CAMPUS TECHNOLOGY

Innovation in Education

The Dark and Light Side of Artificial Intelligence

AI enhances not only the strength and effectiveness of cyber attackers today, but also the tools and methods used to defend against those attacks. As AI's potential for good and bad rapidly evolves, higher education weighs AI's risks and benefits while reshaping and prioritizing cybersecurity strategies to mitigate today's threats.

AI's Security Implications Come Into Focus

Whether powering the latest tools and technology, or transforming an already complex threat landscape, AI increasingly shapes institutions' security postures.

ARTIFICIAL INTELLIGENCE continues to transform every industry and product it touches, and higher education is not immune to the lure and potential AI brings to the table, whether good or bad.

AI-enhanced capabilities within the latest software and other technology tools remove the need for manual, more time-consuming processes or data entry, especially in IT and cybersecurity administration. Yet AI also enables Black Hats to rapidly deliver more pointed malware and phishing e-mails, for example, to an already at-risk group of targets. Fears of student cheating and plagiarism notwithstanding, the potential of AI to do more harm than good within higher education lies well beyond student and recruitment marketing use cases.

"There are effective and constructive implementations of AI, but it can also be used to harness large numbers of devices together, automating attacks and executing against an extensive list of potential targets," said Bill Harrod, federal CTO at Ivanti. Just as the COVID virus continues to mutate and evolve, Harrod believes that AI can similarly power the evolution of a cybersecurity attack.

"It can 'learn' what's effective and what's not, and it can get better," Harrod said. "And there's a real concern that an AI-generated attack could run amok, that the learning algorithm could spawn additional attacks and that get out of control without any sort of kill switch or kill chain, taking on the persona of a wildfire and either burning itself out or succumbing to some sort of control that is finally imposed. We have the challenge of being able to harness AI to improve detection, make it more efficient and effective, but at the same time not allow it to devolve into something that is going to create a more difficult solution to access."

In May 2023, a group of universities and researchers announced the formation of ACTION – the AI Institute for Agent-based Cyber Threat Intelligence and Operation – led by the University of California at Santa Barbara to leverage $20 million in funding from the **National Science Foundation** over the next five years to understand how AI might detect and respond to cybersecurity breaches at scale.

"We've seen tremendous advances in AI over the last few years. There are ways that adversaries trying to attack computer systems can take advantage of that, and there's a lot of fear about using AI tools to automate and scale attacks that used to take a lot of manual effort," ACTION co-principal investigator and University of Virginia Computer Science Professor David Evans **recently told *Government Technology***. "But there are also a lot of opportunities and interest in using these advances to defend systems better and build systems that are resilient to attack."

For now, as campus IT and cybersecurity leaders take stock of their resources and assess risks, AI tops both lists, ready or not. A recent survey of ***Campus Technology*** readers

AI/ML capabilities institutions plan to adopt or roll out within the next 3 years

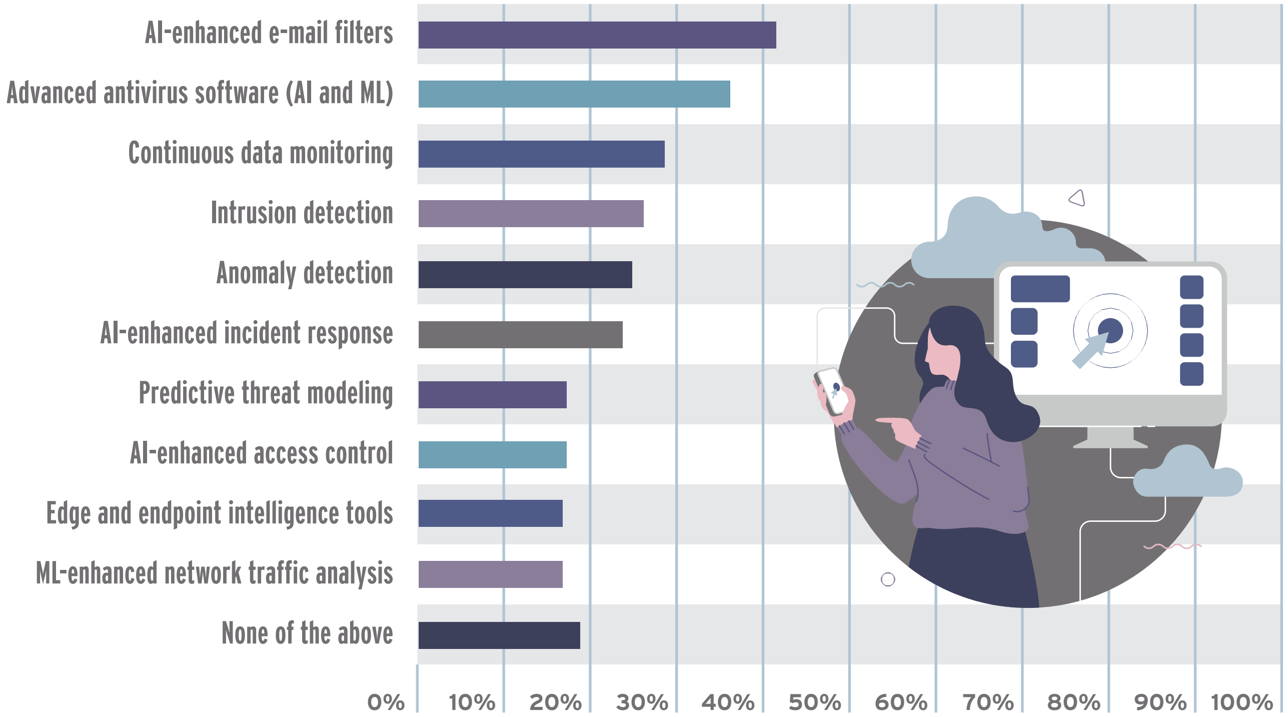

reveals how far along some campuses are when it comes to preparing for or rolling out AI-enhanced security tools and technologies. Respondents' levels of adoption varied, but when asked whether their institutions have already undertaken steps to prepare for advanced AI-powered or -enhanced cybersecurity threats (e.g. increased scale, speed, and complexity of attacks through the use of generative AI or automated malware), nearly 60% said yes.

