright, and a linked install can be forked later without losing configuration.

05 Value Proposition

Five things a partner can say about Botlit that hold up under questioning, each with the thing in the product that backs it.

The model strategy stays the customer's

Botlit does not resell tokens and does not have a house model. The customer registers their own provider accounts, tests each connection before an agent depends on it, and binds different agents to different models. A cheap model can serve low-risk internal questions while a stronger one handles customer-facing work.

Evidence: Eleven provider families in the workspace registry, credentials vaulted and resolved at execution rather than stored in the agent, plus a custom endpoint type for a model server the customer hosts themselves.

Every answer has a receipt

The execution ledger is written by the runtime, not reconstructed from logs, and it covers runs that failed as well as runs that succeeded. That makes the awkward questions answerable: which agent's cost per run tripled and after which published version, which runs answered with nothing retrieved, which bot has not run at all this month.

Evidence: Per-run status, input, output, duration, model, input and output tokens, resolved cost, tool-call count and retrieved-passage count, aggregated per agent and per bot app.

Governance arrives with the product

Botlit is a module of a multi-tenant workspace platform, so roles, permissions, audit trail, tenant isolation, billing, notifications and file storage are already there. A security reviewer is not being asked to trust a young product's own access control; they are looking at the platform's, applied to every operation.

Evidence: A permission rule on every query and mutation across twenty-one API modules, enforced at the gateway; workspace-scoped data throughout; five seeded roles from platform admin to read-only governance auditor.

Guardrails run in the request path

A policy in a document is a hope. Botlit evaluates rate limits first, then budget against the day's consumption, then the incoming message against blocked keywords, patterns and opt-in personal-data heuristics — and evaluates the model's response the same way before it is returned. A blocked response produces one explicit refusal, identical in the app and on Slack.

Evidence: Runtime evaluation of rate-limit, budget, content and agent-trust policies, fail-open by design so a workspace that has configured no policy is unaffected.

Good work becomes a reusable asset

For agencies and platform teams this is usually the commercial argument. A support bot refined on one engagement is captured as a package carrying its agents, published versions, prompt templates and retrieval settings, then installed per client. Harden the prompt once and every linked client can take the improvement.

Evidence: Ten package kinds with linked and forked installs, fork-after-install for a linked package that outgrows the standard build, and workspace-per-client isolation with its own permissions, knowledge and billing.

06 Ideal Customer Profile

A 200 to 2,000-person company where one team fields a high volume of repeat enquiries, already pays a model provider, and has a security or finance function that will ask hard questions before anything goes live.

SEGMENT	TYPICAL SIZE	WHY THEY BUY
Mid-market operations, support and HR teams	100–1,000 employees	They have the enquiry volume to justify automation and no machine-learning team to build it. Templates, knowledge bases and channel connectors get them to a working assistant without hiring, and the cost ledger lets them prove the case internally.
Enterprise platform and IT teams	1,000+ employees	They are being asked to sanction agent deployments across several departments and need permissions, audit and tenant isolation to be structural rather than promised. Bring-your-own-model matters here because model procurement is already a negotiated contract.
Digital agencies and consultancies	5–50 people, several concurrent clients	Workspace-per-client isolation plus linked and forked package installs turn a bespoke build into a repeatable service line, and per-client execution records make it possible to price the engagement on real cost.
Technology startups and scale-ups	50–500 employees, engineering-led	They want an in-product assistant quickly and will drive it through the API rather than the console. Versioned agents, prompt templates and per-run telemetry are what they would otherwise spend a quarter building.
Professional services and financial services firms	100–1,000 employees	Knowledge management and client intake are document-heavy problems that suit retrieval, and these buyers care a great deal about who may see what, which is where the permission model earns its keep.

GEOGRAPHIES

Primary focus is North America and Europe, with India and the wider APAC region served directly from the Chennai team. There is no data-residency guarantee or in-region deployment option yet, so a buyer with a hard residency requirement should be flagged early rather than pursued.

NOT A FIT — DO NOT CHASE

- The buyer needs a SOC 2 report, ISO 27001 certificate or a signed HIPAA business associate agreement before they can sign — none of these exists yet, and no release will create one.
- Voice telephony is a core requirement. Botlit has no voice channel; Voiceflow, Botpress Enterprise and Intercom are the honest answers there.
- The organisation is wholly committed to Microsoft 365 and Azure and procurement prefers adding a Microsoft licence. Copilot Studio will win that on paper and on paperwork.
- They want a polished drag-and-drop conversation canvas today and will judge the product on builder aesthetics rather than governance. The visual builder is not yet at Voiceflow's standard.
- Self-hosting or an air-gapped install is mandatory. Botlit is cloud multi-tenant only, although inference itself can run on a model server the customer hosts.
- There is no repeat-enquiry volume and no document corpus — a team asking a handful of unique questions a week will not see a return.

07 Buyer Personas

Five roles turn up in Botlit conversations. Each cares about a different half of the product, and each has a question you should expect within the first ten minutes.

CTO or VP of Engineering

WHAT THEY WANT

Add AI capability without taking on a bespoke runtime to maintain, keep model procurement negotiable, and get the thing through security review the first time.

WHAT FRUSTRATES THEM

The engineering team is already the bottleneck for internal requests, custom AI builds have a habit of running for months, and single-vendor model dependency is a commercial risk nobody wants to explain to the board.

Open with: *If your model provider doubled its price next quarter, how much of your AI work would you have to rewrite?*

They will ask: What exactly am I locked into, and can we drive all of this from the API rather than your console?

Head of Customer Support or Service

WHAT THEY WANT

Deflect the repetitive half of the queue without degrading the answers customers get, and see what the bot actually said before anyone escalates.

WHAT FRUSTRATES THEM

Volume grows faster than headcount, the previous chatbot embarrassed the team, and bot conversations live somewhere the support team cannot see.

Open with: *What share of your tickets this month were the same handful of questions asked in different words?*

They will ask: What happens when the agent does not know the answer, and how does the conversation reach a human?

Compliance officer or security reviewer

WHAT THEY WANT

Establish who may deploy an agent, what data it can reach, what it may say, and whether any of that can be demonstrated after the fact.

WHAT FRUSTRATES THEM

AI pilots arrive already built, with governance improvised afterwards. Certification questions get answered with roadmap language, and nobody can produce an audit trail on request.

Open with: **If your regulator asked what your AI assistant told a customer last Tuesday, could you show them?**

They will ask: Do you have SOC 2, and where exactly does our data go once an agent answers?

Operations or platform lead

WHAT THEY WANT

Run the agent estate day to day: publish versions, watch executions, catch a failing channel before customers do, and keep spend inside a budget.

WHAT FRUSTRATES THEM

There is no single place to see what the agents are doing, cost is invisible until the invoice, and every new use case means another tool.

Open with: **Where do you go today to find out which of your automations is costing the most and returning the least?**

They will ask: Can I see a failing run in detail, and can I stop an agent spending beyond a limit?

Agency founder or delivery lead

WHAT THEY WANT

Sell an AI assistant as a repeatable service line, deliver it per client without rebuilding, and price it against a cost they can actually see.

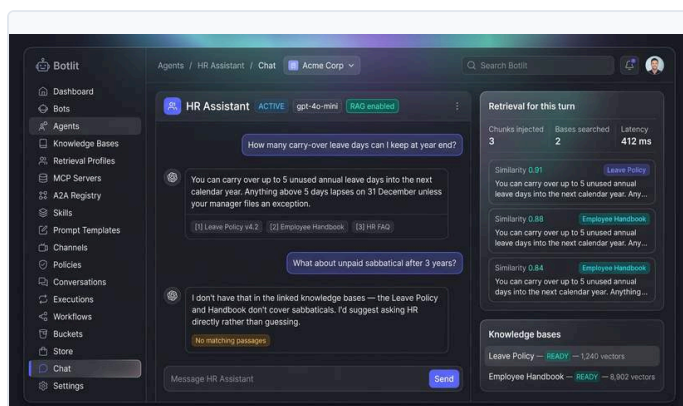
WHAT FRUSTRATES THEM

Every engagement starts from a blank page, client data has to be kept apart by discipline rather than architecture, and margin never improves with repetition.

