AI In Benefits: What You Need to Know to Keep Your Data Safe While Using Artificial Intelligence Tools

By Dorothy Cociu

Your grandmother Peggy is sitting quietly at home watching her favorite TV show when she gets a very frightening phone call. It's someone who says they have you, her cherished grandchild, and Peggy has to pay X dollars in Bitcoin to get you back. Then they let her hear your voice, begging you to please pay so you can go home. We've all heard the stories on the news about scammers using artificial intelligence to create videos, images and audio that sound like our loved ones and use it to trick you (or your grandparent), scare you and, quite often, take your money. Today, it's not just a phone call. It could be a message sent to your grandmother's phone or email, with a video from what looks and sounds just like you, and it's even more terrifying. Unfortunately, Grandma Peggy has no idea it's an AI deep fake. Less than two years ago, it was just something that happened to others, scattered stories in the news. But now, AI tools are everywhere, and literally everyone and anyone can learn to use them. Quite often, large platforms create videos to train you on exactly how to create those video and audio clips that bad actors can use to do just that, with very little effort.

AI is everywhere, whether we want it to be or not. Sure, it can be a great tool to speed up your work, and it can help make your presentations more attractive and engaging, and your videos more compelling to watch. But there are a lot of dangers when working with AI—most often when you don't even realize that AI is being used! can be heart attack symptoms.

The basic dangers of AI finding its way into your data

Combine the basic dangers of using AI with the restrictions on industries like health insurance, where there is a ton of confidential and private information, including medical information, sitting in what you think are very secure areas of your systems, encrypted and protected. And yet, every day, confidential data is leaking, being used in Large Language Models (LLMs) or worse, being spread across the internet. Often times, the leakage happens even though the company has put in protections, because the employees of your company are using unauthorized AI tools that open the doors and allow the AI to work its way into your systems, sometimes without you knowing it. I ask you: do you know if or when your staff is using AI? Do you know what they are using and how or why? Are there restrictions on what is fed into the AI tool with prompts to get the responses that the employee wants? I asked principals of Aditi Group (technology & IT consultants) what words of caution they would share with individuals and companies as they move more and more towards AI tools in their everyday lives.

"First of all, 'AI' is used loosely and can really mean a slew of different things," stated Ted Flittner of Aditi Group. "Some position algorithms as AI. OCR. Voice-to-text is AI. AI Agents can be a simple researcher while others can take actions—do things on their own. Do your research. Read the privacy policies, security practices, certifications and use policies. And understand a little about how AI works. The fact is that AI models are black boxes. Anthropic CEO Dario Amodei is famously quoted as saying 'People outside the field are often surprised and alarmed to learn that we do not understand how our own AI creations work. They are right to be concerned: this lack of understanding is essentially unprecedented in the history of technology.'" For example, you cannot audit a decision you cannot explain. Flittner continued, "AI models are not like software programs that can be debugged and corrected in code. Models are trained. Like people are. Biases in training become part of an AI model or AI agent. Some AI models have hallucinated answers and, in some cases, have fabricated fake data. Others have been shown to act counter to their guardrails."

Flittner suggests "Treat AI like you would when hiring people. Would you hire someone without any background check, references, work history or interview? Would you immediately give them access to company and customer confidential data immediately? Would you give them login credentials to your bank, social media accounts and email? Would you pass their work directly to customers without reviewing it first?

Yet that's what many people have done by racing to try out new AI tools. Consider OpenClaw personal AI agent. In a span of several months, it changed names three times as security researchers found glaring security holes, while users turned over logins to all their social media, email and other accounts. The lure of a tool that can do so much for us makes many people turn blind eyes to security and privacy."

Some additional suggestions from Flittner included the following: "Do research before adopting AI. Use it to enhance work your people do. But keep a human layer—make sure people review AI work before using it or passing it on. If AI does research or makes content, have AI give you the sources and check them."

I also asked some experts to share their thoughts on my Benefits Executive Roundtable podcast, Season 7, Episode 8 (which aired on March 3, 2026).

