

3½ Things

RESEARCH BRIEF · APRIL 2026

AI for Onboarding and Institutional Knowledge

Public examples in AEC and adjacent industries, with implications for a cautious enterprise rollout.

For information purposes only.

PREPARED BY

Jigar Desai, PhD, PE, MBA

DATE

April 23, 2026

OVERVIEW

Executive summary

This brief examines how firms are using AI for onboarding, internal learning, and institutional knowledge management, with implications for a US-based AEC firm considering a pilot. Three findings frame the opportunity.

First, the firms getting practical value are not deploying a standalone “AI onboarding bot.” They are building a three-layer stack: a governed knowledge layer (onboarding content, standards, expert interviews), a structured experience layer (resumes, projects, credentials), and AI interfaces that sit on top of both. Onboarding, internal learning, marketing, BD, and proposals draw from the same substrate.

Second, publicly documented AEC examples are real but still sparse. Shepley Bulfinch, Bora, BWBR, Skanska, VHB, and WSB each represent a different entry point: onboarding ecosystems, expert digital twins, internal learning programs for emerging professionals, safety-domain AI, technical-software tutors, and centralized experience databases for proposals. Outcome metrics in AEC remain largely qualitative. Outside-AEC examples (McKinsey, Super.com, Confluent) publicly report measurable time savings and adoption rates that suggest what mature deployments can deliver.

Third, security and compliance are manageable, not prohibitive. Enterprise-tier AI platforms are SOC 2 Type II attested and typically exclude enterprise data from model training by default. The practical question is not whether AI can be used compliantly, but whether the deployment sits inside the firm’s own governance layer: data classification, access controls, retention, audit, and human review. For a US-based firm, NIST AI RMF and SOC 2 are the most relevant reference frames. EU AI Act obligations generally apply only to firms with EU operations or customers.

The most credible first move is not a flashy agent. It is a narrow pilot with approved content, clear permissions, and one business problem that matters: onboarding, expert-finding, or experience retrieval for pursuits.

Research note

This brief is intentionally selective. It includes only examples that are publicly documented well enough to verify, and it distinguishes between strong public outcome metrics and illustrative public use-case evidence. In AEC specifically, the public record is still thinner than in software and consulting, so several AEC examples come from company-authored posts, vendor case studies, or trade-press interviews rather than independent third-party research. All quantitative claims have been cross-checked against the primary source listed in the references section.

THE ARCHITECTURE

01 How firms are actually doing this

Across the public examples reviewed, the recurring architecture is not a single “magic bot.” It is a stack with three layers. The first is a knowledge layer that holds onboarding guides, standards, policies, training content, expert interviews, and internal question-and-answer content. The second is an experience layer that holds resumes, project sheets, case studies, credentials, staffing history, and lessons learned. The third is an AI interface layer of search, chat, guided question-and-answer, proposal drafting support, and expert-finding tools that sit on top of approved content.

That pattern matters because it explains why some deployments help new hires ramp faster while also helping marketing and BD teams find the right project history and team qualifications without rebuilding the same content in separate systems. The experience layer, in particular, is where marketing, BD, and HR leverage converges, because proposal generation, capability positioning, and credentialing all draw from the same structured data.

02 · WHO IS DOING WHAT

02 Public examples in AEC

The table below separates what is clearly documented from what should be treated more cautiously. Several of the clearest AEC examples come from vendor-published case studies or interviews rather than independent third-party evaluations. That does not make them false, but it does mean they should be treated as verified use-case evidence rather than audited benchmarks.

Organization	What is publicly documented	Publicly visible value	Evidence basis
Shepley Bulfinch	An “I’m New Here” onboarding ecosystem inside the firm’s intranet, plus AI search that answers practical questions long after orientation, including staffing and office-specific expertise questions.	Strong evidence of a real onboarding use case. No quantified public KPI was identified.	Vendor case study (Knowledge Architecture)
Bora Architecture and Interiors	“RoboCorey,” a digital twin of Director of Sustainability Corey Squire, built from his 101- and 201-level best practices, project reflections, and his published book <i>People Planet Design</i> , made searchable through Bora’s Synthesis intranet.	Strong evidence of expert capture and knowledge reuse. No quantified public KPI was identified.	Vendor case study (Knowledge Architecture)
BWBR	Landmark Learning, a decade-long effort to help emerging professionals build judgment, confidence, and technical fluency faster, now paired with an integrated upskilling system using role-based guides, Communities of	Strong evidence of AI-assisted internal learning and onboarding support for emerging professionals. No	Knowledge Architecture interview and