Open with: **How many times has your team built essentially the same support bot for a different client?**

They will ask: Can I keep clients fully separated, and can I push an improvement to all of them at once?

Thirteen situations where Botlit earns its place. The first seven describe what happens inside a single agent turn; the rest are the workflows around it. Anything still on the roadmap is labelled as such.



KNOWLEDGE, SUPPORT OR ENABLEMENT LEAD

Answers grounded in the company's own documents


THE SITUATION

The answers people need are spread across a policy handbook, a help centre and three shared drives. A general-purpose model asked about the refund window will produce something fluent and wrong, which is worse than silence.

WHAT HAPPENS WITH BOTLIT

An agent declares which knowledge bases it uses. On every turn Botlit embeds the question with the same embedding model that base was built with, searches that base's own vector collection, takes the best passages, numbers them and puts them in front of the model with a standing instruction to use them where relevant and admit ignorance otherwise. Bases still indexing are skipped rather than half-searched.

The assistant answers from the organisation's material, and says it does not know when the material does not cover the question.



STARTUP BUILDER OR PLATFORM ENGINEER

Agents that do things, not just describe them


THE SITUATION

The assistant can explain how to check an order status but cannot check one, so every conversation ends by asking a human to do the actual work.

WHAT HAPPENS WITH BOTLIT

The organisation registers the tool servers it already runs. Botlit opens a real session against a server, lists its tools, and offers them to the model on every turn; when the model asks for a tool it is invoked, the result is fed back, and the model is called again until it answers or the iteration budget runs out. Outbound targets are checked against a strict endpoint guard, and locally-executed tool commands are refused by design because a shared platform must not run tenant-supplied shell commands.

Agents complete tasks inside the conversation, with a bounded loop instead of an unbounded bill.



SUPPORT OPERATIONS LEAD

The right specialist takes each message


THE SITUATION

One assistant with a two-page prompt is trying to be a billing expert, a technical expert and an escalation path at once, and it is competent at none of them.

WHAT HAPPENS WITH BOTLIT

A bot app fronts several specialist agents, each attached with a priority, trigger keywords, intent patterns and a fallback. The app declares a strategy — single, round-robin, rule-based or intelligent — and the intelligent strategy scores the message against keywords and patterns first, escalating to a model classifier only when scoring is inconclusive. When the classifier is unsure it falls back to a configured destination rather than guessing.

Each message reaches a specialist that is actually good at it, and the routing decision is on the record afterwards.



ENTERPRISE AUTOMATION ARCHITECT

Agents that hand work to other agents, on a leash

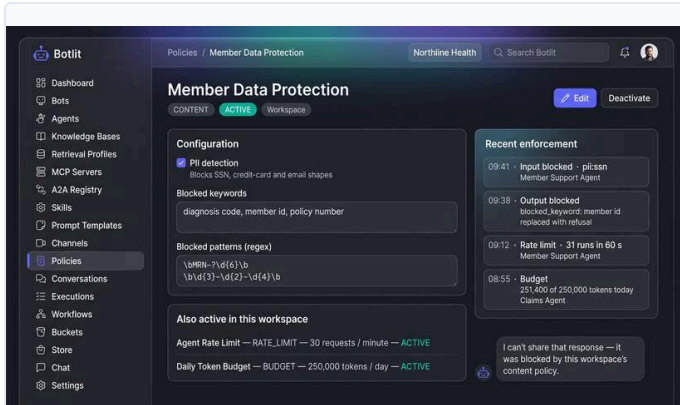
THE SITUATION

A single agent cannot both research a question and act on the answer, but wiring agents together usually means a bespoke orchestration layer nobody wants to own.

WHAT HAPPENS WITH BOTLIT

Delegation runs the target agent for real — its own model, its own knowledge, its own tools — and returns the genuine response to the caller, recording the exchange. It can be requested through the API or by the model mid-conversation. Delegation depth is capped, a cycle is rejected outright so one agent cannot loop back to another, and connections are gated on the trust level the workspace requires.

Specialists collaborate on a question, with bounds and an audit record rather than an unbounded chain.



COMPLIANCE OFFICER OR RISK REVIEWER

Guardrails that stop the model rather than describe it


THE SITUATION

The AI usage policy is a document. Nothing in the request path reads it, so the first time it is tested is when something goes wrong in front of a customer.

WHAT HAPPENS WITH BOTLIT

Policies written once per workspace are read by the runtime on every turn, cheapest check first: rate limit, then budget against the day's consumption, then the incoming message against blocked keywords, supplied patterns and opt-in personal-data heuristics. The model's response is evaluated the same way, and a violation returns one explicit refusal — the same wording in the app and on a channel — rather than leaking a partial answer.

A written rule becomes an enforced control, and a blocked interaction fails in a way a customer can read.



OPERATIONS LEAD, FINANCE PARTNER OR AGENCY OWNER

Know what every agent run costs


THE SITUATION

AI spend goes wrong quietly. A prompt gets longer, a model gets swapped, an integration starts retrying, and the first anyone hears about it is the invoice.

WHAT HAPPENS WITH BOTLIT

Botlit writes an execution record for every run — web chat, channel message, delegated task or workflow-triggered — whether it succeeded, failed, timed out or was cancelled. Each carries the agent and app, the input and output, the status, the duration, the tokens, the resolved cost, the model and structured error detail, alongside counts of tool calls and retrieved passages. The same records roll up per agent and per app.

Cost per agent, per client or per use case becomes a number rather than an estimate.



AGENCY FOUNDER OR PLATFORM OWNER

A copilot that drafts agents and a supervisor that watches them (roadmap)


THE SITUATION

Going from "we want a support bot" to a working one means a provider, a model, a knowledge base, documents, an index, an agent, a prompt, an app, a link, a publish, a connector and a route. That cliff is where most evaluations stall.

WHAT HAPPENS WITH BOTLIT

The design turns that into one sentence: describe the bot and Botlit drafts the whole graph as drafts a human approves entity by entity, preferring a proven template over a fresh guess. A scheduled supervisor then reviews the fleet and raises decision cards — failure-rate regression attributed to the publish that caused it, grounding decay, cost anomalies, stalled indexes, idle agents — with the evidence attached. Nothing auto-applies.

Roadmap, and labelled as roadmap on the public site. Useful to show direction; never sell it as available.



MID-MARKET OPERATIONS LEAD

From signup to a working assistant in one session


THE SITUATION

Previous AI trials ended in a sandbox with nothing in common with how the team would actually run an assistant, so the pilot never converted into anything real.

WHAT HAPPENS WITH BOTLIT

The workspace is created with roles, audit and billing already wired in. A curated template — customer support, HR, finance and others — becomes a real, editable bot app and agent in one action, pre-wired with sensible instructions. The app carries a type, a lifecycle status from draft through active to archived, and a visibility setting controlling who inside the organisation can see it.

The first session ends with a persisted assistant inside the real security model, not a demo sandbox.



AGENCY FOUNDER OR DELIVERY LEAD

Build once, deliver to many clients


THE SITUATION

The team has built a good support bot three times for three clients. None of that work compounds, and client data is kept apart by discipline rather than architecture.

WHAT HAPPENS WITH BOTLIT

Each client gets their own workspace, isolated by construction. The agency's best build is captured as a store package carrying its app, agents with published versions, prompt templates and the retrieval settings that made answers feel right. Fast-follow clients take a linked install that tracks the agency's improvements; bespoke clients take a fork they own, or a fork-after-install when a linked client outgrows the standard build.

One hardened build becomes a service line, with per-client cost records to price it against.



BUILDER OR PRODUCT OWNER

Change an agent's behaviour and be able to undo it

THE SITUATION

The new instructions looked better in testing, then started hedging in real conversations. Nobody can say precisely what changed, and the old wording is gone.

WHAT HAPPENS WITH BOTLIT

Agents and apps are treated like code. You iterate in draft, publish a version with a changelog describing what changed and why, and roll back to the previous published version in one action. Versions move through draft, published and historical states, and each agent accumulates its own statistics, so version seven can be compared with version six on execution evidence.

Behaviour changes become reversible decisions with a paper trail instead of unrecoverable edits.



