



AAAI-D WHITE PAPER

AI Agents in the Dental Practice: A HIPAA Compliance Guide

Published by the American Academy of Artificial Intelligence in Dentistry

Conceived January 2026

Builds on peer-reviewed work on agentic AI in dentistry

(Farooqi et al., *European Journal of General Dentistry*, May 2026)

Revised and published October 2026

Los Angeles, California

Contents

1. The Traffic Light: Know Your Zone
2. Understanding AI Agents: What Makes Them Different
3. Business Associate Agreements: Reading the Fine Print
4. Mapping PHI Flow: Follow the Data
5. Minimum Necessary: The Most Violated Principle
6. Human Oversight: What Agents Should Never Do Alone
7. Auditability: Knowing What the Agent Did
8. Retention, Training Data, and Secondary Use
9. Prompt Injection: The Agent-Specific Threat
10. Recording Consent: The Law Beside HIPAA
11. The Expanding Privacy Landscape
12. Vendor Evaluation: Questions to Ask
13. Implementation Checklist
14. Frequently Asked Questions
15. Why the Academy Published This Guide
16. References

Published by the American Academy of Artificial Intelligence in Dentistry

Conceived January 2026. Builds on the authors' peer-reviewed work on agentic AI in dentistry (Farooqi et al., European Journal of General Dentistry, May 2026). Revised and published October 2026.

Any dentist can now build an AI agent in an afternoon. A scheduling assistant. A patient intake bot. An AI scribe that transcribes visits. An automated insurance verification tool. The platforms — Claude, ChatGPT, Microsoft Copilot, Google Gemini, and a growing ecosystem of no-code builders — have made it that easy.

That ease is the reason for this guide.

An agent is not a chatbot. A chatbot answers questions. An agent takes actions: reading your schedule, accessing patient records, sending messages, updating appointments, connecting to other software — operating on instructions you gave it once, then acting on its own.

Every capability you grant an agent is a door into your practice's data. This guide helps you know which doors are safe to open, which need a second look, and which should stay shut.

IMPORTANT DISCLAIMER — PLEASE READ

This white paper is published by the American Academy of Artificial Intelligence in Dentistry (AAAI-D) for educational and informational purposes only. It does not constitute legal advice, and no attorney-client relationship is created between AAAI-D and any reader by the publication or use of this document.

HIPAA, state privacy laws, and related regulations are complex, jurisdiction-specific, and subject to change. The information contained herein reflects the Academy's understanding as of the publication date noted below and may not reflect subsequent regulatory developments, enforcement actions, guidance documents, or case law.

Every dental practice operates under unique factual and legal circumstances. Readers must consult qualified legal counsel licensed in their jurisdiction before making compliance decisions, implementing AI systems that handle patient information, or taking or refraining from any action in reliance on this document.

AAAI-D makes no warranties or representations, express or implied, regarding the accuracy, completeness, currency, or applicability of the information contained in this white paper. To the fullest extent permitted by applicable law, AAAI-D disclaims all liability for any loss, damage, claim, or expense arising from or related to the use of or reliance on this document.

References to specific products, vendors, or services are for illustrative purposes only and do not constitute endorsement, recommendation, or certification by AAAI-D.

The Traffic Light: Know Your Zone

Before building or deploying any AI agent, locate your use case:

Use Case		Dental Example	What It Takes
●	No patient information involved	Website FAQ assistant (services, hours, insurance accepted)	Business/enterprise tier; confirm vendor data-use terms
●	Internal content from public information	Drafting social posts, patient education handouts	Same as above
●	Anonymized scheduling availability	Showing open slots without revealing patient identities	Same as above
●	Clinical reference without patient data	Asking about drug interactions, material properties, technique guidance	No PHI entered; verify no conversation logging for training
●	Accesses patient names/contact info	Appointment reminders, recall campaigns	Signed BAA; purpose-limited access; audit logging; no consumer AI accounts
●	Transcribes patient encounters	AI scribe during exams or hygiene visits	Signed BAA; patient recording consent; verified storage/retention; all-party-consent state laws addressed
●	Collects clinical information	Digital intake gathering medical/dental history	Full compliance stack: BAA, encryption, access controls, retention policy, security risk analysis
●	Reviews images or radiographs	Agent pre-screening X-rays for findings	BAA with imaging AI vendor; confirm where images are processed/stored; human dentist makes final diagnosis
●	Insurance verification	Agent checking eligibility and benefits	BAA; minimum-necessary data sharing; audit trail of each query
●	Treatment plan communication	Agent drafting treatment explanations for patients	BAA; human review before sending; no autonomous patient communication
●	Any PHI in a consumer AI account	Pasting patient notes into free ChatGPT, Claude, or Gemini	Do not do this — consumer tiers are never BAA-covered
●	Unsupervised clinical documentation	Auto-charting that writes directly to the patient record	Do not do this — require human review and confirmation
●	Autonomous patient messaging	Agent sending treatment recommendations or follow-ups without oversight	Do not do this — restrict to drafts requiring human approval
●	Assumed BAA coverage	"The vendor says HIPAA compliant"	Do not assume — verify the exact product, features, and configuration

●	Broad system access for narrow tasks	Reminder agent with full EHR read access	Do not do this — limit to minimum necessary
●	PHI transmitted to non-BAA subcontractors	Agent using a transcription API, analytics tool, or cloud function without BAA coverage	Do not do this — every handler of PHI needs contractual coverage

The rule: Green means proceed with standard precautions. Yellow means stop — get privacy, IT, and legal review before deploying. Red means stop, period.

Understanding AI Agents: What Makes Them Different

Most dentists have used a chatbot. An agent is a step beyond, and the compliance implications are proportionally greater ^[6].

A **chatbot** receives a prompt and returns text. The interaction is contained. When you close the window, it is over.

An **agent** receives an objective and pursues it across multiple steps and systems. You might instruct it to "follow up with patients who missed their hygiene appointments this month." To do that, it needs to: access your schedule, identify the right patients, retrieve their contact information, compose messages, send them, log the outreach, and possibly reschedule — each step touching a different system, each system a potential PHI exposure point.

This multi-step autonomy creates compliance questions that do not exist with simple chatbots:

- **Scope creep:** An agent granted broad access may retrieve more data than its task requires
- **Tool chaining:** Each connected tool (EHR, email, calendar, cloud storage) becomes part of the PHI flow
- **Persistent memory:** Some agents retain context across sessions, creating ongoing data stores that need governance
- **External content:** Agents that read emails, web pages, or attachments can encounter hidden malicious instructions (prompt injection)

The more tools an agent connects to and the more actions it can take independently, the more rigorous your compliance review must be.

Business Associate Agreements: Reading the Fine Print

A BAA is the legal foundation for any vendor handling your patients' PHI. But not all BAAs are equal, and not all vendor claims mean what they seem to.

What to verify in any BAA for AI tools:

1. **Exact product coverage.** The BAA should name the specific product, tier, and configuration. A BAA covering a vendor's enterprise API does not automatically cover its consumer chatbot, mobile app, or browser extensions.
1. **Feature-level exclusions.** Some vendors exclude specific features from BAA coverage — beta tools, third-party integrations, analytics dashboards, or training-data pipelines. Ask explicitly.
1. **Data residency and processing locations.** Where is PHI processed and stored? If the vendor uses subprocessors (and most do), each one handling PHI needs to be covered.
1. **Retention and deletion terms.** How long is data kept? Can you request deletion? Is deleted data truly removed from backups and logs?
1. **Training data use.** Does the vendor use your data to train or improve its models? Under what conditions? Can you opt out? This must be explicit, not assumed.
1. **Breach notification terms.** What is the vendor's obligation to notify you, and on what timeline?
1. **Configuration requirements.** Many BAAs require specific settings — zero data retention mode, HIPAA-specific endpoints, disabled logging features. Using the wrong configuration can void your coverage.

A real-world example: Anthropic offers a BAA, but it applies only to specific HIPAA-eligible service configurations (enterprise deployments, zero-data-retention API setups). It does not cover consumer Claude products. A dentist who signs up for Claude Pro and pastes in patient notes is not protected, regardless of what the marketing page says.