When asked about the explosion of AI in the past few years, Eric Barricklow, owner of Stellar Cyber Solutions, (BER S7E8 plus panelist on the CAHIP-OC Annual Sales Symposium AI Panel, which I moderated, on March 10, 2026) stated: "The constant use of AI is prevalent; it's everywhere, and I'm not surprised it's exploded into use like it has now, but I am a little surprised that it hasn't imploded a little bit yet, because I used to say 'well, anything you need, there's an app for that..' Now, anything you need, there's an AI for that. But there are like 15 AIs for that! I think we're still in that explosive innovation phase of throwing spaghetti against the wall to see what actually works, and there's a lot of things that are being taken advantage of right now, because people don't know how AI works. Businesses, small businesses especially, are inviting AI into their environment, and it's an unregulated employee with complete access. That's one aspect I'm a little shocked at—how welcoming businesses are at inviting in this unknown, unfathomable technology, to really crawl through all their intellectual property, all their data, all their contacts, and they are just trusting."

As Barricklow said, your employees are using AI, and if you're not monitoring it, not putting up the proper guardrails to protect your data, it can and will leak outside of your protected environment and into the unknown areas of these learning models and possibly onto the internet or on the front page of a newspaper for everyone to have access to.

Who is using AI and what are they using it for?

In my podcast, Miguel Villegas (Mike), founder of iSecure Privacy, discussed further what very common uses are today in the workplace, regardless of the industry you work in. “Today they are using it every day. For everything,” said Villegas. “Productivity with Co-Pilot, drafting emails, summarizing meetings, turning notes into action items, building slide decks and rewriting policies. It’s amazing. It even helps us research benefit plans or gather evidence for claims processing. I use it for writing code. I can write a Python program in less than two minutes, but I still have to de-bug it and everything. But it’s amazing how much of a productivity tool [it can be]. It’s been able to allow me to be much more efficient.”

I asked Barricklow who he sees using AI much more predominantly than others today. “I see the younger generation adopting it much more readily and using it to compose emails, do their LinkedIn queries, their outreach, their cold calling. Also, for code writing and data analytics, and to actually rely on its business decisions with whatever data they input into it.” We’ll get more into why these uses can be dangerous later in this article. Barricklow continued, “Many times they’ll ask the AI [tool] ‘tell me what this means.’ [For] many businesses, there is a disconnect between the old guard, who does it their way, and the newer generation, who are adopting AI to find out what it can do.

But there is also that difference, because I’ve always been told that you have to show your work. If you arrive [at a] business decision, you have to have the algorithms, the calculations and how you derived at that decision in order to stand up in court, in case you go to legal action, or something along those lines. In many cases, in trusting AI to make a decision, you lose that transparency. You lose that ability to defend that decision that you’re making. I see a lot of businesses employing AI or accepting its use without understanding the business risk to it.”

That’s where the disconnect is. That’s where the risk is. Your employees may be doing what they think is good for the company or making their jobs easier by using an AI tool like ChatGPT, Google Gemini or Microsoft Pilot, using a free or individual (not enterprise) version of one or more of these or others, and they’ve never been told, never been trained to know that by using this little AI tool everyone is using, that they’ve put your data and your business at risk.

Biggest Security Risks Organizations Face When Adopting AI Tools

In general, I would say that the biggest risks that organizations face when using AI tools are data privacy, model misuse, lack of governance and, most importantly, the people using the AI tools themselves.

I asked this question of my guests on my AI Security podcast. Adriana Mendietta, a cyber insurance agent and 'Avid AI Agile User,' replied very matter-of-factly. The biggest risk is, according to Mendietta, "The eighth layer, which is the user. Whatever they are pasting into their tool. Whatever their settings are. A free versus an enterprise software. If the company hasn't provided a software and they're just using a good ole free willy internet [product] and pasting information in there, that's a huge risk! It sounds so simple. We can all use a mouse." And she's right. We all use a mouse. We all copy and paste. Easy. Done. Bottom line, "it's always the people," said Mendietta. "Always."

Barricklow continued, "The security risk that AI poses to a large enterprise is different than what it poses to a small business or an individual consultant. It goes back to the overall security posture of that organization. Where larger enterprises typically have their workstations locked down….Users aren't administrators of their computers. They are just normal users. So, if they start to invite or download applications, it's prohibited, or it's really restricted. Everything is protected because there is an air gap in between the user and that end resource in order to prevent prompt injections or to prevent disclosure of information. Their security landscape for an enterprise is much different."

In a smaller organization, obviously, the user is often the administrator, and they are wearing multiple hats all day, and a small business may rely on one person to be office manager, HR manager, receptionist, overall clerical person and more. Many are overwhelmed, and anything they can do to lessen their load will be welcomed, and they will more likely jump right in and start using tools like AI, without necessarily talking to their boss and asking for permission.