ENTERPRISE IT OR PLATFORM OWNER

Use the model contract you already negotiated


THE SITUATION

The organisation has an existing account with a model provider, opinions about where inference may run, and no appetite for a platform that quietly resells tokens at a markup.

WHAT HAPPENS WITH BOTLIT

Providers are registered per workspace across eleven families, each with its own endpoint and key, and each verified with a connection test before an agent depends on it. Keys are vaulted and resolved at execution rather than written into an agent definition. Different agents can bind to different models, and a model server the organisation hosts itself connects through the custom endpoint type.

Model choice stays a procurement decision, and inference costs stay on the customer's own provider bill.



SUPPORT OR INTERNAL COMMUNICATIONS LEAD

Put the assistant where people already are


THE SITUATION

The assistant lives on a page nobody visits. Employees ask in Slack and customers ask on WhatsApp, and neither is going to change habit for a bot.

WHAT HAPPENS WITH BOTLIT

Connectors bind the workspace to web chat, Slack, Teams, Discord, WhatsApp, webhooks or a custom platform, and routes map a channel or thread to a specific bot app. Inbound messages are acknowledged instantly and the agent runs in the background, so a webhook never waits on model latency. Outbound replies are stored and delivered in parallel.

The assistant answers in the tools people already use, without a channel timing out while a model thinks.



KNOWLEDGE CURATOR OR ENABLEMENT LEAD

Turn a document pile into answerable knowledge

THE SITUATION

The answers exist, in documents. Pasting them into a prompt does not scale and produces a different answer depending on who pasted what.

WHAT HAPPENS WITH BOTLIT

Documents go into a knowledge base and are chunked with a chosen strategy — fixed windows, recursive splitting, semantic, Markdown-aware or code-aware — then embedded in batches through the workspace's own provider, with vector dimensions validated at ingest so a mismatch is rejected rather than discovered later as poor recall. Retrieval settings are saved as reusable profiles that can be attached to agents and shared as packages, and semantic search proves retrieval quality before anything goes live.

Knowledge becomes a maintained asset with measurable retrieval quality, not a prompt somebody pasted once.

09 Product Capabilities

What is built, what is partly built, and what is still design. The renderer marks anything that is not shipped, and you should too — the honesty is the point.

Building agents

The bot app is the front door; agents are the specialists behind it. Both are versioned objects rather than editable text fields.

Bot apps

Create, edit and archive user-facing assistants with a type, a lifecycle status from draft to archived, and a visibility setting controlling who in the organisation can see them.

Specialist agents

Internal workers with their own instructions, model binding and declared capabilities such as retrieval, tools, delegation and memory.

Versioning with publish and rollback

Draft, publish with a changelog, and roll back to a previous published version; versions move through draft, published and historical states.

Templates

Curated app and agent blueprints across categories from customer support to finance, instantiated in one action into a real editable object.

Prompt templates

Versionable, variable-parameterised instructions with rendering, duplication and slug addressing, so wording is standardised across a workspace.

Prompt studio in progress

A dedicated authoring surface with live test runs against a real model; the studio page currently runs on sample data and is being wired to the live API.

Draft an agent from a sentence planned

Describe the assistant in natural language and have the platform draft the app, agents, prompts and suggested knowledge as drafts a human approves. Designed, not built.

Knowledge and retrieval

A managed pipeline from documents to answers, with retrieval settings treated as a reusable, shareable asset.

Knowledge bases

Documents are chunked, embedded and stored with per-base vector collections; the index lifecycle runs from empty through indexing to ready, with rebuild on demand.

Five chunking strategies

Fixed windows, recursive splitting, semantic, Markdown-aware and code-aware chunking, with configurable size and overlap.

Retrieval inside the agent turn

Ready knowledge bases linked to an agent are searched on every turn and the best passages are injected into the prompt with an instruction to admit ignorance rather than invent.

Retrieval profiles

Reusable recipes covering search mode, result count and thresholds, attachable to agents and distributable as packages.

Source citations in the answer in progress

Retrieved passages shape the answer today but are not yet returned to the caller as structured citations for display.

Connected knowledge sources planned

Automatic sync from websites and third-party systems rather than document upload.

Models and cost

Bring your own provider accounts. Botlit orchestrates them and records what each run consumed.

Eleven provider families

OpenAI, Anthropic, Google, Cohere, Mistral, Groq, Together, HuggingFace, Azure OpenAI, AWS Bedrock and custom endpoints, registered per workspace with a connection test.

Vaulted credentials

Keys are held in the platform's managed credential store and resolved at execution through a defined priority chain; they are never written into an agent definition.

Per-run cost and token ledger

Every run records tokens in and out, resolved cost, duration, model, tool calls and retrieved-passage count, aggregated per agent and per app.

Usage metering on priced units

Bots created, runs triggered, messages processed and knowledge bases provisioned are metered at the gateway against the plan.

Native adapters beyond the OpenAI wire format

in progress

Providers reachable over an OpenAI-compatible interface run the full runtime today; native protocol clients for the remaining families are the next tranche.

Cost dashboards and anomaly alerts

planned

A dedicated view slicing spend by app, agent and model, plus alerts when an agent deviates from its own baseline.

The agent runtime

What happens between a person's message and the answer. This is the newest part of the product and the part to describe most carefully.

Live chat execution

A send resolves the agent's model and credentials, calls the provider, writes an execution record and stores both sides of the exchange in the platform conversation store.

Multi-agent routing

Four dispatch strategies with keyword and pattern scoring first and a model classifier only as fallback, plus a deterministic fallback destination when scoring is inconclusive.

Tool calling inside the turn

in progress

A real client opens a session against registered tool servers, lists tools, invokes what the model asks for and loops within a bounded iteration budget. Available on the OpenAI-wire provider path; locally-executed tool commands are refused by design.

Agent-to-agent delegation

in progress

The target agent genuinely runs with its own model, knowledge and tools, bounded by a depth cap and outright cycle rejection. Same-workspace and OpenAI-wire path today.

Conversation history replay

in progress

Threads are persisted, and replaying prior turns into the model's context is the next step on this path.

Prompt-injection resistance

planned

Instructions embedded in retrieved documents, hostile tool output or a poisoned package are a known attack surface and the next tranche of guardrail work.

Deployment and conversations

Where the assistant meets people, and where the team goes to see what it said.

Channel connectors and routes

Seven connector types — web chat, Slack, Teams, Discord, WhatsApp, webhook and custom — with routes mapping a channel or thread to a specific bot app.

Instant acknowledgement, background answering

Inbound messages are acknowledged immediately and the agent runs in the background, so a channel never waits on model latency; replies are stored and delivered in parallel.

Shared team inbox

Every bot conversation lands in the platform inbox alongside human messaging, so support staff see agent conversations in the same place as their own.

Per-tenant channel app configuration in progress

End-to-end delivery on Slack or Teams depends on the customer's own platform app being registered for their tenant, which is a setup step rather than a switch.

Embeddable web chat widget planned

The web widget channel type is modelled but the embeddable script for external sites has not shipped.

Human handoff orchestration planned

Escalation, assignment, presence and response-time targets on top of the shared inbox.

Governance, security and audit

The part that decides whether a pilot becomes a rollout. Most of it is inherited from the platform, which is why it is strong for a product at this stage.

Permission rule on every operation

Every query and mutation across twenty-one API modules carries an access rule enforced at the gateway from the composed schema before a resolver runs; the service performs no authorisation of its own.

Workspace tenancy

Every object is workspace-scoped and cross-workspace reads and writes are rejected — a behaviour covered by an automated isolation test suite run against the live service.

Five seeded roles

Platform admin, bot builder, knowledge curator, bot operator and a read-only governance auditor, seeded into every workspace as least-privilege templates rather than flat grants.

Runtime policy enforcement in progress

Rate limit, budget, content and agent-trust policies evaluated on the way in and on the way out, fail-open where a workspace has configured nothing. Authoring an enforceable policy currently requires the API rather than the console.

Data-protection baseline

Encryption at rest, transport security, a published data-processing agreement, subprocessor list, acceptable-use and AI policies, and data subject access, erasure and portability handling.

SOC 2, ISO 27001 and HIPAA planned

Controls are being aligned against SOC 2 criteria, but no accredited auditor has issued a report, no ISO scope is defined, and no business associate agreement exists. None of this can be delivered by a software release.

