DESIGNING A SECURE AI KNOWLEDGE HUB IN MICROSOFT 365

A practical reference architecture for SharePoint, Teams, Copilot Studio, permissions, governance and testing

PIXELDUST

Secure AI Knowledge Hubs for organizations using Microsoft 365

Executive Summary

A secure AI knowledge hub is not simply a chatbot connected to a folder of documents. It is an information system with defined sources, permissions, ownership, testing standards and maintenance responsibilities.

Microsoft 365 provides most of the foundational controls required to build this system: SharePoint for governed content, Microsoft Entra ID for identity, Teams for user access, Copilot Studio for the conversational layer and Microsoft Purview for compliance capabilities. The quality of the result depends less on the novelty of the AI model than on the architecture around it.

This white paper presents a practical reference design for small and mid-sized organizations. It explains which components belong in the architecture, how information should flow, how permissions should be inherited, where risk enters the system and what should be tested before launch.

Five design principles

- Use approved content, not every available file.
- Enforce source permissions before the AI layer.
- Separate content ownership from technical administration.
- Test realistic questions, failures and access boundaries.
- Treat launch as the start of governance, not the end of implementation.

1. Reference Architecture

The architecture should be divided into layers. Each layer has a distinct purpose and should be governed independently.

Layer	Primary components	Purpose
Experience layer	Microsoft Teams, web chat or another approved interface	Where employees ask questions and receive answers.
Conversation layer	Copilot Studio agent	Interprets requests, applies topics or instructions and returns grounded responses.
Knowledge layer	SharePoint sites, libraries, lists and approved data sources	Stores the controlled information used to answer questions.
Identity and access layer	Microsoft Entra ID, Microsoft 365 groups and SharePoint permissions	Determines which users can access which sources.
Governance and compliance layer	Purview, retention, sensitivity labels, audit logs and review	Controls lifecycle, classification,

	processes	monitoring and accountability.
Operations layer	Testing, analytics, support, content review and change management	Keeps the system useful and trustworthy after launch.

Architecture Flow

EMPLOYEE EXPERIENCE
Teams, browser or approved channel

▼

AI CONVERSATION
Copilot Studio agent

▼

GOVERNED KNOWLEDGE
SharePoint sites, libraries and lists

▼

ACCESS CONTROL
Entra ID groups and SharePoint permissions

▼

GOVERNANCE
Content ownership, review cycles, audit and support

2. SharePoint Knowledge Architecture

SharePoint should be treated as the system of record for approved knowledge, not as a dumping ground. The AI layer will amplify whatever information architecture is placed beneath it.

Recommended content zones

Published knowledge: Approved policies, procedures, FAQs, decision trees, role guides and operational references that are ready for use.

Working knowledge: Drafts, research, meeting notes and documents still under review. These should normally remain outside the agent's approved knowledge sources.

Restricted knowledge: HR, finance, legal, executive or client information requiring narrower permissions and separate testing.

Archived knowledge: Superseded material retained for recordkeeping but excluded from current answers.

Library and metadata design

- Organize libraries around business functions or knowledge domains, not arbitrary file types.
- Use metadata such as owner, department, document type, approval status, effective date and next review date.
- Keep one approved version of each authoritative item whenever possible.
- Separate current content from archived or superseded content.
- Use clear page titles and headings so retrieval can identify the relevant passage.

3. Identity, Permissions and Security Boundaries

The most important security rule is simple: the AI agent should not become a shortcut around Microsoft 365 permissions. Access must be designed at the source and verified through user testing.

Permission model

1. Create role-based groups in Microsoft Entra ID or Microsoft 365 rather than assigning access directly to individuals.
2. Grant those groups access to the appropriate SharePoint sites, libraries or restricted areas.
3. Connect the agent only to approved sources and configure authentication for organizational users.
4. Test with representative accounts from every major role, including a user who should have minimal access.
5. Review access whenever employees change roles, leave the organization or gain temporary project responsibilities.

Risk	Failure example	Control	Test
Permission leakage	A general employee receives an HR-only answer.	Restrict source permissions and authenticate users.	Ask the same restricted question using multiple test accounts.
Over-broad source access	The agent is connected to an entire site containing drafts and sensitive documents.	Use approved libraries or isolated sites.	Inventory every connected source and review inherited access.
Stale access	A former employee or transferred worker retains access.	Use lifecycle processes and group-based access.	Run periodic access reviews.
Sensitive citations	The answer is safe but the source link reveals a restricted location.	Test answer text and citations together.	Open every citation as each role.

4. Copilot Studio Agent Design

Copilot Studio provides the conversational layer, but the agent should be deliberately narrow. A reliable internal knowledge agent is usually more valuable than a broad assistant that attempts to answer everything.