Villegas had additional comments on the security risks of AI. He stated on my podcast that the largest risks in adopting AI tools happen when they are not properly controlled or implemented. He said that there are five major risks, based on quotes from the podcast and a post-podcast discussion with him.

• "The first is data leakage. That's an instant compliance and incidence response scenario

• Shadow AI. When employees use unauthorized AI tools for their work—personal accounts, free aps, consumer-grade platforms—your data leaves your controlled environment with no audit trail and no way to know what was shared or retained. You lose control quietly, and often permanently.

• Permissions and over-exposure. Tools like co-pilot will respect access. But if you're using SharePoint or Drives, permissions that have poorly configured permissions means AI tools can inadvertently surface files users weren't supposed to see—or share them externally.

• Hallucinations mean errors. For example, AI-generated errors in claims analysis or compliance guidance can have direct financial and legal consequences. The court case section touches this later, but the two discussions are not connected.

• Third party integration risk. AI itself may be fine, but the connectors, the APIs, the plug-ins, might be misconfigured, or not properly patched or current, and that creates a risk in AI."

The use of free software versus paid software

Barricklow, Mendietta, Villegas and Aditi Group all mentioned the importance of using Enterprise Software rather than free software. As Barricklow said in the podcast, "if it's free, you are the product." Aditi Group has said similar things in podcasts in the past. "The software axiom that if it's free, YOU are the product," said Flittner. "It's been too often true for decades. We've been training Google's language models for decades with Gmail, Android and Google Search. When software companies take massive investor dollars, they need to either get you to pay for the program or use you to develop it." Flittner continued, "Exceptions can be open-source projects, which are often alternatives to Big Tech paid platforms, and often developed specifically to give better personal privacy protections. If it's low development budget, you're probably not the product. OpenAI was originally intended as a non-profit. Following our rush to use, it was converted in 2025 to a for-profit company with massive investor infusions, and now it's valued at over $850 Billion."

Privacy & artificial intelligence

Some of the common uses of AI today, as mentioned above, include meeting assistants, notetaking during meetings, etc. Let's say you are on a Zoom meeting with one of your group clients and you're talking about private issues—issues that are protected under HIPAA privacy laws and others. If you're letting Zoom record your meetings and perform notetaking and summaries of those meetings, you may be violating one or more privacy laws. I think it would be a bit difficult to get a Business Associates Agreement signed by your AI assistant! So, is simply using a tool like Zoom's AI Companion and doing these mundane tasks putting you at risk for HIPAA violations? I think the idea of privacy when it refers to the use of AI is a complicated one, at best.

On my podcast, Mendietta had this to say on the matter: "I think that AI may change the definition of privacy, and the way we look at it. Today, private means I have 100% control and AI might change that, because you're giving some of that control to AI."

The "Big Three" AI tools – ChatGPT, Google Gemini and Microsoft Co-Pilot

Although there are many other AI tools, the most commonly used are ChatGPT, Google Gemini and Microsoft Co-Pilot. I call these the "Big Three." They are literally everywhere. You can't do an online search or do any basic function on your computer without one or more of them staring you in the face and offering to help you with your tasks. I would think that most consumers use ChatGPT or Google Gemini, due to the cost and the ease for individuals to sign up and use. So, what are the differences between these, and what are each best suited for? I asked the experts on my podcast to get the answers. Mendietta gave me the best "short version" of the differences in the Big Three and what they are best used for.

"ChatGPT is my thinker," said Mendietta. "The one that's going to think with me. I'm going to prompt it, we're going to grow ideas, almost create [with me]. Gemini is my analyzer and my researcher. Co-Pilot is my doer. It integrates best with Microsoft products."

A detailed discussion on these differences can be found in the AI security podcast mentioned above.

"The *lure of a tool* that can *do so much* for us *makes* many *people turn blind eyes* to *security and privacy.*"

-Ted Flittner

Important information for Gmail users

It's important that any users of Gmail for your email services—whether a personal account or a business account—understand that in January 2026, Google automatically opted all U.S. Gmail users into Gemini AI features, giving it access to search, summarize and interact with your email content, which could be a serious privacy issue in our industry. They have offered an opt-out, but it's a several step process that I suggest every Gmail user take to be sure they've opted out. If you're not sure, it's likely you're still opted in, and you should contact your IT professional immediately to go into your settings and change them to be sure you're opted out. If your IT professional isn't aware of this, perhaps it's time to find a new IT professional! Larger companies generally do this on a company-wide basis; it's not left up to the individual user. But smaller companies and individuals may do their own IT work.