Single sign-on and directory provisioning planned

Enterprise sign-on and automated user provisioning are planned as part of the enterprise tier.

Reuse, extension and ecosystem

How a good build becomes an asset, and how the product reaches beyond its own console.

Ten-kind package store

Apps, agents, tool packs, prompt templates, skills, retrieval profiles, policy bundles, channel templates, agent-to-agent definitions and knowledge-base templates, installable linked or forked, with fork-after-install.

Buckets

Hierarchical folders for organising apps and agents once a workspace grows past the first ten.

Embedded visual workflows

The FluidGrids workflow builder — the same engine used across the Burdenoff platform — runs inside Botlit, and a workflow can supply its own model credentials into an agent run.

Full GraphQL API

The entire product surface is available through the public gateway with actor identity and permission enforcement, so the console is a convenience rather than the only way in.

Store catalogue depth in progress

The published catalogue is thin at seed level and the community publishing and review workflow is not yet open, so early customers should expect to build their own packages.

Desktop, mobile, CLI, SDKs and editor extensions planned

Specified and documented, but nothing is released — no package is published and no store build has shipped.

10 Pricing & Plans

Quota-based rather than feature-based: every tier is intended to carry the full platform, with tiers differing on scale and on enterprise governance rather than on which features are unlocked. Model inference is separate and always the customer's own cost — Botlit does not resell tokens, so a prospect's total spend is the subscription plus whatever their provider charges. Bots, runs, messages and knowledge bases are metered against the plan; other dimensions in the published table are documented targets, not live caps. Checkout is not enabled yet, so every commercial conversation today runs through us directly.

PLAN	PRICE	WHO IT IS FOR	WHAT IS INCLUDED
Free	\$0 Free forever, for evaluation and small teams	Individuals and small teams proving the idea before anyone signs anything	Up to 10 bots; 1,000 runs per month; 2,000 messages per month; 2 knowledge bases; Bring your own model keys; Per-run token and cost records; Community support
Pro	\$49 / user / month Annual commitment available at \$468 per year, equivalent to \$39 per month	Production teams running assistants for real traffic	Up to 100 bots; 50,000 runs per month; 50,000 messages per month; 20 knowledge bases; Tool and agent-to-agent registries; Channel connectors and routes; Agent and app versioning with rollback; Priority support
Enterprise	Contact us Custom, negotiated on scale and governance requirements	Large organisations with governance, procurement and scale requirements	Unlimited scale, custom quotas; Everything in Pro; Multi-agent orchestration; Single sign-on and enterprise identity; Advanced roles and audit export; Custom tool servers and model endpoints; Custom service level and named success contact

WHAT TO TELL A PROSPECT ABOUT PRICE

- Model inference is never included. The customer pays their own provider directly, which is usually an advantage in the conversation because the cost is transparent and negotiable.
- Bots, runs, messages and knowledge bases are the dimensions actually metered today. Do not describe tool-server counts, delegation counts or interface call limits as hard caps.
- Prices are not yet published on the site and payment collection is not switched on. Quote these figures as the current plan structure and route any commitment through us.
- Enterprise features such as single sign-on and a named success contact are part of the enterprise packaging rather than something a Pro customer can buy as an extra.

11 Competitive Landscape

Seven real alternatives a prospect will name. Each row includes a genuine concession — a table where the competitor never wins is a table nobody believes, and partners who oversell get caught.

ALTERNATIVE	WHAT IT IS	WHERE THEY ARE STRONG	WHERE WE WIN
Microsoft Copilot Studio	Microsoft's agent builder for organisations already standardised on Microsoft 365 and Azure, licensed per tenant with a credit-based consumption model.	Deep Microsoft 365, Teams and SharePoint integration, roughly twenty channels, mature tool and delegation support that is generally available, and a full compliance portfolio including SOC 2, ISO and HIPAA. If the buyer lives in Microsoft and procurement prefers another Microsoft licence, they win on paperwork alone.	No Azure or Microsoft dependency, eleven provider families instead of an Azure-first default, per-run token and cost receipts instead of tenant-level credit dashboards, a forkable ten-kind package store rather than a curated catalogue, and permission enforcement on every API operation rather than a directory add-on.

ALTERNATIVE	WHAT IT IS	WHERE THEY ARE STRONG	WHERE WE WIN
Salesforce Agentforce	Salesforce's agent layer over CRM data, priced per conversation or through flexible credits and requiring an appropriate Salesforce edition underneath.	Unmatched depth against Salesforce records and cases, mature governance including data masking and attribute-based access control, a broad certification portfolio, and generally available tool and delegation runtimes. If every conversation ends in a Salesforce case, they are the right answer.	Agents run without any Salesforce contract, so the entry cost is a subscription rather than an edition upgrade plus per-conversation fees. Direct spend attribution per run replaces credit aggregates, and the assistant can span Slack, Teams, Discord, WhatsApp and custom systems rather than orbiting CRM.
Intercom Fin	A support-focused AI agent bolted onto the Intercom messaging suite, billed per resolution on top of seat costs.	The fastest path to a working support deflection bot for a team already on Intercom, outcome-based pricing that aligns cost with value, the broadest compliance portfolio of this group including a HIPAA business associate agreement, and around ten polished channels.	Botlit is not support-only — the same platform covers HR, IT, sales and operations. The customer chooses the model instead of accepting one proprietary engine, gets token-level receipts instead of outcome-level billing, and is not tied to Intercom seats. Multi-agent routing beats a single agent per workspace.
Botpress	A popular bot builder with a free tier, a self-hostable open-source core and bundled model spend.	Genuinely fast onboarding and time to a first working bot, broad channel coverage including voice on their enterprise tier, bundled model spend that removes token anxiety for smaller teams, and a self-hosted option for data-residency requirements Botlit cannot meet today.	Multi-tenant workspace-per-client isolation rather than a single-workspace model, permission enforcement on every operation, per-run audit with tokens and cost, four real routing strategies so one front door can front several specialists, and a forkable ten-kind package store rather than starter templates.
Voiceflow	A design-first conversational platform with the strongest visual canvas in this category and native voice telephony.	Best-in-class canvas and real-time design collaboration, strong environment promotion with staging and rollback, native voice telephony, and SOC 2 Type II, ISO 27001 and HIPAA compliance today. A team that judges on builder polish or needs voice will prefer them, and you should say so.	Architecture for an agent workforce rather than a single flow: specialist agents behind an app, an agent registry with trust levels, workspace-per-client isolation for agencies, a forkable package store, and per-run cost receipts instead of workspace-level credit aggregates.
OpenAI Responses API and Agents SDK	OpenAI's own toolkit for building agents in code against OpenAI models, priced purely on tokens with no platform seat fee.	The fastest developer path for a team already committed to OpenAI, production-ready built-in file search, category-leading reasoning models, and no platform subscription at all. For a code-first integration inside an existing product, that is hard to beat.	Botlit is the control plane the toolkit does not include: permission enforcement rather than key-scoped access, native channel connectors rather than an API-only surface, a package marketplace, a full execution audit rather than aggregate usage, and no single-provider dependency.
LangChain and LangGraph	The most widely used open-source framework for building agent runtimes, with a hosted observability and deployment layer alongside it.	The broadest model and tool ecosystem in open source, mature graph-based orchestration, deep tracing and evaluation tooling, and the freedom to self-host or embed with no platform dependency at all. A team that wants full control of the runtime should use it.	There is no runtime to host, monitor or scale, and no permission model, audit trail, channel layer or management console to build first. Governance, multi-tenancy, versioning and the package marketplace arrive with the subscription rather than as the next six months of engineering.

THE POSITIONING LINE TO HOLD

Botlit is the agent platform for teams whose AI strategy is bigger than one vendor's model and one vendor's channel — governed by default, and honest about what it costs.

12 Why Us / Differentiators

If a prospect remembers six things about Botlit, make it these. Each is structural — a consequence of how the product is built rather than a feature that could be copied in a sprint.

Governance is inherited, not bolted on

Botlit is a module of a multi-tenant workspace platform rather than a standalone tool that later added access control. Every operation on the API carries a permission rule enforced at the gateway, tenant separation is in the data model, and the audit trail, billing and file storage were already there. That is why a young product can survive a security review at all.