The principle: Never accept "HIPAA compliant" as a product label. Compliance is a property of your specific deployment — the right service, under the right contract, with the right settings, used the right way.

Mapping PHI Flow: Follow the Data

Before deploying an agent, trace every path patient information can take. A typical dental AI agent might touch:

- **Practice management system** (Dentrix, Eaglesoft, Open Dental) — appointment data, treatment plans, billing
- **Imaging systems** — radiographs, intraoral scans, CBCT volumes
- **Communication tools** — email, SMS reminders, patient portals
- **Cloud storage** — document storage, backup systems
- **Transcription services** — audio processing for AI scribes
- **Analytics and logging** — usage dashboards, audit logs, error reporting
- **The AI model itself** — prompts, context, outputs, and any retained memory

Each system that creates, receives, maintains, or transmits PHI is part of your compliance surface. Each needs either a BAA (if a third party) or appropriate safeguards (if internal). A compliant AI model in the middle of a noncompliant workflow does not make the workflow compliant.

Practical exercise: Draw your agent's data flow on paper. Every arrow that carries patient information needs a compliance answer. If you cannot draw the complete picture, you are not ready to deploy.

Minimum Necessary: The Most Violated Principle

HIPAA's minimum necessary standard requires that PHI access be limited to what each task genuinely requires. AI agents routinely violate this principle — not from malice, but from convenience.

It is easier to grant an agent broad read access to your EHR than to carefully scope its permissions. But that convenience creates risk:

- An appointment-reminder agent does not need access to clinical notes, radiographs, or financial records
- An insurance verification agent does not need the patient's full chart — just eligibility-relevant demographics
- A recall campaign agent does not need treatment history beyond the recall interval

Design principle: Start with zero access. Add only what the specific task requires. Review quarterly. The question is never "can the agent access this?" — it is "does the agent need this?"

Human Oversight: What Agents Should Never Do Alone

The defining risk of AI agents is autonomous action. Establish clear boundaries:

Require human confirmation before:

- Writing to the clinical record
- Sending any patient communication
- Changing, cancelling, or creating appointments
- Submitting insurance claims or pre-authorizations
- Transmitting data to any external party
- Deleting or modifying existing records

The framework: Reading is lower risk than writing. Drafting is lower risk than sending. Suggesting is lower risk than deciding. Calibrate oversight to the consequence of error.

An agent that drafts a treatment plan explanation for dentist review is a productivity tool. An agent that sends treatment recommendations directly to patients is a liability.

Auditability: Knowing What the Agent Did

When software acts on behalf of your practice, you need to know what it did. This is both a HIPAA requirement (audit controls) and a practical necessity.

Your agent deployment should log:

- What patient data the agent accessed, and when
- Which tools and APIs it called
- What information it transmitted, and to whom
- What actions it took (messages sent, records modified, appointments changed)
- Authentication events (who authorized the agent's actions)

Ask your vendor:

- Are audit logs available to you, or only to the vendor?
- How long are logs retained?
- Can logs be exported for your own records?
- Who at the vendor can access logs containing your PHI?

If you cannot answer "what did the agent do last Tuesday?" you have an audit gap.

Retention, Training Data, and Secondary Use

What happens to patient information after the agent processes it? This question has more layers than most practices realize:

- **Prompt retention:** Are your inputs stored? For how long? Who can access them?
- **Output retention:** Are the agent's responses logged? Where?
- **Model training:** Is your data used to train or fine-tune the vendor's models? Can you opt out? Is opt-out the default?
- **Vendor personnel access:** Under what circumstances can the vendor's employees view your data?
- **Memory and context:** Does the agent retain information across sessions? Can that memory be inspected, corrected, or deleted?
- **Downstream sharing:** Is data passed to subprocessors, analytics providers, or integration partners?

Get written answers to each question. "We take privacy seriously" is not an answer.

