ONBOARD

How your board's most sensitive documents stay protected when AI is in the room.

How Your Board's Data Stays Protected When AI Is in the Room

Your board materials are among the most sensitive documents in your organization: financial forecasts, M&A discussions, legal strategy, personnel decisions. When AI touches that data, the security architecture behind it determines whether your information stays protected.

This document explains, in plain terms, how OnBoard's AI features work under the hood, how Microsoft Azure protects the data in transit, and why OnBoard's AI architecture is designed so that no board data reaches external systems or third-party AI providers without authorization.

How an AI Request Works at OnBoard

When a director asks OnBoard's AI to summarize a board book section, or when an administrator generates draft minutes, here's what actually happens:

```
Director clicks "Summarize" in OnBoard
        ↓
OnBoard sends the document text to Azure OpenAI Service
(encrypted in transit, over a private connection)
        ↓
Azure processes the request inside your dedicated instance
(no other customer's data is present, no shared model memory)
        ↓
Azure returns the AI-generated summary to OnBoard
        ↓
OnBoard displays the result — with the same permissions
as the original document
        ↓
```

```
Azure does not store the input or output.
No data is used for model training.
```

Every step in this chain is encrypted and authenticated. Azure discards the request data after processing — but OnBoard logs the interaction metadata (who made the request, which document was referenced, and when it happened) so your organization maintains a full audit trail. Let's break down each layer.

Layer 1: Authentication — Who Gets In

OnBoard uses **Microsoft Entra ID** (formerly Azure Active Directory) to authenticate every AI request. This is the same identity system used by Fortune 500 companies to protect their most sensitive cloud workloads.

What this means for your board:

- Every AI request is tied to a verified user identity, not an anonymous API key.
- Role-based access control (RBAC) ensures that only authorized users can trigger AI features. A director can summarize documents they have access to. The AI inherits the same permission boundaries.
- Multi-factor authentication (MFA) and single sign-on (SSO) through providers like Okta and Azure AD add additional layers before anyone reaches the AI features.

Layer 2: Network Security — How Data Travels

OnBoard connects to Azure OpenAI through **private endpoints** — dedicated, encrypted tunnels that never touch the public internet.

What this means for your board:

- When your board book text travels from OnBoard to Azure for summarization, it moves through a private network connection that never touches the public internet.
- There is no public API endpoint that an attacker could discover or target.
- All data is encrypted in transit using TLS 1.2+ (the same encryption standard used by banks for online transactions).

Layer 3: Data Isolation — What Happens Inside Azure

This is the most important layer for boards concerned about confidentiality.

Microsoft Azure OpenAI Service provides these guarantees:

PROTECTION	WHAT IT MEANS
No cross-customer data sharing	Your prompts and documents are processed in isolation. No other organization's data is present during your request.
No model training on your data	Your board materials are never used to train, fine-tune, or improve any public AI model. This applies to Microsoft, OpenAI, and every other provider in the chain.
No data retention by Azure	After the AI generates its response, the input and output are not stored by Azure. The request is processed and discarded.
No access by Microsoft employees	Microsoft personnel do not have access to your prompts or completions. Abuse monitoring uses automated systems, not human reviewers, for enterprise customers.
Geographic data residency	Your data is processed in the Azure region you or your platform provider selects. For organizations with data sovereignty requirements, this means your board data stays in your jurisdiction.

Azure processes the request, returns the result, and discards the input data. No prompts or responses persist after the request completes, and none of it feeds into public model training.

Layer 4: OnBoard's Additional Protections

Azure provides the infrastructure security. OnBoard adds governance-grade controls on top.

Permission Inheritance

When AI generates a summary of a board book section, that summary inherits the same access permissions as the source document. If a director doesn't have access to the compensation committee's materials, the AI can't generate summaries of those materials for them either.

Contained Architecture

OnBoard's AI runs entirely inside the portal. There is no way for a director to accidentally send board data to a public AI model through OnBoard. The system is architecturally contained:

- Prompts travel to Azure OpenAI through private endpoints and responses return the same way.
- There is no browser extension, no copy-paste to external tools, and no path for board data to reach a public AI model.

The architecture enforces this by default. There is no setting to override, and no way for a user to accidentally route data outside the system.

Audit Trail

Every AI interaction is logged: who requested it, what document it referenced, and when it happened. This creates the same defensible audit trail that boards expect from their other governance activities.

Encryption at Rest

All board data stored in OnBoard, including any AI-generated content that's saved (like draft minutes), is encrypted at rest using **AES-256 encryption**. This is the same encryption standard used by government agencies for classified information.

Certifications and Compliance

OnBoard and its underlying Azure infrastructure maintain the following certifications and compliance standards:

STANDARD	TYPE	WHAT IT COVERS
SOC 2 Type II	Certification	Independent verification that security, availability, and confidentiality controls are in place and operating effectively over time.
ISO 27001	Certification	International standard for information security management systems. Covers risk assessment, access control, incident management, and more.
ISO 27701	Certification	Privacy-specific extension of ISO 27001. Demonstrates compliance with data privacy regulations including GDPR and CCPA.