The bot app and the agent are different objects

Most platforms give you one bot with one enormous prompt. Botlit separates the user-facing front door from the specialists behind it, connected by routing links that carry priorities, keywords and fallbacks. A support bot can front billing, technical and escalation specialists without the user ever knowing, and without anyone writing orchestration code.

Eleven provider families, keys resolved at execution

Provider credentials live in the platform's credential store and are supplied only at run time, never written into an agent definition. Swapping a model binding is a configuration change, not a rewrite, and inference stays on the customer's own account — which also means the platform has no incentive to route work to an expensive model.

A receipt for every run, including the failures

The execution ledger is written by the runtime rather than reconstructed from logs, and it covers runs that failed, timed out or were cancelled. It carries tokens in and out, resolved cost, duration, model, tool calls and retrieved-passage count. Very few products in this category record the retrieved-passage count at all, and it is what makes silent retrieval failure detectable.

Ten kinds of package, linked or forked

Not a template gallery — an artefact registry that treats governance policies, retrieval settings and knowledge-base blueprints as distributable objects alongside whole apps and agents. Linked installs keep tracking the author's improvements; forked installs are owned outright; a linked install can be forked later without losing configuration.

Security decisions that are stated rather than hidden

Tool servers that would execute a locally-supplied command are refused outright, because a tenant-authored command in a shared platform is a remote-execution primitive. Outbound tool targets are checked against a strict endpoint guard. Delegation is depth-capped and cycle-rejecting. These are the sort of choices a reviewer asks about, and Botlit has written the reasoning down next to the code.

A visual workflow engine that is genuinely part of the product

The FluidGrids builder used across the Burdenoff platform runs inside Botlit rather than being embedded from a third party, and a workflow can supply its own configured model credentials into an agent run. Multi-step automation and conversational agents are two halves of the same product instead of two vendors.

13 Customer Proof & Credibility

There are no customers to point at, so credibility has to come from somewhere else: the depth of what has actually been built, the governance the product inherits, the evidence that it is exercised rather than merely written, and the fact that the team labels its own gaps in public.

Use the honesty as the proof — a prospect who catches an overclaim will discount everything else you said.

TWENTY-ONE GOVERNED API MODULES

The public interface spans bot apps, agents, routing links, chat, executions, knowledge bases, providers, tool connections, agent-to-agent registry and messaging, channels, policies, prompt templates, retrieval profiles, skills, the store, templates and buckets — every field carrying an enforced permission rule.

TESTED AGAINST THE LIVE SERVICE, NOT A SANDBOX

An automated end-to-end suite runs the core object lifecycles, tenant-isolation regressions, app and agent lifecycle, knowledge documents, retrieval profiles, channels, policies, store reads, execution metrics, prompt templates and skills against the live deployment, and the results are recorded including the failures.

CROSS-TENANT ISOLATION IS VERIFIED, NOT ASSERTED

A dedicated regression suite attempts cross-workspace reads and writes and confirms they are rejected while same-workspace operations continue to work. This is the single test a security reviewer cares most about and it is run rather than described.

ELEVEN MODEL PROVIDER FAMILIES

OpenAI, Anthropic, Google, Cohere, Mistral, Groq, Together, HuggingFace, Azure OpenAI, AWS Bedrock and custom OpenAI-compatible endpoints, each registered with the customer's own key and verifiable with a connection test before an agent depends on it.

DATA-PROTECTION BASELINE PUBLISHED

A data-processing agreement, subprocessor list, acceptable-use policy, AI policy, privacy and cookie policies are live on the site, with data subject access, erasure and portability handling implemented rather than promised.

LEAST-PRIVILEGE ROLES SEEDED AUTOMATICALLY

Five workspace roles — platform admin, bot builder, knowledge curator, bot operator and a read-only governance auditor — are seeded into every workspace as templates with narrowed permission sets, so nobody has to hand-build an access model.

PUBLIC ROADMAP LANGUAGE THAT UNDER-PROMISES

The product's own site labels the copilot and supervisor as roadmap, states that source citations are not yet returned, says the tool loop runs on one provider path, and explains where the visual builder is behind a competitor. That candour is unusual and it is worth pointing at directly.

Be straight about this. Botlit is pre-launch. There are no customers, no revenue, no reference accounts, no published usage figures and nothing to quote from a user. The site runs in waitlist mode and payment collection is not switched on. Several capabilities a competitor would list as shipped are partial here: tool calling and agent-to-agent delegation run only on the OpenAI-compatible provider path, retrieved passages are not yet returned as citations, some configuration is API-only rather than available in the console, and the copilot and supervisor are design work. On compliance, there is no SOC 2 report, no ISO 27001 certificate and no HIPAA business associate agreement — controls are being aligned but no accredited auditor has been engaged, and no software release can change that. What a prospect should weigh instead is the depth of what is built, the governance inherited from an established multi-tenant platform, the automated evidence that isolation and lifecycles work against the live service, and the fact that the team publishes its own gaps. If any of the missing items is a hard requirement, say so early and walk away — that judgement is worth more to us than a pursued deal that dies in security review.

14 Demo Information

There is no self-serve trial yet, so a demo is a live conversation rather than a link. That is an advantage: you get a qualified prospect in front of the person who built the product, and we do the technical explaining.

HOW TO GET A DEMO

Book a session on the 1-1 calendar link or email Vignesh directly. He runs the walkthrough against a real workspace, and early-access onboarding is arranged from the same conversation. Prospects can also join the waitlist themselves at www.botlit.ai/waitlist, but a partner introduction gets a faster and better-prepared session.

Runs about Thirty minutes for the standard walkthrough; forty-five to sixty if a security or compliance reviewer is joining, because those questions deserve unhurried answers.

WHERE TO LOOK

- **Book a 1-1 walkthrough** — <https://calendly.com/vignesh-burdenoff/1-1>
- **Product site** — <https://www.botlit.ai>
- **Use-case stories** — <https://www.botlit.ai/use-cases>
- **Documentation** — <https://docs.botlit.ai>
- **Application** — <https://app.botlit.ai>

The walkthrough, in order

- 1 A bot app created from a template, with its specialist agents behind it**
Start here, because it reframes the whole thing. The app is the front door and the agents are specialists — that separation is what lets one support bot handle billing and technical questions properly instead of one prompt trying to do both.
- 2 The agent's instructions, its model binding and its version history**
This is a versioned object, not a text box three people edit. Publish with a changelog, roll back when a change makes answers worse. Ask the prospect what happens today when a prompt change goes wrong.
- 3 A knowledge base with its documents, chunking settings and semantic search**
The answers already exist in their documents. Show retrieval being validated before anything goes live — proving retrieval quality up front is what separates this from pasting text into a prompt and hoping.
- 4 A chat with the agent, answering from that knowledge, and refusing when it should**
The moment worth waiting for is the refusal. Ask something the documents do not cover and let the assistant say it does not know. Every support leader in the room has been burned by a bot that would rather invent.
- 5 The execution record for that exact message**
The same conversation, with tokens in and out, cost, duration, model and how many passages it retrieved. This is the slide that wins the finance conversation, and the retrieved-passages count is how you catch retrieval failing silently.
- 6 The provider registry with several model families configured**
Their keys, their contracts, their choice — and a connection test before an agent depends on it. This answers the lock-in objection before it is raised, and it is the one thing almost no competitor in this bracket can match.
- 7 A guardrail policy blocking a message, and the refusal a user actually sees**
A written policy that runs in the request path. Point out that the wording is identical in the app and on Slack, so a blocked interaction never turns into two different failure modes.
- 8 A channel connector and route sending the assistant to Slack**
Where the assistant meets people. Inbound messages are acknowledged instantly and the agent answers in the background, so the channel never times out waiting for a model.

15 Sales Talking Points

Eight things to say, and why each one lands. Vary which you lead with — the finance argument and the security argument reach different rooms.

Your keys, your model contracts, your choice of where inference runs.

Lock-in is the fear underneath most of the objections in this category, and it is the one place Botlit is straightforwardly stronger than the incumbents. It also reframes price: the subscription is the platform, the model spend stays theirs.

Every run leaves a receipt — tokens, cost, duration, model, tool calls and retrieved passages.