Where We Are Versus Where We Want to Be

Automation (26%), zero trust (25%), and cloud-native security (23%) are the top cybersecurity frameworks, approaches, or strategies respondents said their institutions have adopted or started to adopt within the last year or more, according to the same survey. Looking ahead at frameworks, approaches, or strategies their institutions plan to adopt within the next year or two, AI-enhanced threat detection (28%) gained more responses.

What Lies Beneath

Respondents' greatest concerns around AI and cybersecurity risks are focused on generative AI-enhanced malware and ransomware (42%), impersonations or deep fakes (41%), and AI data privacy risks (38%).

Those results mirror results **reported by Pew Research Center**, which found that Americans' concerns around AI, by and large, focus heavily on maintaining control of AI, along with data privacy and safety, with 53% of respondents concerned that their information is not being kept safe or private.

Top 3 concerns around AI and cybersecurity risks (from most to least concerned)

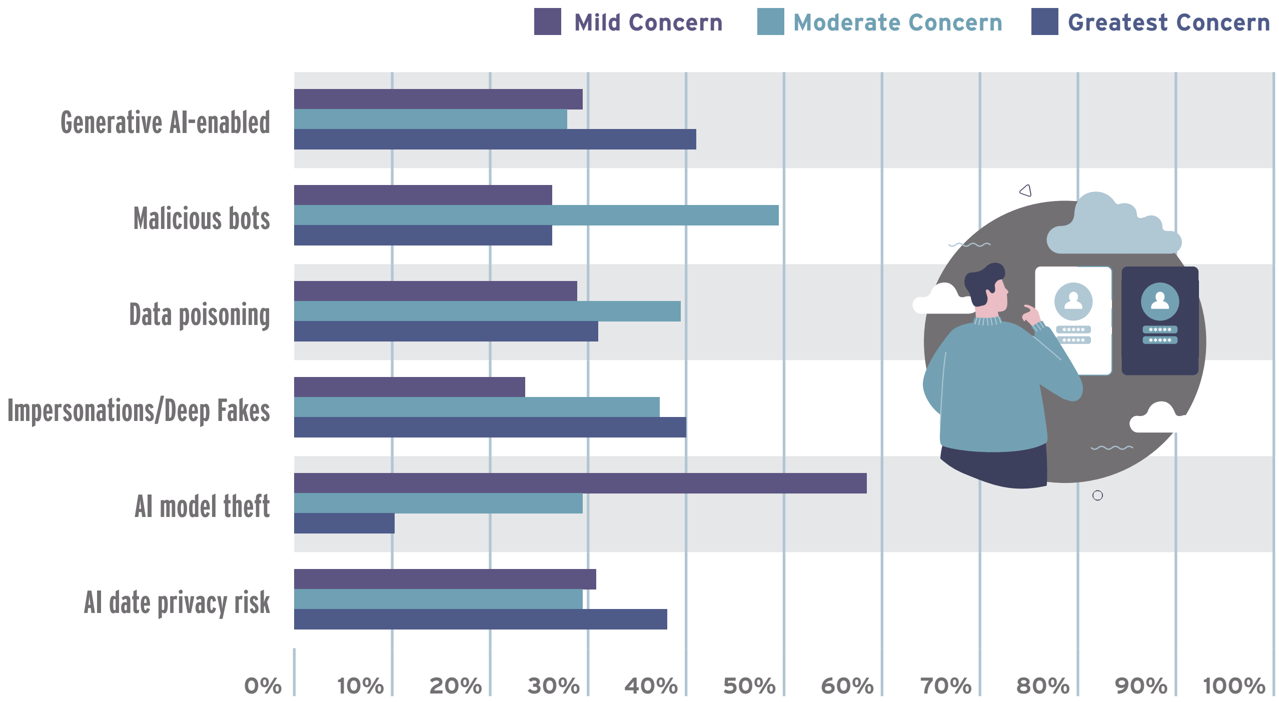

Cybersecurity Leaders Ready for Disruption

Artificial intelligence will power broad evolutions in cybersecurity, not just in risk assessment but also in advancing zero-trust strategies and boosting UX.

AS MORE MAINSTREAM or consumer-accepted use cases for artificial intelligence and machine learning come sharply into focus, black hats continue to gain new potential to deliver novel threats through leveraging AI.

For all of the good that AI and ML allow higher education to accomplish, the number and reach of cybersecurity risks grows just as rapidly alongside those capabilities.

Indeed, "the rise of artificial intelligence is a double-edged sword for CISOs," **Katell Thielemann**, VP distinguished analyst at Gartner, acknowledged during the opening keynote for the 2022 Gartner Security & Risk Management Summit. "Enterprises are facing a deluge of automated cyber attacks, which are exponentially rising in velocity, variety, and complexity. However, AI is simultaneously supporting security teams in detecting and responding to threats, fundamentally changing organizations' defense paradigms."

Risks come in many forms, whether through the rapid rise of endpoints and blind spots within an institution's network, slow or otherwise delayed patch management, deep fake phishing scams created through generative AI tools, overworked and under-resourced security teams – the list goes on. Colleges and universities must protect their infrastructure, data, and constituents through agile cybersecurity strategies and cutting-edge technologies that can move quickly to thwart potential attacks.

"Attackers are weaponizing AI just as fast as organizations augment their defenses with it, meaning that it's not enough for cybersecurity technologies to evolve – strategy and leadership approaches must change, too," Gartner Distinguished VP Analyst Andrew Walls said.

> Zero-trust frameworks enhanced by AI-backed threat reporting and analysis represent one such strategic shift, allowing institutions to keep pace with the scale and complexity of AI-powered cyber threats.

Zero-trust frameworks enhanced by AI-backed threat reporting and analysis represent one such strategic shift, allowing institutions to keep pace with the scale and complexity of AI-powered cyber threats. New cybersecurity tools developed specifically to enhance zero-trust approaches over traditional VPNs, for example, deliver real-time insight into global threats, detect anomalies, and ease end-user experience so that users don't seek out less secure workarounds. These tools also power edge computing and ensure threats are detected wherever users are working or data is collected, isolating risks and tamping them down before they grow too powerful to control.