I asked Flittner about what he suggests to clients about the Gmail opt-out feature. "Google users, and that's most people, should seriously consider opting out of Gemini or at least opting out of having your data used to train it. This includes Google Docs, Drive, Google Photos, Gemini in Chrome browser, Gemini on your Android phone. The fact that Gemini features are 'on' by default and that opting out is a little convoluted hints at how valuable our data is to Google."

Flittner continued, "Workspace paid customers have more options—since you're already paying. Like allowing company data to train only company's Gemini agents. So, you can train it on specific data types to improve usefulness for your company, without leaking out to the public. At the other end of the spectrum, Workspace admins can disable Gemini features for all users across the board."

So, what do you do to opt out? Flittner stated, "If you want to opt out, just make a quick web search for the specific steps. Google makes others show us how to turn off Gemini. They give us articles like "Manage & delete your Gemini Apps activity."

AI policies

It's important that companies create very specific policies on the use of AI tools for their employees. What can they use, and what can they not use. If using AI, the company should purchase the software and train employees on their AI policies. I asked Aditi Group why it's important for companies to address AI in their policies and discuss with and train their employees the dangers of using AI without the proper security measures in place. Flittner asked, "Would you want your staff to just pull in anyone they find on the street and put them to work, giving them access to customer and company data? The fact is, your employees are experimenting with any number of different AI offerings for their personal lives. If you don't have policies and training, then someone is going to be using AI without your knowledge and could be putting your company in danger."

Barricklow also mentioned the importance of policies and training in the BER podcast S7 E8.

AI use in claims and claims auditing

AI use is now very commonly used in the claims and claims auditing arena. Since it's still relatively new, and many insurance companies and TPAs are adopting these tools without a lot of research, I wanted to summarize some of the important considerations for them when finding an AI vendor.

Using AI in claims processing and claims auditing is becoming very popular. Soon it will likely be the norm. What, in general, are the advantages of using AI in the claims auditing phase? To assist me with this part of the article, I asked Angel Onuoha, the CEO of Avelis Health, an AI claims auditing firm, to provide some important information for our readers.

"AI is best at speed and consistency," stated Onuoha. "It can scan huge volumes of claims and flag the same error patterns every time (overpayments, coding inconsistencies, missing documentation and outlier pricing), without getting tired." He continued, "The real advantage is prioritization: AI can audit everything and help you focus humans on the small slice of claims that actually has recovery value."

Onuoha's company uses a Large Language Model AI program (LLM). I asked him what, in general, can LLMs do to enhance claims processing or auditing? "LLMs turn messy inputs, EOBs, itemized bills, clinical notes, contracts, plan docs, into structured facts you can audit against. They extract, normalize and explain," stated Onuoha. "They're also strong at translating audit findings into defensible communication: clear rationale, citation to the relevant language, and a clean request for records or corrections."

TPAs, insurance companies, etc. must be concerned with AI Security. I asked Onuoha what should a TPA or Insurance Company look at when selecting an AI claims or claims auditing vendor to assure that their raw claims data is not being used to further train the AI Vendor's Model? How can they be assured that the claims file isn't being shared across the internet? Onuoha graciously provided this valuable information: "Ask for a written guarantee that your data is not used to train any foundation model. It should be in the MSA/BAA, not a marketing page." He continued, "Look for clear technical controls: encryption in transit and at rest, least-privilege access, audit logs and strong vendor policies on data retention and deletion. If they use third-party model APIs, confirm they're using enterprise/privacy modes where prompts aren't retained or used for training, and that data stays in approved regions."

AI claims auditing firms like Avelis Health can receive from the TPA or carrier many types of claims files, CSV files or even raw claims data files. Because I deal with a lot of TPAs and I specialize in privacy and security, I asked Onuoha how these TPAs or carriers can be assured that data from TPA A is being protected and not shared with TPA B and TPA C, etc.? "This is a core architecture question: tenants must be logically isolated," Onuoha responded. "You want true tenant separation at the storage layer and strict access controls, so one client's data can't ever appear in another client's environment."