It converts an AI conversation into a business conversation. Finance can attribute spend, operations can spot the agent whose cost per run tripled, and an agency can price a client engagement on real numbers.

The assistant is told to say it does not know rather than invent.

Everyone in the room has been embarrassed by a confidently wrong bot. Naming the failure mode first, then showing the refusal in a demo, does more for credibility than any accuracy claim.

Every operation carries a permission rule enforced before the request reaches the code.

Security reviews are where AI pilots die. This is a structural answer rather than a policy document, and it is the reason a product at this stage can be considered by an enterprise at all.

One front door, several specialists behind it.

It explains why Botlit is not just another chatbot builder, and it maps directly onto how support and IT teams are already organised. Most prospects immediately name the three specialists they would want.

A good build becomes a package you install again — linked or forked.

For agencies and platform teams this is the commercial hook. It turns delivery work into an asset and makes the second engagement cheaper than the first, which is exactly what their margin problem needs.

Guardrails are evaluated on the way in and on the way out, cheapest check first.

Compliance officers want a control, not an intention. The ordering detail — rate limit before budget before content, so a runaway integration is stopped before it spends anything — signals that this was designed rather than described.

Pre-launch, and we would rather tell you the gaps than discover them together in a security review.

Saying it before they find it is the single most persuasive move available with an early product. It also efficiently disqualifies buyers who need certification today, which saves everyone a quarter.

The eleven objections that actually come up, in the words prospects use. Answer the hard ones plainly — a partner who hedges on compliance loses the room permanently.

"You do not have SOC 2. We cannot buy software without it."

That is correct and I am not going to talk around it. There is no SOC 2 report, no ISO certificate and no HIPAA agreement, and none of those can be produced by a software release — they require an accredited auditor over months. Controls are being aligned against SOC 2 criteria and the underlying evidence exists, but if certification is a gate on your signature then today is not the day. Tell me what your window looks like and I will introduce you when there is something real to show.

PROOF · The certification roadmap is published, states the current status of each certificate plainly, and the product's own site says in progress rather than certified.

"You are a small company. What happens to us if you disappear?"

A fair question for any early vendor. Two things reduce the exposure. The model inference runs on your own provider account with your own keys, so the expensive dependency was never ours. And your knowledge is your own documents, your configuration is reachable through a documented interface, and conversations are stored in a system you can read out. You would lose the management layer, which is real, but not your data or your model relationship.

PROOF · Bring-your-own-provider architecture with credentials the customer holds and revokes, plus a full API surface over every object in the product.

"We could build this ourselves. We already have engineers using these models."

You could build the agent loop in a fortnight — that part is genuinely easy now. What takes quarters is everything around it: an access model applied to every operation, tenant isolation, an audit trail, versioning with rollback, channel connectors that acknowledge instantly and answer in the background, a cost ledger, guardrails in the request path, and a console non-engineers can operate. That is the part your team will still be maintaining in two years.

PROOF · Twenty-one API modules covering the management plane, with permission enforcement, versioning, the execution ledger, seven channel types and the package store already built.

"Why not just use Microsoft Copilot Studio? We already pay Microsoft."

If you are entirely inside Microsoft 365 and Azure, and procurement prefers another Microsoft licence, they may well be the right answer and I will say so. Where it changes is model strategy and cost visibility: Copilot Studio is Azure-first and bills in credits at tenant level, so you cannot attribute spend to an agent or negotiate model pricing separately. If your AI plans are bigger than Azure, that ceiling arrives sooner than people expect.

PROOF · Eleven provider families with the customer's own keys, and per-run token and cost records rather than tenant-level credit dashboards.

"Our last chatbot made things up and we had to switch it off."

That is the failure mode this product is built around. On every turn the agent's own knowledge bases are searched, the best passages are put in front of the model, and it is instructed in the same breath to say it does not know rather than invent. Ask for the demo and deliberately ask something the documents do not cover — watching it decline is more convincing than any accuracy figure I could quote.

PROOF · Retrieval injected inside the agent turn, only from bases whose index has finished building, with graceful degradation and a logged reason when a base cannot be searched.

"How do we know what this will cost us? AI bills are unpredictable."

Two costs, kept separate on purpose. The subscription is a published plan. The inference is billed by your own provider on your own account, so you keep whatever rate you negotiated and we have no incentive to route you to an expensive model. On top of that, every run records its tokens, its cost and its model, so an agent whose cost per run doubled after a change is visible in a day rather than at the end of the month.

PROOF · Per-run cost and token ledger aggregated per agent and per app, and a plan structure of free, \$49 per user per month, and custom for enterprise.

"Where does our data go, and will it be used to train models?"

Your knowledge bases and conversations serve your agents; they are not used to train foundation models. Inference goes to the provider you configure, which means you choose the jurisdiction and the terms — including pointing the platform at a model server you host yourself over a compatible interface. Data is encrypted at rest and in transit, and credentials sit in a managed store and are supplied only at run time, never written into an agent.

PROOF · Published data-processing agreement, subprocessor list and AI policy, plus a custom endpoint type for self-hosted inference.

"Half of this looks like roadmap. What is actually working today?"

Working today: building and versioning agents, knowledge bases with retrieval inside the live turn, eleven provider families with vaulted keys, routing across specialists, the execution ledger, channel connectors, the package store, buckets, workflows and the permission model. Partial today: tool calling and delegation run on the compatible-provider path only, citations are not yet returned to the caller, and some configuration is API-only. Roadmap: the copilot and supervisor. I would rather you had that list from me than assembled it yourself.

PROOF · The product's own site labels roadmap items as roadmap, including on the pages that would be most tempting to overclaim.

"We need this to answer phone calls too."

It will not. There is no voice channel — the connectors are web chat, Slack, Teams, Discord, WhatsApp, webhooks and custom integrations. If voice telephony is central to what you are solving, Voiceflow or Botpress Enterprise are the honest recommendations and I would rather say that now. If voice is a later phase and text is the current problem, this is still worth a conversation.

PROOF · Seven channel connector types, none of which is telephony.

"What support do we actually get? We cannot be the ones debugging your product."

At this stage support is unusually direct: you deal with the people who build it, not a first-line queue. The free tier is community support, Pro is priority support, and enterprise arrangements include a named contact and an agreed service level. Be aware that the flip side of a small team is a small team — if you need round-the-clock coverage across time zones today, say so and we will be straight with you about it.

PROOF · Support tiers are part of the published plan structure, and the product owner takes demo and onboarding conversations directly.

"We are a regulated business. Can we run this on our own infrastructure?"

The platform itself is cloud multi-tenant only — there is no self-hosted or air-gapped install, and that is not a near-term option. What can stay on your infrastructure is the part most people are worried about: inference runs against whatever model endpoint you configure, including a model server you host yourself. If the requirement is that no data leaves your network at all, this is not the right product today.

PROOF · Custom OpenAI-compatible provider type for a self-hosted model server; workspace-scoped tenancy with an automated cross-tenant isolation test suite.

17 Discovery / Qualification Questions

Eleven questions that separate a real Botlit opportunity from a polite one. Listen for volume, for an existing model relationship, and for whoever will have to approve it.

ASK THIS

LISTEN FOR

How many enquiries does your team handle in a month, and roughly what share are the same handful of questions?

A four-figure monthly volume with a large repeat share is the strongest single signal. Low volume or genuinely unique questions means the return will not be there and you should say so.

Have you tried a chatbot or AI assistant before? What happened to it?

A switched-off bot is a gift — it means budget existed and the failure is specific. Listen for whether it was accuracy, cost, integration or governance that killed it, because that tells you which half of the pitch to lead with.

Which model provider do you already have an account or contract with?

An existing relationship makes the bring-your-own-model argument concrete and lowers perceived cost. Nobody at all means you are also selling the concept of model procurement, which lengthens everything.

Where do the answers actually live today — a handbook, a help centre, a wiki, someone's head?

Documents that exist and are reasonably current mean a knowledge base can be built in days. Answers that live only in people's heads mean a content project has to happen first, and that changes the timeline honestly.

Who has to approve this before it goes live in front of customers?

If security or compliance sign-off is named, surface the certification gap immediately. If nobody is named, ask again in a different way — an unnamed approver is the most common reason a promising deal stops moving.