Organization	What is publicly documented	Publicly visible value	Evidence basis
	Practice content, and AI search on approved content.	quantified public KPI was identified.	KA Connect 2025 talk
Skanska USA	A suite of internal AI tools collectively called Expert Sidekicks, including a general-purpose Skanska Sidekick and a Safety Sidekick trained on Skanska-specific safety data, the EHS Manual, and OSHA construction standards, delivered on mobile and desktop. Leadership publicly describes a “walk before you run” rollout.	Clear evidence of phased internal AI deployment in construction. The public record supports the named tools and describes Expert Sidekicks as a growing suite.	Company press release and Construction Dive interview
VHB	A Microsoft Copilot Studio-based “Bentley Copilot” embedded in Microsoft Teams, trained on Bentley documentation, manuals, and forum content, used as a real-time tutor for engineers learning Bentley software.	Strong evidence of AI-assisted technical onboarding. Public commentary describes faster learning and one reported instance of automating a design task estimated at thousands of hours into a single afternoon.	VHB news release and Bentley Systems blog interview
WSB	After growing from 450 to more than 1,600 employees through acquisitions, WSB centralized over 900 profiles of people and project data into Flowcase as a single searchable hub.	Publicly reported outcome of a 20 percent reduction in proposal preparation time. A proposal manager at the firm has publicly stated that tailoring ten resumes now takes under an hour, where it previously took days.	Vendor customer case study (Flowcase)

Industry-wide AEC adoption context

For a broader view, the 2026 Bluebeam AEC Technology Outlook, based on a global survey of more than 1,000 AEC technology decision-makers, found that only 27 percent of AEC firms currently use AI for automation, problem-solving, or decision-making. Among those that do, 68 percent report savings of at least 50,000 US dollars, and 46 percent report saving between 500 and 1,000 hours using AI tools. Ninety-four percent of current AEC AI users plan to expand investment in the coming year. The gap between adoption and outcome is the more revealing statistic: AEC firms that have started report meaningful ROI, but the sector as a whole is still in early innings.

03 · OUTSIDE AEC

03 Outside-AEC examples with stronger outcome metrics

These examples matter because they show what a more mature internal-knowledge deployment can look like once an organization has enough structured content, adoption support, and governance in place. The outcome metrics here are publicly reported by the companies themselves, often through vendor case studies.

Organization	What is publicly documented	Publicly visible value	Evidence basis
McKinsey	Lilli, the firm's internal generative AI platform, trained on roughly a century of proprietary documents, used for knowledge search, materials creation, proposal drafting, and expert-finding.	McKinsey reports that 72 percent of the firm is active on Lilli, generating over 500,000 prompts per month, with up to 30 percent time savings on knowledge-search and synthesis tasks.	Company article
Super.com	Glean deployed as a centralized knowledge hub across existing tools (Google Drive, Slack, Confluence, GitLab) to reduce time spent searching for documentation.	Publicly reported outcomes include employees finding information 20 minutes faster per day, over 1,500 hours saved monthly, and 20 percent shorter onboarding times.	Vendor customer case study (Glean)
Confluent	Glean deployed across more than 30 internal data sources to surface documents, past deal cycles, help content, and messages for distributed teams.	Publicly reported outcomes include more than 70 percent of employees actively using Glean, 5 to 10 minutes saved per support ticket in investigation time, and faster onboarding for new hires.	Vendor customer case study (Glean)

04 · TAKEAWAYS

04 Implications for an AEC firm

A few threads run through the stronger public examples. First, the firms getting meaningful value start with governed content, not the entire file universe. Approved onboarding guides, standards, FAQs, software how-to content, and benefits or policy documents are the safest starting ground. Second, the experience database

earns its keep early. The fastest visible win for marketing and BD teams usually comes from structured resumes, project sheets, credentials, and case studies that can be searched and reused across pursuits. WSB's experience is a representative example.

Third, capturing a few expert domains before trying to scale firmwide produces better results than trying to capture everything at once. Bora's RoboCorey and VHB's Bentley Copilot both illustrate the value of narrowing the problem to a high-value knowledge domain first. Fourth, security review belongs at the front end of design, not as a late-stage veto. Public vendor certifications help, but internal permissions, retention settings, auditability, and human review still determine whether a deployment is viable in practice. Fifth, the disciplined firms pilot in layers. A cautious sequence usually moves from internal question-and-answer on approved content, to expert-finding, to searchable experience data for proposals, and only then to more advanced drafting or agentic workflows.

05 · SECURITY AND COMPLIANCE

05 Security and compliance considerations

SOC 2 is often invoked as a blocker for enterprise AI. For enterprise-tier platforms, it typically is not. Claude for Enterprise, ChatGPT Business and Enterprise, Microsoft 365 Copilot, Google Gemini for Workspace, Glean, Flowcase, and Knowledge Architecture's Synthesis all publicly list SOC 2 Type II attestations, with enterprise data typically excluded from model training by default.