Prompt Injection: The Agent-Specific Threat

Traditional cybersecurity focuses on keeping unauthorized people out. AI agents introduce a different vector: unauthorized instructions hidden in content the agent legitimately reads.

An agent that processes patient emails, reads web pages for drug information, or parses intake forms can encounter **prompt injection** — malicious instructions embedded in otherwise normal content. Examples:

- An email containing hidden text instructing the agent to forward patient data to an external address
- A webpage with invisible instructions to modify the agent's behavior
- An intake form field containing instructions disguised as patient information

The more authority an agent has, the more damaging a successful injection becomes. This is not hypothetical — the National Institute of Standards and Technology (NIST) has identified prompt injection as a key AI risk category in its Generative AI Profile, and 2026 federal initiatives are developing agent-specific security controls ^[2].

Mitigations:

- Limit agent permissions to the minimum required (see above)
- Require human confirmation for consequential actions (see above)
- Treat external content as untrusted input, not instructions
- Keep agents updated — vendors are actively developing injection defenses
- Include prompt injection in your security risk analysis

Recording Consent: The Law Beside HIPAA

If your agent records patient conversations — as AI scribes do — HIPAA is not your only legal obligation.

In **all-party consent states** (including California, Florida, Illinois, Pennsylvania, Washington, and others), recording a conversation without the consent of all parties can create civil and criminal liability independent of HIPAA. A California class action filed in January 2026 targets an ambient AI documentation tool on exactly these grounds.

What this means for dental practices:

- Obtain clear patient consent to record before using any AI scribe or ambient documentation tool
- Consent to record is separate from HIPAA authorization — you need both where both apply
- Post visible notices in addition to obtaining individual consent
- Understand that consent requirements vary by state — a workflow compliant in Texas may not be compliant in California

The Expanding Privacy Landscape

HIPAA is the federal floor, not the ceiling. Dental practices deploying AI agents should be aware of additional legal frameworks:

State privacy laws. The California Consumer Privacy Act (CCPA) as amended by CPRA, Washington's My Health My Data Act (which covers health-adjacent data beyond HIPAA's scope and provides a private right of

action), and similar laws in Colorado, Connecticut, Virginia, and other states may impose requirements beyond HIPAA — including data minimization, purpose limitation, and consumer rights to access and delete.

FTC enforcement. The Federal Trade Commission has pursued health data misuse under its Section 5 authority, including a July 2026 action against a telehealth company for sharing health data with advertising platforms. The FTC has also warned against "AI washing" — overstating AI capabilities or privacy protections.

The HIPAA Security Rule update. HHS proposed a major strengthening in January 2025: mandatory encryption of electronic PHI, multi-factor authentication, annual security risk analyses, asset inventories, and incident response plans. The rule has not been finalized as of late 2026, but the direction is unmistakable. Build your AI workflows as though these requirements are coming.

Vendor Evaluation: Questions to Ask

Before selecting an AI agent vendor for your practice, get written answers to:

1. Will you sign a BAA? Which exact products, features, and configurations does it cover?
1. Where is our data processed and stored? Who are your subprocessors?
1. Is our data used for model training? Can we opt out? Is opt-out default?
1. What are your data retention and deletion policies? Can we request complete deletion?
1. What audit logs are available to us? How long are they kept? Can we export them?
1. What happens in a breach? What is your notification timeline and process?
1. What security certifications do you hold (SOC 2, HITRUST, ISO 27001)?
1. How do you address prompt injection and agent-specific security risks?
1. What are your authentication and access control capabilities? Is MFA supported?
1. Can we configure data residency restrictions (e.g., US-only processing)?

If a vendor cannot or will not answer these questions in writing, that is your answer.