ASK THIS

LISTEN FOR

Do you know what your current AI or automation tooling costs per month, and can you break that down by team?

An answer of no, or a number they clearly do not trust, is the opening for the cost ledger. A precise breakdown means they already care about this and the argument will land quickly.

Where would this assistant have to live for people to actually use it?

Slack, Teams or WhatsApp are all supported and are a good sign. A phone line, an in-product widget embedded on their public site, or a channel not in the list are all things to flag before enthusiasm outruns capability.

Do you need one assistant, or several that would each know a different part of the business?

Several is the better fit and lets you introduce the app-and-agents split. One simple assistant is still viable but competes with cheaper tools, so check what else they are evaluating.

If a prompt change made answers worse, how would you find out and how would you undo it?

Hesitation here is the versioning conversation writing itself. Teams that already treat prompts casually are the ones that most need publish, changelog and rollback.

Is there a data-residency, self-hosting or certification requirement I should know about now?

Ask this deliberately early. A hard requirement for SOC 2, HIPAA, an air-gapped install or in-region data is a disqualifier today, and finding it in week one is a good outcome, not a failed conversation.

If you serve several clients or business units, do they need to be kept fully separate?

A yes points at workspace-per-client isolation and often at the package store as well. It usually also means the deal is larger than the first conversation suggested.

18 How to Identify a Good Opportunity

Use these before you spend time. Green signals compound — two or three together is worth an introduction the same week. Red signals are not soft; they mean walk away and keep the relationship.

STRONG SIGNAL — PURSUE

- High volume of repetitive enquiries, in support, HR, IT or a service desk, with headcount that will not grow to match
- A chatbot they already switched off, or one they are visibly unhappy with
- An existing model provider contract, and someone who worries aloud about being locked to a single vendor
- Finance or leadership asking what the AI actually costs and getting no clear answer
- An agency or platform team rebuilding a similar assistant for each client or business unit
- A documented body of knowledge — handbook, help centre, policy library — that people currently search by hand

Introduce us straight away.

WORTH A CONVERSATION

- Interested but their approver has not been named yet — keep going, but get the name before promising a timeline
- They want an embeddable widget on their own public website, which is modelled but not yet shipped
- Tool calling or agent-to-agent delegation matters to them and their preferred provider is outside the compatible-endpoint path today
- Compliance certification is wanted but not yet a contractual gate — worth an early conversation about their timeline
- They have a strong preference for a polished drag-and-drop builder and will judge on that first
- The knowledge exists but is scattered and out of date, so a content project precedes any deployment

Qualify further before an introduction.

POOR FIT — DO NOT CHASE

- A SOC 2 report, ISO 27001 certificate or signed HIPAA agreement is required before signature
- Voice telephony is a core requirement rather than a later phase
- Self-hosted, on-premise or air-gapped deployment is mandatory
- Strict in-region data residency is contractual — there is no residency guarantee today
- They are fully committed to Microsoft 365 and Azure and procurement wants another Microsoft licence
- Low enquiry volume with no repeat pattern and no document corpus to draw on

Politely park it; suggest another product if one fits.

THE ONE-QUESTION TEST

"How many enquiries do you handle in a month, and what share of them are the same questions asked in different words?" A four-figure volume with a large repeat share means there is real money in deflection and the rest of the conversation is worth having. Anything less and the return will not justify the change, however much they like the product.

19 How to Introduce Us

Keep the introduction short and about their problem, not our architecture. You are handing over a conversation, not closing a deal — and being straight about the pre-launch stage from the first sentence is what makes the introduction land rather than sour later.

IN CONVERSATION

""You know that thing you said about the same five questions coming in all day? There's a product called Botlit I think you should look at. It's an AI assistant platform, but the part that makes it different is that it answers from your own documents and records exactly what every answer cost — so it's not a black box when finance asks. It's pre-launch, so this is an early-access conversation rather than a purchase. Want me to introduce you to the founder?""

EMAIL — SUBJECT: AI AGENTS YOUR SECURITY TEAM WOULD ACTUALLY APPROVE

Hi [Name],

You mentioned [team] is fielding the same questions over and over, and that the last assistant you tried never made it past security review.

There's a product I think is worth twenty minutes: Botlit. You build an assistant, ground it in your own documents, run it on whichever model provider you already pay for, and put it on Slack or Teams. Every answer is recorded with the tokens and cost behind it, and every operation runs under your existing roles and permissions.

It's pre-launch, so this is an early-access conversation rather than a purchase order — and there's no SOC 2 report yet, which I'd rather you heard from me.

Happy to introduce you to Vignesh, who builds it: vignesh@burdenoff.com.

[Your name]

SHORT MESSAGE / LINKEDIN

Saw your post about support volume outgrowing the team. Worth a look at Botlit — AI assistants that answer from your own documents, run on your own model keys, and record the cost of every answer. Pre-launch, so early access rather than a product you can buy today. Happy to introduce you to the founder if it's relevant.

WARM THREE-WAY INTRODUCTION

When you make the introduction, copy vignesh@burdenoff.com and include three things: what the prospect's repetitive-enquiry problem is in their own words, who has to approve an AI deployment there, and whether any certification or residency requirement is already on the table. That last one saves everybody a month. Add the 1-1 link (<https://calendly.com/vignesh-burdenoff/1-1>) so they can book directly, and mention the phone number +91-7358445777 if the conversation is time-sensitive.

Please do not: Do not describe Botlit as generally available, or imply anyone is using it in production today. It is pre-launch with no reference customers, and a prospect who finds that out later will discount everything else you said.; Do not quote a compliance certification. There is no SOC 2 report, no ISO certificate and no HIPAA agreement, and no phrasing makes that acceptable to imply.; Do not promise a feature from the roadmap section — particularly the copilot, the supervisor, source citations, the embeddable widget or voice.; Do not negotiate price, discount or contract terms. Quote the published plan structure if asked and route the commercial conversation to us.; Do not forward internal comparison notes or a prospect's requirements to a third party. Send them to us instead.

20 Lead Handoff Process

The introduction is the deliverable. Do not try to run discovery on our behalf — a short, well-informed handover gets a better first call than a long one, and protects your commission registration.

1. Register the opportunity before the introduction

Email vignesh@burdenoff.com with the company name, the contact and the reason you think it fits. Registration is what establishes your claim to the commission, and it has to happen before the prospect contacts us independently.

2. Send the qualification summary

One short paragraph covering the repetitive-enquiry problem in their words, their approximate volume, who approves an AI deployment, and any certification or residency requirement already mentioned. This is what makes the first call useful rather than exploratory.

3. Make the warm introduction

A single email introducing the prospect to Vignesh, with the 1-1 link so they can book directly. Keep yourself on the thread; you are not expected to run the meeting, but staying visible keeps you part of the relationship.

4. Let us take the technical and commercial conversation

We handle the demo, the security questions, the pricing and the paperwork. If the prospect asks you something you are unsure about — particularly on compliance or on what is shipped — say you will find out rather than guessing.

5. Agree who owns follow-up

For a referral introduction we take it from the first call. For a consulting or implementation engagement, agree up front which of us drives the evaluation, because that decision changes the commission tier that applies.

6. Stay on the thread through the decision

We report back after the first call and again at the decision. If the prospect goes quiet, tell us what you know — a partner's read on why a deal stalled is usually better than ours.

SEND US THIS — IT IS ALL WE NEED

- Company name, approximate headcount and what they do
- The named contact, their role, and how you know them
- The problem in the prospect's own words — which team, which enquiries, roughly what volume
- Who has to approve an AI deployment there, by name or by role
- Any certification, data-residency or self-hosting requirement already mentioned
- Which tools they use today for support, messaging or automation
- Whether they already pay a model provider, and which one
- Your read on timing — active project, budgeted, or exploratory

WHAT HAPPENS NEXT, AND WHEN

We acknowledge a registered introduction within one working day and aim to hold the first call within five working days. If a prospect is time-sensitive, say so in the registration email or call +91-7358445777 and we will move faster. You get a written update after the first call and again at the decision point, whichever way it goes.