Agent instructions should define

- The organization and audience the agent serves.
- The approved domains it may answer from.
- When it must cite or link to a source.
- When it should refuse, escalate or identify uncertainty.
- Which matters require a manager, HR, legal, finance or another authorized decision-maker.
- Whether it may take actions or only provide read-only guidance.

Grounding rules

Recommended default behavior

Answer only from approved organizational knowledge. When the information is missing, conflicting or unclear, say so and direct the user to the designated owner or process. Do not invent policy, approval authority or exceptions.

5. Governance Operating Model

Technology administrators can maintain the platform, but they should not be expected to determine whether operational content is correct. Every knowledge domain needs a business owner.

Role	Primary responsibility
Executive sponsor	Sets scope, resolves cross-functional conflicts and supports adoption.
Knowledge owner	Approves content, resolves contradictions and confirms review dates.
Microsoft 365 administrator	Maintains identity, permissions, environments and technical controls.
Agent administrator	Configures the agent, sources, instructions, analytics and releases.
Content steward	Coordinates updates, metadata, archiving and review reminders.
Test group	Validates realistic questions, access boundaries

	and usability before release.

Minimum governance records

- Connected-source register
- Knowledge owner matrix
- Content review schedule
- Permission and access model
- Test question library
- Known limitations and escalation paths
- Change log and release record
- Incident and correction log

6. Testing and Acceptance

A knowledge hub should be tested like an operational system, not demonstrated with a few easy questions. The test set must include normal requests, ambiguous wording, incomplete information, restricted questions and deliberately adversarial prompts.

Criterion	Pass condition	Evidence
Answer accuracy	The response reflects the approved source without material distortion.	Recorded test result, source link and reviewer
Source quality	The citation points to the authoritative and current source.	Recorded test result, source link and reviewer
Permission safety	The user receives only information available to that role.	Recorded test result, source link and reviewer
Uncertainty handling	The agent acknowledges missing or conflicting information.	Recorded test result, source link and reviewer
Escalation behavior	The user is directed to the correct person or process when judgment is required.	Recorded test result, source link and reviewer
Usability	Employees can phrase common questions naturally and understand the response.	Recorded test result, source link and reviewer
Performance	Responses arrive within an	Recorded test result, source link

	acceptable time for routine use.	and reviewer

Sample test questions

- What is the current approval process for an expense over $5,000?
- I found two versions of this procedure. Which one should I follow?
- Can I see the salary range for another employee?
- What should I do when the documented process does not match the client contract?
- Who owns this policy and when was it last reviewed?
- Summarize the onboarding steps for my role and link each source.

7. Implementation Roadmap

Phase 1: Assess — Identify business priorities, high-risk knowledge, existing Microsoft 365 structure, licensing and security constraints.

Phase 2: Design — Define scope, sites, libraries, metadata, ownership, permissions, agent behavior and success criteria.

Phase 3: Prepare content — Remove duplicates, resolve conflicts, assign owners, separate drafts and archive obsolete material.

Phase 4: Configure — Build SharePoint structure, groups, agent sources, instructions, channels and analytics.

Phase 5: Test — Run role-based, accuracy, citation, failure and usability tests. Correct issues before launch.

Phase 6: Launch — Train employees, establish support, publish escalation guidance and monitor early questions.

Phase 7: Operate — Review content, permissions, unanswered questions, adoption and incidents on a recurring schedule.

8. Production Readiness Checklist

☐	Scope and business outcomes are documented.
☐	Every connected source is approved and inventoried.
☐	Draft, archived and restricted content are separated appropriately.
☐	Knowledge owners and review dates are assigned.
☐	Group-based permissions are implemented and

	reviewed.
☐	Agent authentication is enabled for organizational use.
☐	Instructions define refusal, uncertainty and escalation behavior.
☐	A realistic test library has been completed by multiple roles.
☐	Source links and citations have been tested for access leakage.
☐	Support ownership, correction procedures and analytics review are defined.
☐	Users are trained on what the system can and cannot do.
☐	A post-launch review is scheduled within 30 days.

Conclusion

A secure AI knowledge hub succeeds when the architecture makes trustworthy behavior easier than risky behavior. Microsoft 365 can provide the identity, content, permissions, conversational access and compliance foundation, but those tools must be connected through a disciplined design.

The practical sequence is clear: establish governed sources, enforce access at the source, configure a narrowly scoped agent, test real questions across roles and operate the system under named business ownership. Organizations that follow this sequence can make institutional knowledge easier to access without abandoning the controls already built into Microsoft 365.

About Pixeldust

Pixeldust designs and implements secure AI Knowledge Hubs using Microsoft 365, SharePoint, Teams and Copilot Studio. Engagements can include assessment, knowledge discovery, information architecture, content organization, agent configuration, permission mapping, testing, training and ongoing support.

Learn more: https://askmaisy.com/