Keeping Track of It All

"When I talk with CIOs and CISOs, particularly in higher ed, the top three cybersecurity concerns that they focus on are, first of all, ransomware and protecting their students' and employees' data," said Bill Harrod, federal CTO at Ivanti. "Second, they're concerned with how to manage and protect the devices they issue as well as any other devices that connect to their networks. Third is how are they managing access to the networks and applications they're responsible for."

Ivanti's Neurons platform includes a cloud-native patch management solution, enhanced by AI, which helps teams prioritize vulnerabilities, as well as an industry-first AI tool that enhances zero trust frameworks.

"Zero trust has become the buzzword around cybersecurity," Harrod says, "and a lot of what zero trust really boils down to is good cyber hygiene: knowing what's on the device, who's on your network and what they have access to. Part of the zero trust model is defining very fine-grained zones of trust and strong authentication that controls who has access to applications and data on the network and how they move from one zone of trust to another."

How to Get There

As AI-powered cyber threats continue to propagate at exponential speed and scale, higher ed CISOs and security teams must look beyond the security tools they've deployed to successfully manage risks proactively. Gartner advises security leaders to prepare for this next evolution of cybersecurity by embracing practices of continuous foresight, recognizing weaknesses, and accepting that the attack is never over.

As security budgets grow, institutional leaders "will expect a highly strategic approach to security investment that results in demonstrable returns typified by fewer breaches and greater enterprise resilience," **Gartner stated**. "Pursuing multiple models of the future will enable security leaders to build an investment strategy that is flexible enough to respond to new threats with agility."

"The most **effective CISOs** don't try to do it all," Gartner's Thielemann said. "Play to your strengths as a leader, and then augment your teams with those who complement your weaknesses."

Time to Take the Leap

Higher ed institutions on the fence about AI-powered cybersecurity must consider what could happen if they fail to embrace the AI tsunami.

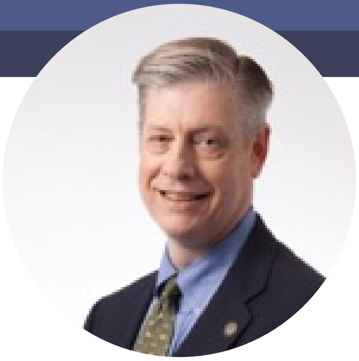

BILL HARROD, CISSP
Public Sector CTO, Ivanti

CAMPUS TECHNOLOGY recently spoke with **Bill Harrod**, Public Sector CTO at Ivanti, about the duality of artificial intelligence in cybersecurity, and how higher education security teams can harness the potential that AI-powered tools and resources bring to bear.

What specific capabilities or power do modern, AI-enhanced cybersecurity tools deliver that weren't available perhaps five years ago, or even two years ago?

Artificial intelligence has become the latest buzzword in emerging technology, and there's certainly a lot of concerns around what we call "generative AI" – particularly around things like privacy and confidentiality, as well as the validity of generative AI. The idea of deep fakes and information being developed or perpetrated by generative AI, and it being erroneously or purposefully, maliciously inaccurate is certainly a concern. But on the other hand, we use AI and machine learning to power our automation engines. It allows us to be able to evaluate information, particularly indicators of network interruptions or preemptive equipment faults or other service degradations, and then use that hyper automation and some predictive analytics to diagnose and even remediate issues before any human or manual intervention is required. This really is what makes AI such a powerful two-edged sword.

Five years ago, automation was mostly accomplished by writing and executing a script. Was it faster than a manual process? Sure, but it was not adaptive. Today, machine learning and AI allow routines and playbooks to be dynamic and updated in near real time.

The attack surface for an institution has changed dramatically. We now rely more heavily on cloud-based applications. There are more IoT devices than there were five years ago. Smartphones, tablets, and personally owned devices make up a much larger portion of the compute power and the endpoint devices on a network. All of this comes down to a radically different risk profile for most organizations.

The pandemic allowed us to change where we worked, and that has become much less geographic-centric or geographic-specific. One of our slogans at Ivanti is that that we help "make the everywhere workplace possible," and that really is about allowing users to bring the technology that they know and like, and making it secure and usable on an institution's network, or from a remote location, by more users.

> One of our slogans at Ivanti is that that we help "make the everywhere workplace possible," and that really is about allowing users to bring the technology that they know and like, and making it secure and usable on an institution's network, or from a remote location, by more users.
>
> – BILL HARROD, IVANTI

It also opens a larger, more sophisticated attack surface and more attack vectors. Institutions need a strong, comprehensive asset inventory. They need to know what devices should have access to the network, and then automate the discovery of those endpoints, whether IoT or smart devices. And all of those devices should be included in a comprehensive asset inventory and a way of managing network devices – an IT service management capability – and then understanding the risk profile for all of those devices. How do we leverage risk-based vulnerability management to control that risk across the enlarged enterprise?

What is the reality that IT and security teams face today and what else is available to make their work more efficient?

Higher ed institutions certainly have a challenge with constrained IT budgets and personnel resources, and it often comes down to that comprehensive inventory of assets: What do you need to protect, how should you protect it, and what are you protecting it from? They must know not only what the devices are but what the applications are that they're trying to protect, and the operating systems, and a way of prioritizing that vulnerability remediation.

Risk equates to vulnerability, times threat, times cost. If you can take any of those vulnerability or threat or cost variables and make it zero or nearly zero, then your risk declines significantly. One of the ways that Ivanti helps our clients is by continually tracking known vulnerabilities, and enriching that knowledge with intelligence around what is actively being exploited in the wild, what has been weaponized.

The highest-risk vulnerabilities are those that have remote code execution, when somebody outside of an organization triggers a vulnerability to execute code remotely and/or privileged escalation. Once they get something loaded onto the network, they can manage it externally, but also then change their privilege to "administrator" and have access to run additional code or move from one room to the next, from that zone of trust. That's really where risk-based vulnerability management and actionable threat intelligence come in. Some solutions are focused on just that: being able to create that funnel so that we can take a large number of vulnerabilities across a large number of applications and devices, and narrow it down to actionable intelligence that can be accomplished.