SOC 2 is an attestation of a vendor's operating controls. It is not a blanket approval of any given workflow, and it does not substitute for the organization's own data classification, access controls, retention settings, or audit policies. The practical question is whether the deployment sits inside the firm's governance layer.

Relevant reference frames for a US-based AEC firm

For a firm headquartered in the United States, the most relevant reference frames are the NIST AI Risk Management Framework and its Generative AI Profile (voluntary but increasingly expected by enterprise buyers), SOC 2 Type II and ISO 27001 for vendor evaluation, and the emerging ISO/IEC 42001 standard for AI management systems, which is starting to appear in enterprise vendor assessments.

State-level activity in the US is accelerating. The Colorado AI Act is currently scheduled to take effect on June 30, 2026, after being delayed from February 1, 2026, and remains subject to active amendment efforts in the 2026 legislative session. California's automated decision-making rules, covering employer use of AI in hiring, were finalized in 2025. More than 20 US states now have general privacy laws in effect. A firm operating in Pennsylvania and the Mid-Atlantic should monitor these state frameworks, particularly where AI touches hiring, housing, or credit decisions.

The EU AI Act has extraterritorial reach but applies to US firms only when they provide AI systems to EU customers, place AI outputs on the EU market, or have EU operations. For a US-based AEC firm using AI internally for staff onboarding and knowledge management, it is not directly applicable absent those conditions. It remains a useful directional signal for where enterprise AI governance expectations are heading.

Public posture of relevant enterprise tools

The table below summarizes publicly stated security and compliance positions from each vendor’s own documentation. These summaries are informational. Actual contract terms, retention settings, and approved data sources must still be verified against each organization’s internal governance.

Platform	What the official public record says	Why it matters for a cautious rollout
OpenAI ChatGPT Business, Enterprise, Edu, and API	Business data is not used for training by default. The SOC 2 Type 2 report covers the API Platform and ChatGPT business product services. ISO/IEC 27001 and 27701 certifications are tied to the API Platform, ChatGPT Enterprise, and ChatGPT Edu. Consumer ChatGPT (free and Plus) is not covered under SOC 2. On August 29, 2025, OpenAI renamed ChatGPT Team to ChatGPT Business.	Tier and plan selection matter for the exact scope of applicable attestations.
Anthropic Claude commercial products	SOC 2 Type II and ISO 27001 attestations for the Claude API and enterprise products. The SOC 2 Type II report covers a 12-month observation period across all five Trust Service Criteria. Anthropic does not train on commercial customer data by default.	Useful when a firm wants a commercial product posture that separates default enterprise use from optional data-sharing programs.
Google Workspace with Gemini	Workspace data stays in Workspace and is not used to train or improve underlying models outside Workspace without permission. Gemini for Google Workspace is SOC 1, SOC 2, and SOC 3 compliant.	Useful when the firm wants AI grounded inside existing Workspace identity, permissions, and admin controls.
Microsoft 365 Copilot and Copilot Chat	Prompts, responses, and Microsoft Graph data used by Copilot are not used to train foundation models, and enterprise data is not available to OpenAI under enterprise data protection. Microsoft 365 supports enterprise compliance requirements including SOC 2 and ISO 27001 attestations, FedRAMP authorizations, and HIPAA and HITECH support through Business Associate Agreements. There is no formal HHS-approved HIPAA certification; Microsoft positions its services as enabling customer compliance.	Useful when the firm wants AI grounded inside the Microsoft 365 identity, permissions, and compliance boundary.
Glean	Trust Center publicly lists SOC 2 Type 2 and ISO certifications. Public case studies describe measurable onboarding and knowledge-access gains at Super.com and Confluent.	Relevant for enterprise search and agent layers that sit across many systems.

Platform	What the official public record says	Why it matters for a cautious rollout
Flowcase	SOC 2 Type II and ISO 27001 certifications. Customer data is not used to train AI models. Focused on resume, CV, and case-study management for professional services and AEC firms.	Relevant for the experience-database layer that marketing, BD, HR, and proposal teams can share.
Knowledge Architecture (Synthesis)	SOC 2 and GDPR compliant, with a Trust Center powered by Drata providing real-time transparency into security posture, attestations, reports, and policies.	The platform behind many of the AEC case studies cited in this brief, including Shepley Bulfinch, Bora, and BWBR.

