
Controlled, Auditable AI Knowledge for Operators, End Users, and Partners

The governed control layer between your enterprise data and the large language models you deploy.

01

EXECUTIVE
OVERVIEW

Making AI Reliable When Real Data and Real Risk are Involved

Large language models like ChatGPT, Claude, Gemini, Grok, and Llama are powerful. The business challenge is not whether these models can generate useful output. It is making AI reliable, governed, and safe when real data, real customers, and real risk are involved.

Sprinklenet Knowledge Spaces is the control layer between your knowledge and the AI experiences you deploy. It lets you bring your own data, apply governance and guardrails, and deploy assistants and AI-powered experiences that users can trust. It is designed for business leaders, technical teams, and everyday end users who need fast, accurate answers grounded in approved information.

Model choice becomes a procurement decision, not a rebuild, when your knowledge lives in a control layer you own.

The rest of this paper describes what Knowledge Spaces solves, how the control layer is structured, how a deployment reaches production in about six weeks, and how partners build their own products on top of it.

02

WHAT IT SOLVES

The Constraints Every Organization Hits

Organizations run into the same six constraints when they try to operationalize AI. Knowledge Spaces is built to resolve each one.



Trust

Answers grounded in approved sources, not guesswork.



Control

Define what the AI may use and how it may behave.



Traceability

Track what changed, who changed it, and when.



Data Ownership

Your data stays unambiguously your data.



Sharing Without Leakage

Collaborate without exposing source materials.



Integration

Connect AI to real systems and pipelines.

03


THE CONTROL LAYER

The models stay interchangeable underneath. Your application stays yours on top.

Where Knowledge Spaces Sits

You pipe your data in once. Inside Knowledge Spaces you set the rules: which models see what, who can ask what, what gets logged, and how every organization and user is managed. The intelligence you build sits on top; the sources sit underneath; the governance is in the middle.

FIG. 1 THE CONTROL LAYER




04

STRUCTURE


Organizations contain Knowledge Spaces; Spaces are used by Bots under explicit rules.

Structure, and How It Works

FIG. 2 CONTAINMENT MODEL



How It Works, in Five Steps



Enterprise Integration

Connect an existing application API to a Knowledge Space so operational data becomes part of what the bot can use, for example a CRM or order-management system. If a needed API does not exist, Sprinklenet can help design and build it so critical systems safely participate.

05

BUILD ON TOP

Built for Partners and Product Builders

Knowledge Spaces is not just an internal tool. It is a foundation for partners and clients who want to build their own solutions for end users, customers, and markets.

- Each customer gets their own Organization and private Knowledge Spaces, in a repeatable model.
- The builder maintains a Core Space of trusted frameworks, playbooks, and reference knowledge.
- Combine the builder's Core Space with each customer's private Space, without mixing customer data.
- Offer controlled knowledge sharing so customers opt into collaboration without exposing materials.
- Deploy branded chatbots, embedded assistants, APIs, and structured-output workflows that can be monetized.

Knowledge Spaces turns AI into a governed capability you can productize and deploy at scale, delivering trusted experiences to your entire workforce or millions of customers while maintaining clear data boundaries.

06

IN PRACTICE

Use Case Examples

Private Equity Portfolio Intelligence

01

SCENARIO

A firm wants portfolio-wide insight while maintaining strict boundaries between companies, teams, and external audiences.

HOW IT WORKS

A Space per portfolio company holds its SOPs; analyst bots combine firm-wide and approved company Spaces; company bots stay restricted.

OUTCOME

Faster insight across the portfolio, with less risk of data cross-contamination.

Real Estate Brokerage Deployment

02

SCENARIO

A brokerage wants every agent to have a compliant client-facing assistant while the broker keeps guardrails and brand standards.

HOW IT WORKS

The brokerage maintains compliance Spaces; agents use approved Spaces without exposing source docs; MLS and market feeds connect for live context.

OUTCOME

A consistent, compliant client experience, scalable to thousands of assistants.

06

IN PRACTICE

Continued

Government Contracting Compliance

03

SCENARIO

Contractors need a trusted AI experience for FAR, CAS, and internal-policy alignment, without hallucinations or uncontrolled outputs.

HOW IT WORKS

Specialized compliance bots ground answers in the regulations; teams pressure-test decisions; traceability supports audit readiness.

OUTCOME

Faster answers, fewer errors, and better defensibility and governance.

Enterprise Marketing Planning

04

SCENARIO

A brand wants better campaign planning by combining internal performance data with curated partner insights, keeping both parties controlled.

HOW IT WORKS

The brand builds internal Spaces; the partner shares selected Spaces without raw file handoff; bots combine them under framing rules.

OUTCOME

Faster planning cycles and better decisions from combined intelligence.

07


TIME TO VALUE

A typical deployment reaches production in about six weeks.

Six Weeks to Production

Knowledge Spaces is designed for autonomy. Sprinklenet offers full-service onboarding, and the platform is intuitive enough for your team to manage directly after brief training. Whether you drive the process or we do, a typical deployment reaches production in about six weeks.

FIG. 3 DEPLOYMENT TIMELINE



Deploy via Embed or API

Knowledge Spaces supports embedded experiences and API-based delivery, so you can wrap your own branded UI around the intelligence, across websites, applications, and other digital interfaces. Developer documentation for the Knowledge Spaces API is published at sprinklenet.com/developers.

Frequently Asked Questions

How is Knowledge Spaces different from using ChatGPT or Gemini directly?

It adds governance, traceability, and data control. You get a high-quality experience grounded in approved knowledge with explicit constraints, rather than a general tool that may respond beyond what your organization can trust.

Are we locked into a specific LLM vendor?

No. Knowledge Spaces is LLM-agnostic. It makes switching models easy as the market evolves, and lets you pick the best model per use case on speed, accuracy, and cost.

Do we hand our data over to LLM vendors?

No. Only the relevant context is retrieved from your approved sources and sent to the model to generate a specific response. For regulated and enterprise workloads we route to models under no-retention, no-training terms, and confirm the exact model terms as part of your deployment.

How do LLM keys and billing work?

Bring your own API keys for preferred providers, which typically simplifies billing and cost governance. If you do not bring keys, Sprinklenet can provide model access and pass through usage costs, with professional-services support for setup.

Can we control what sources a bot is allowed to use?

Yes, granularly. Restrict a bot to specific documents, live data via API pipelines, approved URLs, or the open web, and define exactly which combination each bot can access.

Can we share knowledge with partners without exposing documents?

Yes. Another Organization can use a shared Space in a bot without receiving the underlying source materials, and sharing can be revoked at any time.

GET STARTED

Ready to Control Your AI Knowledge?

Whether you need an internal compliance bot, a partner collaboration space, or a customer-facing assistant, Knowledge Spaces provides the governance layer you need.



Knowledge Spaces

LEARN

sprinklenet.com/knowledge-spaces

BUILD

sprinklenet.com/developers

CONTACT

sprinklenet.com/contact