In terms of institutional IT security strategies, what needs to change or how should teams adapt to keep up with more rapidly evolving technologies and capabilities?

It's incredibly important to keep up with evolving technologies, to remain agile and not get locked into a legacy technology solution. Over the last five years we've seen an incredible move to the cloud. It is even more secure today with secure cloud computing environments, so it's important to be able to adapt to those evolving solutions that are moving to the cloud – but it's also important to do it in a way that provides the best end-user experience. People are going to use technology that they're comfortable with, that is similar to the technology they use for themselves on a personal basis, on a daily basis. If you're going to try and force people through a particular cumbersome way of accessing applications and data, then you lose the mass and people will either find a way around it or they'll just adopt a different solution or look for it someplace else.

Can Data Mind the Data?

In the data-rich world of higher ed, leveraging AI and automation to understand and manage data risks or recovery is par for the course. The next level: AI and ML tools that access and leverage that data for other purposes.

A NEW REPORT FROM FORRESTER, "**Bring-Your-Own-AI Hits the Enterprise**," reveals that more than a quarter of global IT decision-makers indicate that 51% to 75% of their employees will likely use generative AI technology by the end of 2024. Higher education security and IT leaders must consider similar potential within their own environments, where students, faculty, administrators, and other constituents stand in for "employees" in the enterprise, but the risks to data security are just as great, if not more critical given the nature of the data potentially exposed.

Another recent survey, "**Generative AI Through the Eyes of Gen Z**," notes that, of the 43% of respondents who indicated they have used generative AI tools to "help with schoolwork," 51% were college students.

While this somewhat-nascent toolset has endured a lot of hype that has yet to come to fruition, many campuses have taken a "wait-and-see" approach rather than fully embracing or preparing for AI's potential in the classroom or as part of institutional tasks. Yet when it comes to managing risk, there can be no waiting.

"Immature data governance, concerns about algorithmic bias, and ineffective data management and integration pose the greatest challenges to the implementation of AI in higher education," D. Christopher Brooks wrote in the June 2022 edition of ***EDUCAUSE Review***.

As higher education's both known and dark data stores continue to grow exponentially larger and faster – mountains of student data, and data culled from research, for starters – all of that data increasingly finds itself up for grabs by untested, unsanctioned, or just plain unknown AI tools. As a result, data management and backup and recovery solutions and security are more important than ever before.

Data Management and Security Challenges Exposed Anew

Data management and governance, security, and integrity continue to rankle higher education IT and cybersecurity teams – and have done so at least since the era of Big Data, if not longer. The open nature of higher education systems, in many cases, is the cause of delayed or otherwise incomplete data governance planning; however, new risks and potential for solutions and tools that leverage AI and ML capabilities have renewed urgency for effective and complete data management.

"It comes down to a risk-benefit analysis conversation, really," said Christian Westervelt, senior manager, technical sales engineering for healthcare, public sector, and education at Veritas. "Does inviting AI to accomplish a task, whether it's pattern recognition, code development – does it make sense or does it open up a potential risk vector?

"As with any new technology, in many cases it's about getting back to the basics. Even with all of the capabilities that are available and accessible to our customers, I'm still surprised by the core conversation around recovery, recovery point, and recovery time. Those fundamentals have not changed in many, many years. What's changed is the 'how.' How can we leverage that technology?

Common Challenges to the Implementation of AI source: **EDUCAUSE Quick Polls**, June 2021

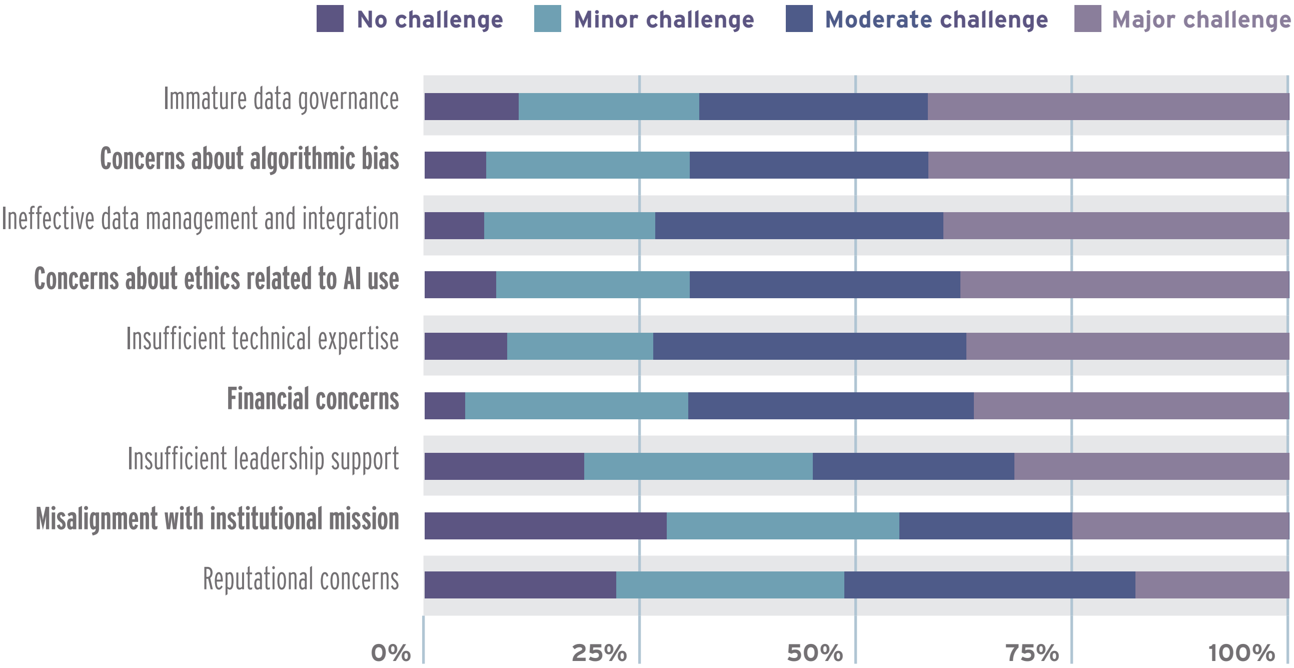

Make that determination but focus on the basics. Get the plan. Make sure it's documented. Make sure it's tested. That's the rock on which we can move forward," Westervelt advised.

