KnowBe4

10 Questions Every CISO Should Ask About AI-Powered Human Risk Management Tools

KnowBe4, Inc. | 33 N Garden Ave, Suite 1200, Clearwater, FL 33755
Tel: 855-KNOWBE4 (566-9234) | Email: Sales@KnowBe4.com

0625US

INTRODUCTION

Talk of artificial intelligence (AI) has never been more common. Human risk management (HRM) and security awareness vendors of all sorts are leveraging generative AI powered by large language models (LLMs) to enhance and broaden their platform capabilities and product offerings.

Put simply, everything seems to have the letters “AI” (or even “agentic AI”) attached to it now. But how do you differentiate between too good to be true and what would actually help you get your job done?

The variety of features can be dizzying, so we’re here to help. Here are some important questions to ask regarding AI-based functionality when considering your HRM or security awareness initiatives.
Read on for sample questions to ask and key phrases and concepts to focus on.

1 Do We Really Need AI to Solve Our Problem? If So, Why?

This is the first and most basic question to ask of any vendor touting an AI-based platform or feature set. Ideally, you would have a use case already in mind. This would help you decide what to look for in an answer that would mean the most for your organization and responsibilities.

Key points to look for in a vendor's reply should focus on benefits, not just features. Instead of describing simply what their platform can do, get them to answer how these capabilities will help you. A strong understanding of what you and your organization are trying to solve with even considering an AI-based approach will help you guide a conversation with a vendor. Some example benefits to look for include:

- **Improved employee engagement through dynamic, personalized training recommended based on user profiles**
- **Provide immediate protection against emerging threats while delivering contextual coaching to your users**
- **Time savings through automation and agentic execution of tasks**
- **Clear insights on program successes and opportunities through AI-powered data and trend analysis**

Look for responses that go beyond the hype and show a deep understanding of both AI's capabilities and the specific challenges in HRM. The vendor should be able to articulate concrete ways in which AI addresses real problems, enhances effectiveness and provides value that traditional methods cannot. You need the vendor to convince you whether the AI component is a genuine asset or just a marketing gimmick.

2 Can You Explain How Your Advertised Machine Learning/AI Works?

Any vendor should be willing to talk about what's under the hood of their tool or platform. Consider researching some basic AI concepts and principles on your own to familiarize yourself with some of the terminology. Some topics to address include:

- **How data is input and processed**
- **How the model has been trained (more on that later)**
- **How the AI adapts and evolves**
- **How AI model performance is measured and validated**

Look for responses that demonstrate deep expertise in both AI and cybersecurity, address potential challenges and limitations, and show how the AI is specifically tailored for security awareness applications. Vendor transparency is key here so you can assess the suitability, reliability and potential risks of implementing their AI-based tool.

3 Do You Have Domain Expertise in Our Industry?

Be prepared to ask the vendor how they have helped others in your specific industry. A variety of cybersecurity risks are common to many fields, but don't settle for one-size-fits-all answers.

For example, healthcare institutions will have different requirements for training and data security than those in the finance sector. Get as specific as you can with the vendor and consider asking for case studies or customer success stories featuring other organizations in your industry.

The concept of AI in the HRM space may be relatively new, but a vendor's ability to address many different risks across multiple industries will ideally translate into their AI-based platform or tool. Consider inquiring about the backgrounds of any key experts leading AI initiatives your vendor mentions.

4 Who Is Using This Feature/Product? Where Are They Using It And How Are They Using It?

AI as a broad concept has had the last few years to pass the "innovators" phase of the technology adoption life cycle. It's starting to flirt with "early adopters" and even the "early majority." For technology vendors touting AI-based platforms and tools to be successful, they likely would have already had to have existing happy customers that have found real value in what they're selling. This means it's in your best interest to ask about them.

Case studies, customer success stories, peer reviews, etc. that underscore the value of the AI features you're considering is critical. Look for responses that go beyond generic claims and offer specific, verifiable examples of how the tool is being used successfully in various contexts. This information will help you assess whether the tool is a good fit for your organization, understand its potential impact and identify best practices for implementation. Vague or evasive answers to these points would be a red flag, suggesting either limited real-world usage or a lack of transparency.

5 What Is Your Delivery Model? On-Prem? The Cloud? On Your Server? Distributed?

The answer to this question should be to-the-point but also comprehensive. Look for a clear and concise view of how the solution is architected, deployed, secured and maintained. Especially when it comes to an AI-powered tool, this information is vital for assessing the fit with your infrastructure, the potential risks and benefits, and the long-term viability of the solution in your environment.

Watch out for vague or evasive answers to these points. These red flags could suggest either a lack of maturity in the product or a lack of transparency from the vendor.

6 Who Is Training (or Who Trained) the AI Tool/Platform? Who Will Continue to Train?

Don't be afraid to dig a little deeper with this question. After all, if their AI functionality is one of the vendor's primary selling points, they should be able to provide concrete details on how their AI instance is trained.