Deal registration: Register by emailing vignesh@burdenoff.com with the company name, the contact and a one-line reason before any introduction is made. Registration is acknowledged in writing and holds for ninety days, renewable while the opportunity is genuinely active. An account that was already in an active conversation with us is not eligible, and we will tell you that immediately rather than later.

21 Partner / Referral Commercials

Botlit follows the standard Burdenoff partner programme, which is consistent across every product in the portfolio. Commission is paid on collected net subscription revenue — not on inference costs, which the customer pays their own model provider directly, and not on taxes or third-party pass-throughs. The tier that applies depends on how much of the engagement you carry, from a registered introduction through to full implementation delivery or resale.

PROGRAMME	YOUR ROLE	FIRST-YEAR	RENEWAL / ADD-ONS
Referral Partner	Registered introduction, no sales ownership	10% of first-year collected net subscription	0% unless separately approved · 3%, capped at \$10,000
Consulting Partner	Introduction plus evaluation support	12% of first-year collected net subscription	0–3% if active in the success plan · 3%, capped at \$15,000
Implementation Partner	Introduction plus implementation delivery	15% of first-year collected net subscription	3–5% if active in the renewal success plan · 5%, capped at \$25,000
Authorised Reseller	Approved resale, first-line commercial ownership	Up to 15% reseller margin on the approved quote	Up to 8% if the account stays active and supported · Up to 5% on dedicated tenant or managed services, if approved
MSP / Strategic Partner	Bundled or managed offering	Custom	Custom · Custom

THE RULES THAT MATTER

- Register the opportunity in writing before any introduction. Registration establishes the claim; an unregistered introduction to an account we are already talking to does not earn commission.
- Registration holds for ninety days and is renewable while the opportunity is genuinely active. We will tell you within a working day if an account is already in play.
- Commission is calculated on collected net subscription revenue only — excluding taxes, third-party pass-throughs, and the customer's own model inference spend, which never flows through us.
- The tier is agreed before the evaluation starts, not after the deal closes. If you take on implementation delivery partway through, we adjust in writing at that point.
- Discounting is ours to approve. A partner-quoted discount that we have not agreed is not binding on the customer or on the commission calculation.
- Refunds, chargebacks and unpaid invoices are netted against commission, because commission follows collected revenue rather than booked revenue.
- Partner-branded resale under the Authorised Reseller tier requires written approval and a signed agreement before any customer-facing quote is issued.

GETTING PAID

Commission is calculated on revenue actually collected and paid quarterly in arrears, within thirty days of the quarter end, against a partner invoice. Payment is in the currency of the customer's contract. Because the product is pre-launch and checkout is not yet enabled, early partners should expect the first commission cycles to run alongside the first paid contracts rather than on a fixed calendar.

These are the standard Burdenoff partner programme terms, published here for guidance. They are not an offer and they are not binding on their own. The terms that actually apply are those in your signed partner agreement and in the quote issued to the customer, and where the two differ the signed documents govern. Rates, caps and eligibility can change for opportunities not yet registered.

22 Contact Information

YOUR CONTACT

Vignesh T.V.
vignesh@burdenoff.com
+91-7358445777

BOOK A 1-1 CALL

<https://calendly.com/vignesh-burdenoff/1-1>

Use this for anything — qualification, pricing, a joint call with your prospect, or a live demo.

EVERY BURDENOFF PRODUCT

<https://burdenoff.com/products>

If Botlit is not the right fit, the answer may well be on that page.

WHICH CHANNEL FOR WHAT

- **A live opportunity** — email with the handoff details in section 20, and copy the phone number in if it is time-sensitive.
- **A question you cannot answer** — book the 1-1 call; it is usually faster than email.
- **A prospect who wants to meet us** — send them the calendly link directly, or ask for a joint call.

23 Useful Links & Sales Assets

ASSET	WHERE	WHAT IT IS FOR
Botlit website	https://www.botlit.ai	The public product site — positioning, features, use cases and comparisons. Safe to share with anyone.
Botlit application	https://app.botlit.ai	Where an early-access customer signs in and works. Not a self-serve signup today.
Documentation	https://docs.botlit.ai	Getting started, core concepts, agents, interface reference, tool integration and security sections. Useful for a technical prospect.
Use-case stories	https://www.botlit.ai/use-cases	Twenty-five written scenarios with roadmap items labelled. The best link to send after a first conversation.
Comparisons	https://www.botlit.ai/vs	Head-to-head pages against the named alternatives, including where each competitor genuinely wins.
Security and compliance	https://www.botlit.ai/security	The current posture stated plainly, with in-progress items marked as in progress. Send this before a security conversation rather than after.
Frequently asked questions	https://www.botlit.ai/faq	Covers models, data handling, availability and pricing questions in the product's own words.
Join the waitlist	https://www.botlit.ai/waitlist	For a prospect who wants to register interest themselves. A partner introduction still gets a better first session.
Book a 1-1 call	https://calendly.com/vignesh-burdenoff/1-1	Direct calendar for Vignesh T.V. — include this in every introduction so the prospect can book without another round trip.
All Burdenoff products	https://burdenoff.com/products	The wider portfolio. Useful when a prospect's problem turns out to be a different product's problem.

WHAT YOU CAN SHARE, AND WITH WHOM

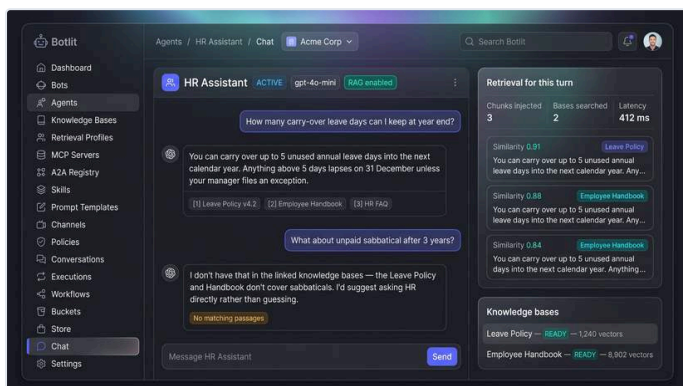
- Share freely in public: anything on www.botlit.ai, docs.botlit.ai and the use-case, comparison, security and FAQ pages. It is written to be public and it is kept honest about maturity.
- Share with a named prospect only: this playbook, the cheat sheet, the plan structure and pricing figures, and the competitor concessions. They are written for someone selling on our behalf and read badly out of context.
- Needs our written sign-off before use: any co-branded material, any press or blog mention, any customer-facing quote, and any use of the Botlit name or logo in your own marketing.
- Never share or imply: customer names, usage figures, testimonials or references — none exist. Never state or imply a compliance certification, and never forward a prospect's requirements or internal notes to a third party.
- Always current: check the site before quoting a capability. Where this document and the live site disagree, the site is the public claim and it wins.

24 Product Mockups — Complete Visual Appendix


7 SCREENS

Every screen in the Botlit design set — 7 in total, none omitted. Use these to show a prospect what the product looks like before a live demo.


PRODUCT SCREENS • 7 SCREENS




001 Botmanor AI 01 Grounded Answer Chat




002 Botmanor AI 02 MCP Tool Loop




003 Botmanor AI 03 Routing Strategy




004 Botmanor AI 04 A2a Delegation Trace



005 Botmanor AI 05 Policy Enforcement



006 Botmanor AI 06 Execution Ledger



007 Botmanor AI 07 Copilot Draft And Ops Feed



Found someone who needs Botlit?

You do not need to be technical to spot a Botlit opportunity: someone answering the same questions all day, with no idea what their AI is costing, is the whole qualification. Make the introduction early and let us carry the technical conversation. The honest parts of this document are the ones that will earn you trust.

EMAIL OR CALL

Vignesh T.V.

vignesh@burdenoff.com

+91-7358445777

BOOK A 1-1 WITH VIGNESH

<https://calendly.com/vignesh-burdenoff/1-1>

ALL BURDENOFF PRODUCTS

<https://burdenoff.com/products>



Burdenoff Consultancy Services Pvt. Ltd. · Botlit is available at www.botlit.ai.

This playbook is confidential and intended for Burdenoff partners. Figures, plans and capabilities are current at the date of issue and may change; final commercial terms are those set out in your partner agreement and in the quote issued to the customer.