SOURCES

References

- AICPA & CIMA. (n.d.). System and Organization Controls: SOC suite of services. <https://www.aicpa-cima.com/resources/landing/system-and-organization-controls-soc-suite-of-services>
- Anthropic. (n.d.). What certifications has Anthropic obtained? Anthropic Support. <https://support.anthropic.com/en/articles/10015870-do-you-have-a-soc-2-or-hipaa-certifications>
- Bentley Systems. (2025, December 12). This engineer gave AI a job: It learns the software so he doesn't have to. Bentley Blog. <https://blog.bentley.com/insights/this-engineer-gave-ai-a-job-learn-the-software-so-he-doesnt-have-to/>
- Bluebeam. (2025, October 28). New Bluebeam report shows early AI adopters in AEC seeing significant ROI despite uneven adoption [Press release]. <https://press.bluebeam.com/2025/10/new-bluebeam-report-shows-early-ai-adopters-in-aec-seeing-significant-roi-despite-uneven-adoption/>
- Construction Dive. (2025, May 21). With AI, "walk before you run": Skanska USA exec. <https://www.constructiondive.com/news/ai-skanska-sidekick-safety-project/748732/>
- Flowcase. (n.d.). Data processing. <https://www.flowcase.com/data-processing>
- Flowcase. (n.d.). WSB case study. <https://www.flowcase.com/customers/wsb>
- Glean. (n.d.). Confluent customer story. <https://www.glean.com/resources/customer-stories/confluent>
- Glean. (n.d.). Super customer story. <https://www.glean.com/resources/customer-stories/super>
- Google Workspace Updates. (2024, August 16). SOC compliance for Gemini. <https://workspaceupdates.googleblog.com/2024/08/gemini-soc-compliance.html>
- Knowledge Architecture. (2025, July 10). KM 3.0 case study: Improving onboarding effectiveness at Shepley Bulfinch. <https://www.knowledge-architecture.com/blog/km-30-case-study-improving-onboarding-effectiveness-shepley-bulfinch>
- Knowledge Architecture. (2025, July 10). KM 3.0 case study: Scaling expert knowledge at Bora. <https://www.knowledge-architecture.com/blog/km-30-case-study-scaling-expert-knowledge-bora>

- Knowledge Architecture. (2026, January 15). Redesigning learning for the next generation of AEC talent: Dan Hottinger and Kari Shonblom of BWBR. <https://www.knowledge-architecture.com/blog/redesigning-learning-for-the-next-generation-of-acc-talent-dan-hottinger-and-kari-shonblom-of-bwbr>
- Knowledge Architecture. (n.d.). Synthesis: The intranet + AI search solution for AEC firms. <https://www.knowledge-architecture.com/synthesis>
- McKinsey & Company. (n.d.). Rewiring the way McKinsey works with Lilli, our generative AI platform. <https://www.mckinsey.com/capabilities/tech-and-ai/how-we-help-clients/rewiring-the-way-mckinsey-works-with-lilli>
- Microsoft. (2026). Data, privacy, and security for Microsoft 365 Copilot. Microsoft Learn. <https://learn.microsoft.com/en-us/microsoft-365/copilot/microsoft-365-copilot-privacy>
- National Institute of Standards and Technology. (2024). Artificial intelligence risk management framework: Generative artificial intelligence profile (NIST AI 600-1). <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence>
- OpenAI. (2025, August 29). ChatGPT Business rename FAQ. OpenAI Help Center. <https://help.openai.com/en/articles/12111915-chatgpt-business-rename-faq>
- OpenAI. (n.d.). Enterprise privacy at OpenAI. <https://openai.com/enterprise-privacy/>
- OpenAI. (n.d.). Trust Portal. <https://trust.openai.com/>
- Reed Smith. (2025, August 28). Colorado AI Act delayed, further revisions likely. <https://www.reedsmith.com/our-insights/blogs/viewpoints/10210wf/colorado-ai-act-delayed-further-revisions-likely/>
- Skanska USA. (2025, April 30). Skanska launches “Safety Sidekick” AI tool to advance job site safety across the U.S. [Press release]. <https://www.usa.skanska.com/who-we-are/media/press-releases/298531/Skanska-launches-Safety-Sidekick-AI-tool-to-advance-job-site-safety-across-the-U.S>
- VHB. (2025, March 31). VHB’s AI-powered Bentley tool shared with industry leaders. <https://www.vhb.com/news/vhb-shares-ai-powered-bentley-copilot-tool/>

For information purposes only. Reader and user must verify information for relevancy and accuracy as things may change since they were included in this brief. Prepared as a research brief. This document reflects public information available as of April 23, 2026. Contract terms, retention settings, and internal governance controls must still be verified against each organization’s own policies before deployment decisions are made.