I next asked Onuoha what types of firewalls can be put in place, and what should TPAs or carriers be looking for to be sure their data is secure and staying within the walls of the AI Claims vendor? "From a buyer's standpoint," Onuoha stated, "the easiest proof is documentation: security overview, pen test summary, and a clear diagram showing where PHI flows and where it does not."

Given laws like HIPAA Privacy & Security and various other laws, including CA state laws such as CCPA/CPRA and CHIP-OC, I asked Onuoha what TPAs and carriers should be looking for to assure the AI vendor they are using is compliant with all of the applicable state and federal privacy & security laws? "For HIPAA, they should sign a BAA, have documented Security Rule controls, and be able to show risk assessments, incident response procedures, and access auditing," Onuoha shared. "For California, you want clarity on data handling and subcontractors: who touches the data, where it's processed, retention timelines and breach notification processes. If they claim compliance, ask for evidence: SOC 2 (or a roadmap with dates), penetration testing and written policies."

Onuoha, by the way, joined Barricklow and I on the AI in Benefits Panel at the CAHIP-OC Annual Symposium's 9am session on March 10 in Lake Forest. That session was standing room only and had the highest session rankings in many years for the event. Attendees called it the "best session of the day," and some said the "best session they'd been to in many years." A few said it was the "best session they'd ever attended at a CAHIP-OC event." So, obviously, these guys know what they are talking about and "wowed" the audience.

AI in the courts

There have been several court cases in the past two years that have identified the true risks of using AI in law research. On more than one occasion, the AI tool couldn't find case law that supported what its user wanted, so it simply hallucinated responses for their brief filings, which looked very real to the attorneys or staff of attorneys, and no one bothered to verify the information. Besides being horribly embarrassed, the law firms are now in serious trouble for using AI and are updating AI policies to include a "cite checking policy."

Summary

AI is everywhere. We can't ignore it or push it aside for a later day to think about it. Hopefully this article has enlightened you a bit when it comes to AI security.

I asked Flittner if he had any final comments about AI and security and he responded: "AI is the wild west right now. Developers are racing to get products into the marketplace. Users are giddy with the hype around the possibilities. Everyone needs to take time to read the details and demand proper security. Many AI products have been shown to send private data to companies and countries that we wouldn't want. Others store and transmit data insecurely. Remember to treat these tools like you would people. Don't just hire a stranger and hand over the keys to the kingdom or trust them implicitly. Do your due diligence and keep people in the process."

As you finish reading this article, think back to Grandma Peggy and think about how frightened she would be if something like that happened, and she thought you were the victim of a crime and were in danger, not knowing that it was her that was the victim, or potential victim, of a would-be crime. As the caring grandchild of Peggy, you could sit down with her and tell her about what's happening in the world today with AI, and how bad actors can easily create deep fakes and copy your voice, face and other traits, to scare her and try to steal her money. Law enforcement suggests each family have codes and key questions and answers only your family members know of, so that in the event something like this happens in your family, you will know how to react and verify.

Also, be sure everyone knows to never pay anything unless you are able to have a live conversation with the "victim" and be sure to ask them the question or use the key phrases you've practiced at home.

As far as the health insurance industry goes, we are required to comply with privacy and and security laws. We are mandated to take all necessary steps to ensure our PHI and PII are protected. If you're using AI without the proper security steps in place, you are in violation of the law It's as simple as that. And your company could be largely at risk.

Dorothy M. Cociu, RHU, REBC, GBA, RPA, *is a nationally recognized expert in employee benefits with more than 30 years of experience in self-insurance, compliance, and health plan administration. She is the author of The ABC's of HIPAA Compliance and a leading authority on HIPAA Privacy and Security, ARRA, HITECH, and the Affordable Care Act. Dorothy has served as a consultant, expert witness, and independent fiduciary for employers, third-party administrators, and the U.S. Department of Labor. A longtime instructor for the Certified Employee Benefit Specialist (CEBS) Program, she is also a nationally sought-after speaker who provides education and training on healthcare legislation, compliance, and self-funded health plans.*

dmcociu@advancedbenefitconsulting.com

Author's note:

Special thanks to the individuals and firms featured in this article for their support, including
Eric Barricklow (ebarricklow@stellarcybersolutions.com)
Adriana Mendietta (adriana@mendieta.net)
Miguel (Mike) Villegas (villegasmo@isecureprivacy.com);
Angel Onuoha (angel@avelishealth.com).
Aditi Group can be reached at info@aditigroup.com.