Can AI Improve Resiliency?

An institution's ability to protect, recover and maintain its data no matter the event or disruption is the keystone of so-called data resiliency. A resilient environment prevents data losses, minimizes down time, and helps teams or institutions wholly recover from any such attack, event, or disruption. Today, "AI, data governance and the cloud are inextricably linked to data resilience," advised author Guy Pearce in the ***ISACA Journal***. "Recognizing that data are at the heart of each of these concepts means that the converged technologies need to be the focus of an organization's data resilience strategy."

Deep insight into what is happening in an environment at any given time is required for proper recovery, Westervelt said. "There's the dark data assessment process – half of the battle is that they don't know everything about their environment, about their users, about the activity that's taking place. We need to lift those blinders off, let them see or gain access to information about the data they already have, what they're storing or protecting, so that ultimately, should the need arise, they can recover from it."

When the rapid growth of new threat vectors converges with another trend – fewer resources to handle the IT, data management and security workloads – "and with a lot of people working from home, you have a lot less control of the data," said Vishal Kadakia, solution systems engineer at Veritas. "You don't always know where it resides."

Tools or solutions that leverage AI and automation can help fill the gap, and "that's one of the reasons why it's so prevalent now," Kadakia added.

> It comes down to a risk-benefit analysis conversation, really. Does inviting AI to accomplish a task, whether it's pattern recognition, code development – does it make sense or does it open up a potential risk vector?
>
> – CHRISTIAN WESTERVELT, VERITAS

"Separating the data and the infrastructure – because they're two different platforms – we can look at the data, the unstructured data, and understand where it resides, who has access to it, the last time it was looked at, whether it has any sensitive information. From an infrastructure standpoint, we need to know what it's doing: Has it been compromised? Are you protecting everything that you need to?" Kadakia said. "When you say, 'AI,' what does that mean for you? Is that an application? How are you leveraging it? It really invites further conversation, which is always a good thing."

Back to Basics: Data Protection and Recovery in the Era of AI

The fundamentals of data protection remain constant in the face of emerging AI-powered threats.

CAMPUS TECHNOLOGY RECENTLY spoke with **Christian Westervelt**, senior manager, technical sales engineering for healthcare, public sector, and education, and **Vishal Kadakia**, solution systems engineer, both at Veritas, about the critical importance of data protection in the current era of rapid AI evolution.

We talk a lot about the potential harm and risk of AI, but in many cases this technology can also help in security and data protection. What are we seeing as far as advances in AI and cybersecurity practices, especially in data management and protection?

Kadakia: We're trying to understand how we can use AI from a cybersecurity perspective against ransomware, and some of the things we can do in early detection. Can we leverage that to do any type of alerting? Can we provide a system that can give an early warning to say, "Hey, you may want to take a look at this. Your data may be compromised." We're seeing that, across our platform, – and it doesn't matter which portfolio we're talking about – leveraging AI and machine learning to help with those types of things, help teams where their data is at risk, where they might not even know data is at risk, with dark data assessment, with data management from start to end. Pattern recognition is key: leveraging our solution suite to detect anomalies, things that either shouldn't be occurring or stand out in the normal day-to-day data management.

There are things that we've always been doing. We have been capturing metadata. Can we use machine learning or AI to look at that metadata to give that early warning? Can we also use machine learning on user behavior patterns to discover if there are any anomalous user behaviors occurring outside of the norm? Reporting and intelligence on the metadata have been part of our core solution offerings for a long time. The newest aspect is machine learning, leveraging the benefit of pattern recognition. We've also branched out into looking at the user behaviors, the access they have – that's a different solution offering, but that's been done for many, many years.

How can teams move from a reactive state to a more proactive stance against new threats?

Westervelt: The biggest piece is a plan, and this is true not just for higher ed but for all customers. You have to have a plan. Very importantly, it has to be a documented plan and it has to be a tested plan. And that last piece is the critical element, where you may have a documented plan or an organization may have a plan, but they haven't tested it. The time of a crisis that is not the time to try to validate whether we've covered all of the bases. Make sure you've accounted for – as much as possible – inevitable scenarios: no access to network, to data, whatever it may be.

Kadakia: The other aspect of that is knowing where your data is. Understanding not just where but also what's in there. Is there any PII information? Is there any classified or very personal information that needs to be super protected? That's where we're seeing teams wanting to have the backups, wanting to have the immutable storage where no one can modify those data. I think a lot of higher education is understanding that we need to invest in our infrastructure and these plans.

What are the technologies or infrastructure components that these teams are relying on now for that protection?

Westervelt: It all comes down to the cloud, which leads to all sorts of different in-roads and conversations. How are they leveraging cloud infrastructure? How can they best leverage it, whether that's hosting mail, access, or other areas for data storage? And also importantly when it comes to recovery, what kind of value-add can be achieved through the cloud?

> Pattern recognition is key: leveraging our solution suite to detect anomalies, things that either shouldn't be occurring or stand out in the normal day-to-day data management.
>
> – VISHAL KADAKIA, VERITAS

Kadakia: A lot of teams look at isolated recovery environments, where they have a location that's kind of like a vault – it's locked down, but they know that what they're putting in there is clean information, therefore they know that it can't be compromised in any way. Indelible WORM storage, where it's write once, read many, that cannot be modified. Ironically, if you look in the protection realm, tape was a great medium because it was offline and it could not be modified. In many cases where organizations are leveraging disk-based protection, we now make offerings and options for indelible storage so that users can prove that content has not been altered or modified. That's a huge step forward and a huge part of the recovery process in the event of ransomware. Three or four years ago that was a nice-to-have, but today it's a given that we need immutable. Automated recovery and orchestration is also critical because teams don't necessarily have the resources to do all of that on a manual basis. Whether that's from an on-prem to cloud or bringing it back down, most teams need an automated method to do that. Also, in the past it would be, teams needed a file recovered or an application in the dataset recovered. Now they're requesting entire data centers and all of their apps at scale recovered.

What are the essential questions higher ed teams should be asking right now around these technologies and capabilities?