Implementation Checklist

Before any AI agent touches patient information:

- ☐ Confirm whether the agent will access PHI — if yes, consumer AI accounts are prohibited
- ☐ Execute a BAA covering the exact product, features, and configuration
- ☐ Verify required privacy settings are enabled (zero retention, HIPAA endpoints, etc.)
- ☐ Map all connected systems and confirm PHI handlers are covered
- ☐ Apply minimum necessary access — start with zero, add only what is needed
- ☐ Configure human confirmation for consequential actions
- ☐ Enable and verify audit logging
- ☐ Document retention, training-data, and secondary-use policies in writing
- ☐ Address recording-consent requirements under applicable state law
- ☐ Assess prompt injection risk proportional to agent authority
- ☐ Conduct or update your HIPAA security risk analysis
- ☐ Train staff on the agent's capabilities, limitations, and escalation procedures
- ☐ Establish an incident response plan specific to AI agent failures
- ☐ Schedule quarterly access and configuration reviews

Frequently Asked Questions

Can I use ChatGPT to help write patient letters if I remove names?

De-identified data is not PHI under HIPAA, but true de-identification is harder than removing names. Dates, locations, rare conditions, and combinations of details can re-identify patients. If there is any doubt, use a BAA-covered service.

My EHR vendor offers an AI feature. Is it automatically compliant?

No. Confirm that the specific AI feature is covered under your existing BAA or requires an amendment. Ask the vendor directly.

Can an AI scribe replace my hygienist's clinical notes?

An AI scribe can assist with documentation, but the licensed clinician remains responsible for the accuracy and completeness of the clinical record. AI-generated notes require human review before becoming part of the record.

What if a patient asks whether AI was involved in their care?

Transparency is good practice. Patients have a right to understand how their information is used. Consider disclosing AI involvement in your notice of privacy practices.

Does HIPAA apply if I only use the agent for scheduling?

If the agent handles patient names, phone numbers, or appointment details linked to individuals, that is PHI and HIPAA applies. Truly anonymized availability (open slots with no patient linkage) is not PHI.

Why the Academy Published This Guide

Joanna Stern, the Wall Street Journal columnist, recently described AI in healthcare as "an amplifier" — surfacing more information without resolving what to do with it. In dentistry, that amplification cuts both ways. The same technology that helps a clinician catch what the eye misses can also magnify confusion, overreach, and broken trust when deployed without guardrails.

The Academy's position is straightforward: the technology is not the problem. The absence of shared standards for its responsible use is.

This guide is our contribution to building those standards. We will update it as the technology, the regulations, and the threat landscape evolve.

References

1. Anthropic, "Covered Models under a Business Associate Agreement (BAA)," Claude Help Center, updated July 1, 2026. <https://support.claude.com/en/articles/15455031-covered-models-under-a-business-associate-agreement-baa>
1. National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework: Generative AI Profile," NIST AI 600-1, July 2024. <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-ai-profile>
1. U.S. Department of Health and Human Services, "HIPAA Security Rule to Strengthen the Cybersecurity of Electronic Protected Health Information," Notice of Proposed Rulemaking, 90 FR 898 (January 6, 2025).
1. HIPAA Journal, "HIPAA Updates and HIPAA Changes in 2026," updated 2026. <https://www.hipaajournal.com/hipaa-updates-hipaa-changes/>
1. Joanna Stern, *I Am Not a Robot: My Year Using AI to Do (Almost) Everything* (2026). <https://www.goodreads.com/book/show/241094437-i-am-not-a-robot>
1. Farooqi O, DeNucci DJ, Moharrami M. Agentic Artificial Intelligence in Dentistry: A Practice-Focused Perspective for the Dental Team. *European Journal of General Dentistry*. 2026;15(03): 303-305. doi:10.1055/s-0046-1822820

Published by the American Academy of Artificial Intelligence in Dentistry.

For questions about this guidance, contact info@aaai-d.org.

Last updated: October 2026.

This document is subject to the Important Disclaimer set forth at the beginning of this white paper, which is incorporated herein by reference.

How to Cite This White Paper

American Academy of Artificial Intelligence in Dentistry. *AI Agents in the Dental Practice: A HIPAA Compliance Guide*. Los Angeles, CA: AAAI-D; October 2026.

American Academy of Artificial Intelligence in Dentistry. *AI Agents in the Dental Practice: A HIPAA Compliance Guide*. Los Angeles, CA: AAAI-D; October 2026. Available at: <https://aaai-d.org/white-papers/>