STANDARD	TYPE	WHAT IT COVERS
HIPAA	Compliance	For boards in healthcare — OnBoard provides security controls designed to align with HIPAA standards, offering a HIPAA-ready environment that facilitates the customer's compliance.
GDPR	Compliance	For boards with European operations or directors — OnBoard meets EU privacy standards for personal data processing.

Certifications are verified by independent auditors on a regular cadence. Compliance standards are met through architectural controls, contractual commitments, and the certified frameworks above.

How OnBoard Handles Meeting Recordings

The most common security question boards ask is about meeting recordings. Here's how it works:

Recording Is Optional

OnBoard's Minutes AI uses a meeting recording to generate draft minutes. The recording is initiated by your board administrator — OnBoard never records automatically. If your board isn't ready for recording, Book AI, Agenda AI, and Assist AI work without any recording at all.

The Recording Lifecycle

STEP	WHAT HAPPENS	SECURITY DETAIL
1. Recording	Administrator invites the OnBoard Recording Assistant to the Zoom/ Teams call	Joins as an attendee — only when invited
2. Transfer & Processing	Recording is captured from the meeting and transferred to OnBoard's secure infrastructure, where it is transcribed	Once OnBoard takes custody, the source copy is promptly deleted. Transcription uses the same private endpoints, encryption, and data isolation as all other AI features.
3. Draft Minutes	Minutes AI generates a first draft	Output inherits source document permissions

STEP	WHAT HAPPENS	SECURITY DETAIL
4. Review	Administrator reviews, edits, and approves minutes	A person always approves before anything becomes official
5. Deletion	Recording, transcript, and outlines are deleted as a set	Your retention policy governs the timeline

Key Protections for Recordings

- **Recordings are transferred to OnBoard's secure infrastructure for processing.** Once OnBoard takes custody of the recording, the source copy is promptly deleted. All transcription and AI processing occurs within OnBoard's contained environment.
- **Only the administrator has access.** Directors do not see raw recordings or transcripts. Access follows the same permission model as all board materials.
- **Recordings are deletable.** Most organizations delete recordings within 7–14 days of minutes approval. The approved minutes become the sole official record.
- **No model training on recordings.** The same Azure guarantee applies: your recording data is never used to train any public AI model.

> **Best practice:** Adopt a recording retention policy before enabling Minutes AI. A simple resolution — *"Recordings shall be deleted within [7/14] days of minutes approval"* — addresses the most common general counsel concern.

For a comprehensive FAQ on recording privacy, discoverability, and recommended board actions, see our **Recording Privacy & AI Governance FAQ**.

What This Means for Your Board

A summary for directors:

1. **Your data never trains any public AI model.** This applies across every provider in the chain, including Microsoft and OpenAI, and will continue to apply as models evolve.
2. **Your data never leaves the secure environment.** Private connections, encrypted transit, no public endpoints.

3. **Azure discards your data after processing.** No prompts or responses persist after the request completes.
4. **Your permissions still apply.** If you can't see a document, the AI can't summarize it for you.
5. **Every AI interaction is logged**: who requested it, which document was referenced, and when it happened. Your general counsel can audit the full trail.
6. **The certifications are real.** SOC 2 Type II, ISO 27001, ISO 27701 — independently verified, regularly renewed.

The Real Risk Is Ungoverned AI

69% of directors report using AI for board work. Many of them are pasting confidential materials into public tools like ChatGPT, Gemini, and Claude, where the data may be used for model training, stored indefinitely, or accessible to the provider's employees.

Directors are already using AI. The question is whether that usage happens inside your governance framework or outside it.

OnBoard keeps AI inside the system, where data is encrypted, access is controlled, permissions are enforced, and every interaction is logged. The architecture enforces these protections by default.

Questions Your Board Should Ask Any AI Vendor

If you're evaluating AI-enabled board management tools, ask these questions:

1. Is our data used to train any AI model? *(The answer should be no.)*
2. Where is our data processed geographically? *(You should be able to specify the region.)*
3. Is the AI connection private or does it traverse the public internet? *(Private endpoints are the standard.)*
4. Do AI outputs inherit the same access permissions as source documents? *(If not, AI creates a permissions bypass.)*
5. Can we audit every AI interaction? *(Logging is a prerequisite for governance.)*
6. What independent security certifications do you hold, and when were they last renewed? *(Certifications without a recent audit date are decorative.)*

For a live demonstration of OnBoard's AI Suite and security architecture, request a demo.

Technical details in this document are based on Microsoft Azure OpenAI Service documentation and OnBoard's published security practices as of February 2026. Consult your IT security team for organization-specific requirements.