Westervelt: I really think it gets down to a risk benefit analysis conversation. Does inviting AI to accomplish a task – whether it's pattern recognition or code development – does it make sense or does it open up a potential risk vector?

It gets back to having that conversation. Our technical teams like to be seen as the trusted advisors for our customers. If there's a doubt or you're unsure, have that conversation and determine whether it makes sense for the organization. Is the risk too great by introducing an AI or third-party entity? Where's that being run? There are so many different potential upsides, but there are also some challenges there as well. Even with all of the new capabilities that are available and accessible to our customers, I'm still surprised by the core conversations we have around recovery, the recovery point and the recovery time. Those fundamentals have not changed in many, many years. It's just the "how" that's changing.

AI Demands Greater Focus on the Details

A well-managed environment is a secure environment — but given today's constantly shifting threat landscapes, how can teams ensure they're managing well?

IN THE GLOBAL FIGHT AGAINST cybercriminal activity, higher education finds itself consistently one of the major fronts. Adversaries and nation-states target college and universities primarily due to their lucrative data stores and wealth of research. Higher education's traditionally open environments and culture of sharing make the sector an even more attractive target.

While recent progress has been made to shore up data and address threats through automation and updated security frameworks, the sophistication and volume of security threats made possible today through advances in AI demand ongoing process evolution and proactive risk management. Shifting to zero trust or assume breach frameworks, penetration testing, and red teaming each bring merits that should be considered or adopted by higher education security and IT leaders. And while those tactics or approaches commonly top cybersecurity lists of best practices today, they don't negate the importance of tried-and-true tools like endpoint protection and response (EDR), for example, or performing due diligence when evaluating and purchasing third-party services and solutions.

"The adversaries are really targeting education and, especially, research areas," said Jeff Stewart, Field Chief Technology Officer at SolarWinds. "We have to make sure we're putting extra protections in those places that are targeted the most. Keep threat intel in mind, always, and understand what the large adversaries are doing."

Risks Versus Gains

The global average cost of a data breach in 2023 was $4.45 million, representing a 15% increase over the last three years, according to the latest data breach report from Ponemon Institute and IBM, "The Cost of a Data Breach Report 2023." The report also found that organizations that use security AI and automation extensively save on average $1.76 million compared to organizations that don't.

"Where we were five years ago is very different today from a security perspective," Stewart said. "The tooling and the sharing of data between the tools has helped a great deal. We stop a lot of threats automatically today that we would not necessarily have caught a few years ago. It's quietly protecting us from a lot more than what we would have seen before."

AI "is really a force multiplier, and knowledge multiplier," Stewart added. "Future-state, we will get a force multiplier and that will help some of the shortages we're seeing of highly skilled folks – and that's in a lot of different fields, security being one of them."

The list of offenses that AI enables adversaries to commit keeps growing. Generative AI today allows black hats to swiftly deliver flawless phishing e-mails or even deep fake images and videos that are more difficult to detect than ever. Combatting such threats requires alternative approaches that some higher education security teams may only just be starting, including threat modeling and red teaming.

AI also enables teams to define more detailed assumptions and threat models, and rapidly zero in on data that may be at particularly high risk, high-target assets, and other key vulnerabilities, all of which can serve as the starting point for protection – all the more critical as AI improves attack sophistication on the other side.

Generative AI today allows black hats to swiftly deliver flawless phishing e-mails or even deep fake images and videos that are more difficult to detect than ever. Combatting such threats requires alternative approaches that some higher education security teams may only just be starting, including threat modeling and red teaming.

"AI can help us with the assumptions," Stewart said, "and with making sure things are done correctly and appropriately. That's where the future may lead us: You may not need an expert to design a threat model; you may be able to use your AI to codify a good threat model, codify what we should be thinking about it. That would be our kind of Nirvana, where we ensure teams are working with the most modern, up-to-date knowledge."

Advanced Protection Strategies

Red teaming, where teams adopt an adversarial approach or position to evaluate, assess, and challenge IT systems and infrastructure, may be performed by a third party or internal groups. The practice encourages thinking around external perspectives and helps security teams fortify and prioritize protection of data and assets, and improve an institution's overall security posture. The practice differs from penetration testing, which tends to be performed within agreed parameters and times. Red teaming tends to take place over more extended timeframes and through less obvious means, simulating real attacks. Automated red teaming and penetration testing can also be performed through cloud-based software solutions.

> Where the most advanced AI-enabled tools as well as cybersecurity protections require reliance on cloud-based tools and solutions, assessing and navigating third-party risks has never been more important for higher education security and IT leaders.

Organizations "must invest in strengthening their cybersecurity infrastructure and training their staff to handle potential threats," advised the World Economic Forum in a June 2023 report, "**Cybersecurity and AI: The Challenges and Opportunities**." That includes "regular security audits, incident responses planning and promoting a security-first culture. Incorporating AI in cybersecurity strategies can also play a crucial role in identifying threats and improving response times."

AI developers, in turn, "have a unique responsibility to design systems that are robust and resilient against misuse," the Forum's report continued. "Techniques like differential privacy and federated learning can be used to protect data. At the same time, efforts like OpenAI's work on AI and cybersecurity research are vital to staying ahead of evolving threats."

Working with Hybrid Environments and Third-Party Providers

Where the most advanced AI-enabled tools as well as cybersecurity protections require reliance on cloud-based tools and solutions, assessing and navigating third-party risks has never been more important for higher education security and IT leaders.

"We don't know how much duct tape and chewing gum is holding it all together," Stewart said. "We may see an interface to a service, but you need to ask, 'What visibility will we get?'"

Within the public sector, the FedRAMP standard provides government customers with a great deal of visibility into how cloud-based services will operate. While higher education may not need to leverage FedRAMP-certified services, institutions can leverage lessons from FedRAMP to inform the list of questions they should ask any provider, Stewart advised, including:

- **Are you SOC 2 compliant?**
- **Do you have ISOs?**
- **Do you have controls in place?**
- **How do you store my data?**
- **Where is my data capped?**
- **How do you protect that data?**
- **What open source components do you rely on?**
- **How are they updated?**

"Really develop that checklist of what your third-party program looks like and tier it so that you know what you'll require, from a basic service that you're not sharing any of your data or access with up to other services that you're sharing all of your research data with, and which a nation-state would want. That's where I'm going to my elevated protection plan," Stewart said. "Make sure you understand and tier your protection based on the risks that a service presents to you."