Answers to this question should demonstrate a thorough, ongoing and ethical approach to AI training. The vendor should be able to clearly articulate not just who initially trained the AI, but how it continues to evolve and improve over time. What they tell you is important for assessing the long-term viability and effectiveness of the AI-based tool in our ever-changing security landscape.

7 Are There Any Humans in the Loop?

While this might seem like a simple yes or no question, it can be a deal breaker. In many AI systems, particularly those with higher autonomy or agentic capabilities, having humans in the loop is crucial for oversight, decision validation and ethical considerations.

However, in some advanced AI applications, especially those designed for high levels of independence, human intervention may be minimal or not needed during operation, though oversight remains important in areas like accountability and safety. As these systems mature, the focus shifts from managing the platform to reviewing results, allowing humans to concentrate on analyzing outcomes and strategic decision-making rather than ongoing operational supervision.

The vendor should be able to discuss specific, meaningful ways in which humans are involved, their qualifications, and how this human element contributes to the overall effectiveness and trustworthiness of the solution. An AI-based HRM platform needs to be trusted to be brought into your organization, so a human-in-the-loop approach should be crucial.

8 Can the System be Audited?

This line of questioning flows from the importance of maintaining real, live humans as part of the AI development process. Regular validation tests on the vendor's side need to be done to ensure any generative AI capabilities are as accurate as they can be.

Look for a combination of technical measures (like audit trails and explainable AI), procedural safeguards (like third-party audits and continuous monitoring), and transparency in methodologies. Ask how the vendor ensures the integrity and reliability of their AI system, and how they empower clients to verify this for themselves. Sound auditability and validation procedures need to be part of any AI-based tool that will play a role in your organization's HRM program.

9 Who Is Liable If Your Solutions Fail? What Happens If This Backfires?

No process, platform or tool is perfect. Potential failures and missteps are a part of life, especially with emerging technologies like AI. Don't shy away from discussing failure scenarios with your vendor in objective terms.

Look for transparency about liability terms, robust risk management practices and a partnership approach to handling potential failures. The vendor should be able to explain not just what happens if things go wrong, but how they work to prevent failures and mitigate impacts. This information is crucial for assessing the overall risk of implementing the AI-based tool in your HRM program and for ensuring that they have appropriate protections in place.

10 Where Is Our Data Stored? How Secure Is It?

Given your role as CISO, this line of questioning should be a no brainer. Don't be afraid to ask for specifics and get technical as you need to get. After all, the point of any HRM platform or program is to ensure a strong cybersecurity culture. Such a tool lacking the proper data security protocols should be a non-starter.

Look for specifics on physical and logical security measures, compliance with relevant standards and a clear understanding of the unique security challenges posed by AI-based systems. The vendor should be able to articulate not just where any data of yours is stored, but how it's protected at every stage of its lifecycle. You need to be able to trust that the user data that would go along with providing an HRM platform is secured in a way that aligns with your own data protection policies and regulatory requirements.

Conclusion

Technology vendors will bend over backward to convince you their approach to AI is the best one. In this way, the AI revolution we are in the middle of is not dissimilar from the advent of other major technological advances in our recent history. The same concepts that served you well then will serve you will now:

Trust but verify

If it seems too good to be true, it probably is

Ultimately, they're trying to sell you something

Keep this advice and these questions in mind as you seek out a vendor or vendors to support you in your journey toward a strong security culture.

About KnowBe4

KnowBe4 empowers workforces to make smarter security decisions every day. Trusted by over 70,000 organizations worldwide, KnowBe4 helps to strengthen security culture and manage human risk.

KnowBe4 offers a comprehensive AI-driven "best-of-suite" platform for Human Risk Management, creating an adaptive defense layer that fortifies user behavior against the latest cybersecurity threats. The HRM+ platform includes modules for awareness and compliance training, cloud email security, real-time coaching, crowdsourced anti-phishing, AI Defense Agents, and more.

As the only global security platform of its kind, KnowBe4 utilizes personalized and relevant cybersecurity protection content, tools and techniques to mobilize workforces to transform from the largest attack surface to an organization's biggest asset.

Free Phishing Security Test
Find out what percentage of your employees are Phish-prone with your free Phishing Security Test

Free Automated Security Awareness Program
Create a customized Security Awareness Program for your organization

Free Phish Alert Button
Your employees now have a safe way to report phishing attacks with one click

Free Email Exposure Check
Find out which of your users emails are exposed before the bad guys do

Free Domain Spoof Test
Find out if hackers can spoof an email address of your own domain

KnowBe4

KnowBe4, Inc. | 33 N Garden Ave, Suite 1200, Clearwater, FL 33755
855-KNOWBE4 (566-9234) | www.KnowBe4.com | Sales@KnowBe4.com