Where Will AI Take Cybersecurity Next?

Evolving threats powered by AI spawn advanced thinking and approaches to cybersecurity, along with new skillsets and priorities.

JEFF STEWART
Field Chief Technology Officer, SolarWinds

CAMPUS TECHNOLOGY recently spoke with **Jeff Stewart**, Field Chief Technology Officer at SolarWinds, about managing risk in today's AI-fueled world, and where the technology could take cybersecurity next.

What are some of the more significant threats within the higher education environment today?

Stewart: When you look at higher ed, there's always been a challenge there. Just the nature of higher education lets you have uncontrolled students, uncontrolled machines, researchers who want to do whatever they want to do and very demanding around having that ability – it's very difficult to manage. Those are some of the biggest threats that we see in higher education. And the threat from nation-states, which are really focused a lot on research, on collecting intellectual property – those are some of the hardest threats to protect against.

> Higher education has a lot of requirements for flexibility, to allow people to work together, to allow data to be shared between multiple entities, and often security doesn't necessarily take priority as the top requirement over data sharing and research.
>
> – JEFF STEWART

Higher education also has a lot of requirements for flexibility, to allow people to work together, to allow data to be shared between multiple entities, and often security doesn't necessarily take priority as the top requirement over data sharing and research. Each of those presents unique challenges for higher education. We have to share across borders, across environments, and we've got to protect the data, but we also need to share the data with the right people. That's a very complex environment to protect and secure.

What must be in the mix of protection and tools or approaches to security in this AI-enabled world? How is that evolving?

Stewart: All of our classic tools still apply from a protection perspective. We can't just throw out what we've got and go to brand new. You need classic protections at the endpoints (EDR), protections of the firewall, protections from the high-volume, low-sophistication attacks. Absolutely, those should always be in place, so you start there. How do you design protections for your most critical assets? How do you assume breach on those critical assets? How do you make sure you are designing the appropriate safeguards to minimize the impact of an insider?

Tools are great, but we should not discount thinking, designing, and really looking at how we take the scenarios that are in play and put protections in place against them. Red teaming really helps teams understand how to ensure that what they're doing is truly secure. Assumptions can really get us. Assumptions are one of the

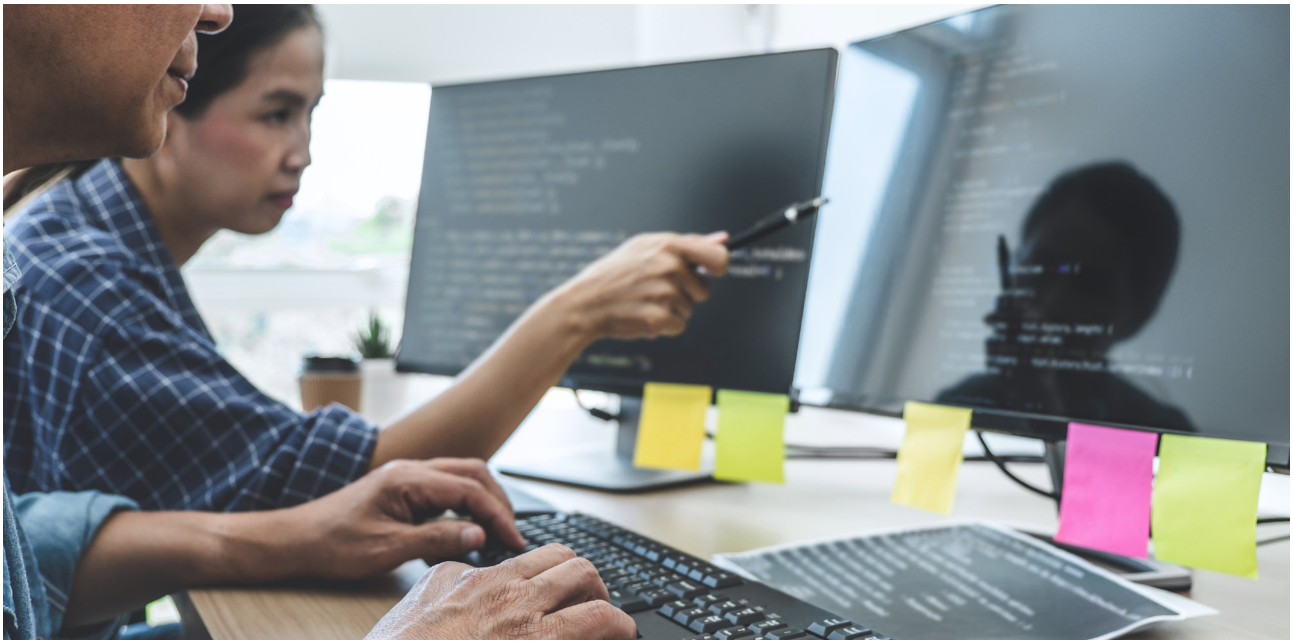

security guys' nightmares, but AI has a place there. AI can help us with the assumptions, with making sure things are done correctly and appropriately. It can help us do appropriate red teaming. It can help us design threat models.

Are there new skills required of these teams that would be using these capabilities and what does that look like?

Stewart: Artificial intelligence engineers can understand how your organization or team effectively wants to use AI and give it the right inputs, the right training model, and also ask the model questions in the right way, then interpret the results and test them. That's a different skillset than what we've needed before. It will be a skillset that evolves, like how we evolved from paper to computers, or from calculators to spreadsheets. These people will need to have the appropriate skills to ask the AI questions, and the appropriate skills to communicate well, as well as to say: "Well, now I need a human to read this." That will be that next shift: What do I trust? What don't I trust? Where do I get help to validate?

What are the tools that institutions can leverage to gain more visibility into their infrastructure and identify potential attacks?

Stewart: We've seen a lot of advances in threat intelligence, and I think we've seen a lot of advancements in endpoint detection response. There are a lot of heuristic models inside of EDRs, a lot of sandboxing and explosion of code inside of firewalls, and moves to more standard models that we see in AWS, Azure, and the cloud system. There's a lot of movement to SaaS solutions with security built in that really does a good job of protection on the back end.

Our old protection model of "I have to do this all on my own" has shifted to, "I can offset some of this," so we can offset to Microsoft 365 or some of our other SaaS tools, and I don't need to be able to do all of the care and feeding and protection of the system. I do have to do the identity management and the configuration, but once that's there I can pull away a lot of the stuff that I would have had to do years ago, so that's shifting our model.

Make a Thoughtful and Strategic Plan

As with any new technology implementation, proper planning is required before signing on to any AI/ML-enhanced tools.

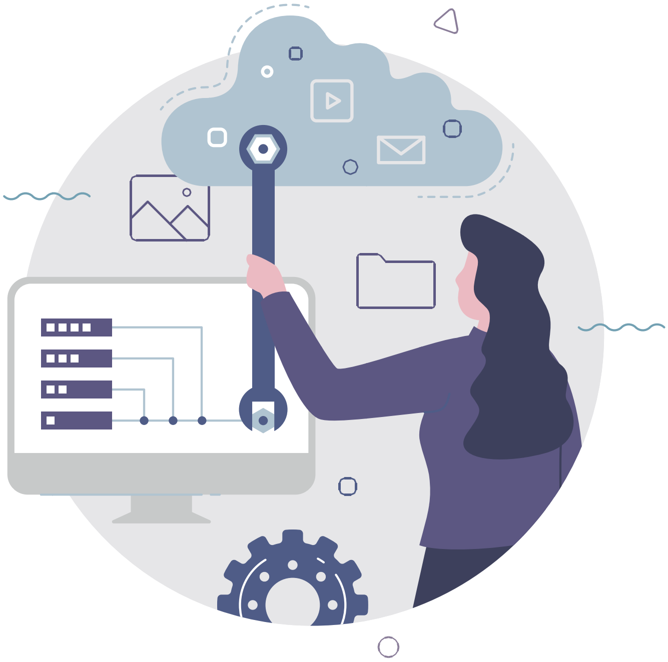

AI'S DOUBLE-EDGED SWORD can be navigated, with support from appropriate partners and through proper user education, ensuring everyone understands the risks and benefits, as well as where to turn for help.

Institutions must work diligently to ensure any vendor partner they engage with is clear and open about their software supply chain, and provides visibility into all of the services and components that make up any given toolset, advised Tim Brown, CISO at SolarWinds.

"We don't necessarily know all of the components that make up a service," Brown said. "We don't know how much duct tape and chewing gum is holding it together. You need to develop a checklist of what your third-party risk program looks like and what you're going to ask each vendor. You need to tier it, so that with a basic service, for instance, you won't be sharing any of your data or access, but you still need to know if they meet general hygiene."

Mitigate Known and Unknown Risks

Beyond vendor partner management and due diligence, it's just as critical that institutions put proper internal procedures in place to aid users weighing whether to leverage AI and spell out how, where, and in what capacity it can be developed. Setting detailed institutional or departmental guidelines for using AI, setting controls around development, as well as implementing and specifying tools for monitoring or otherwise managing the AI in place are all essential steps to take before allowing AI deployments or development to go unchecked. Once deployed, establish ongoing management and testing protocols and schedules for all AI systems throughout their lifecycle, from development through deployment and maintenance.

> Setting detailed institutional or departmental guidelines for using AI, setting controls around development, as well as implementing and specifying tools for monitoring or otherwise managing the AI in place are all essential steps to take before allowing AI deployments or development to go unchecked.

Consult frameworks or guidance published by organizations like NIST, which in January published the AI Risk Management Framework along with the companion NIST AI RMF Playbook and other materials to help cybersecurity and risk management leaders codify or establish their own protocols for successful AI implementation.

Keep current with rapidly changing news and trends around the technology, as well as unintended use cases or consequences of uses or challenges within higher education more broadly.

"Despite the real and potential promise of generative AI applications in higher education, several risks remain," Gartner analyst Tony Sheehan **wrote recently for EDUCAUSE**. Beyond all of the security risks spelled out in this report, broader risks to using generative AI in education and research include hallucinations, or false answers generated without appropriate understanding or context; subpar training data, where insufficient, obsolete, or sensitive information or biases lead to biased or incorrect responses; copyright violations; and deepfakes that appear realistic but may actually be fake content and other fraud and abuse, particularly where bad actors exploit AI to craft fake content or spamming and phishing texts or images, according to Sheehan.

All Eyes Wide Open

At the end of the day, AI holds a solid place in forward-looking cybersecurity strategies. Threat detection, vulnerability management, access control and incident response technology all rely increasingly on baked in AI capabilities that are now table stakes when it comes to an institution's ability to meet the daily complexities and volume of cyberthreats teams now face.

"We are in the very early days of seeing how AI is going to affect education," Ellen Wager, a partner at North Coast EduVisory, **recently told *Campus Technology***. "Some of us are going to stay focused on the basic research to test hypotheses. Others are going to dive into laboratory sandboxes to see if we can build some new applications and tools for ourselves. It's going to be hard to keep up, to filter out the noise on our own. That's one reason why thinking with colleagues is so very important."

Key Types of AI Risks	Risk Category	Executive(s) Responsible	Action Plan for AI Risks		
Regulatory	Adhere to regulations	CIO/CTO and CRO	Understand the continuously evolving regulatory landscape	Enable collaboration between AI practitioners and legal, risk and security members to evaluate use case feasibility and acceptable risks.	Create an AI governance office, which serves an independent audit committee to review results.
Reputational	Secure and safe	CIO/CTO	Acknowledge the threats against AI posed by both malicious and benign actors in your organization.	Bolster security across enterprise security controls, data integrity and AI model monitoring.	Leverage external resources to help secure your AI systems.
Competencies	Technical debt	CIO/CTO	Align AI strategy with cloud strategy and explore cloud as foundation for AI.	Create a technology roadmap to modernize data and analytics infrastructures to align AI goals and timeline.	Create a startup accelerator program to reduce technical debt and innovate incrementally.

Source: Gartner