AI ADOPTION AND DEVSECOPS: STAYING AHEAD WHILE STAYING SECURE

A global survey of software developers and cybersecurity teams

INFORMATION WEEK

Sponsored by JFrog

Executive Summary

Over the past decade, software developers and security teams have grown to understand the importance of integrating application security into the software development lifecycle. This is achieved through disciplines like development security operations (DevSecOps) and the use of a consistent, security-conscious continuous integration and continuous deployment (CI/CD) delivery model. These teams have made significant progress in effectively tracking and securing code from design through testing to production or cloud deployment.

The overall maturity of these disciplines is improving as more teams embed security, visibility, and control into their software development processes, allowing for better governance over the use of third-party and open source software components.

The rise of artificial intelligence and machine learning (AI/ML) is spurring more widespread use of large language models (LLMs) in development practices, creating a tectonic shift for how software engineering, security leaders, and executives manage their users and systems.

The software engineering world is at the beginning of what's expected to be a rapid AI adoption cycle. But data from this report indicates that AI/ML-related development is still operating in a silo outside the purview of mainstream application security (AppSec) programs—and often outside the software engineering team altogether. Development organizations are increasingly challenged to bring the same level of security, visibility, and control over AI/ML components as they apply to the rest of their software supply chain.

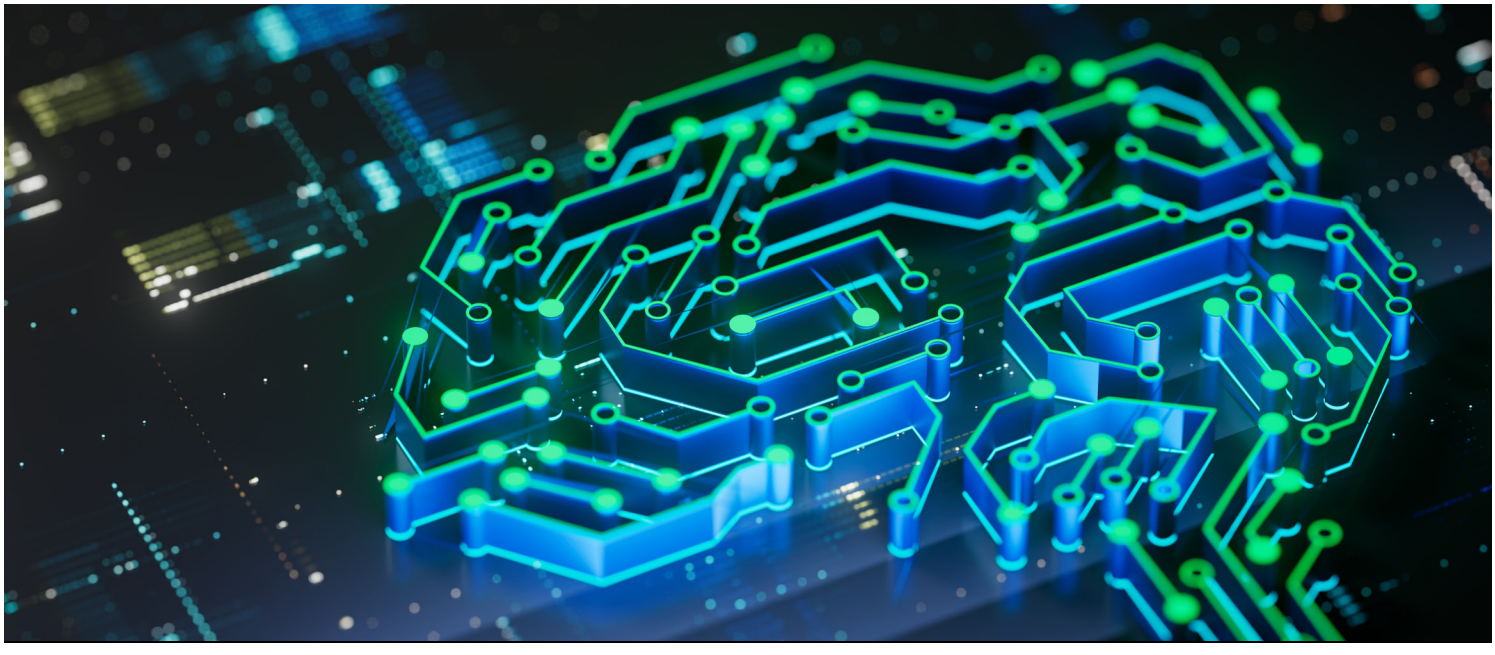

Insights from the survey include:

Lack of business-wide AI security confidence

- 79% of firms say security concerns are slowing the use and/or integration of AI/ML features into software
- The top three AI security concerns within companies are data exposure through LLM usage, malicious code in AI models, and AI bias
- 64% of organizations are either not at all confident or only somewhat confident in their ability to comply with new and emerging regulation around AI usage in software

LLM/AI policies are still lacking

- Only 48% of companies have an enterprise-wide policy for LLM or Generative AI (GenAI) technology usage
- Two out of five companies lack any kind of governance framework for how developers use LLMs
- 58% of companies either have no policy in place or don't know if they have a policy that sets rules for how developers use open-source AI models or components
- 60% of companies don't have a policy for how developers source or license their training data

Enforcement is even spottier

- 68% of firms can't or aren't sure if they can detect where source code comes from when developers use LLM/GenAI tools in their workflows
- 68% of respondents report they have no way to enforce AI component usage or depend on manual review to do so
- 59% say they have no mechanism or rely on manual review to enforce policies about training data

AI supply chain visibility is muddled

- 49% of firms have no reliable way to control usage of ML models in their apps
- Less than a quarter of organizations have a single source of trust for all software components, including AI models
- More than two-thirds of organizations have no reliable method for tracking open-source packages in their software containing transitive dependencies to ML models

MLSecOps is still emerging

- Only 21% of organizations have a formal machine learning, security, and operations (MLSecOps) team or program
- Just 9% of organizations currently conduct both red teaming and threat modeling for AI software

AI/ML Adoption Trends and Attitudes

While most industry analysts are optimistic about the use of AI and ML within enterprise software in the coming years, adoption is nascent. Software engineering is still in the very earliest stages of its journey toward AI maturity. The inclusion of AI/ML features hasn't reached a majority tipping point, with most practices indicating experimental and isolated deployments. Approximately 41% of survey respondents say that under one in ten applications their organization develops today contains or makes calls to machine learning models or AI services **(Figure 1)**. Meanwhile, only 7% of organizations report the same for half or more of their applications.

Nevertheless, those numbers are likely to grow. Almost a quarter of firms have policies that allow their developers to integrate LLM/GenAI into at least some of their applications, with a sliver of firms—6%—giving their developers carte blanche to integrate LLM/GenAI into any of their applications **(Figure 2)**.

In spite of the engineering excitement around AI-related software development, security is causing many firms to slow their pace of adoption. An overwhelming majority of firms—79%—believe security concerns are slowing rollout of AI/ML technology and integration of AI/ML features into software made by the organization.

Figure 1

Machine Learning in Apps

Thinking of the applications your teams develop today, what percentage contain or make calls to machine learning models or services?

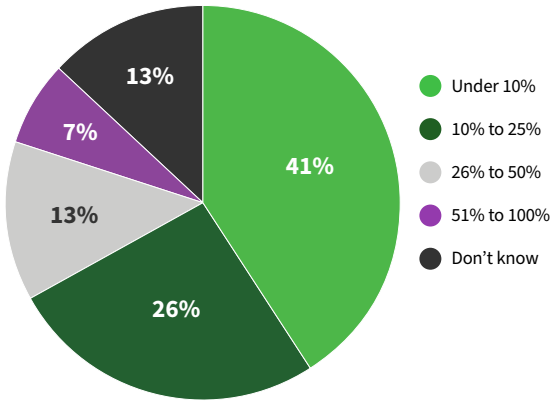

Data: InformationWeek survey of 210 cybersecurity and IT management, May 2024

Figure 2

Policy for Developers' Use of LLMs/GenAI

Do you have a policy specifically for how developers use LLMs/GenAI?

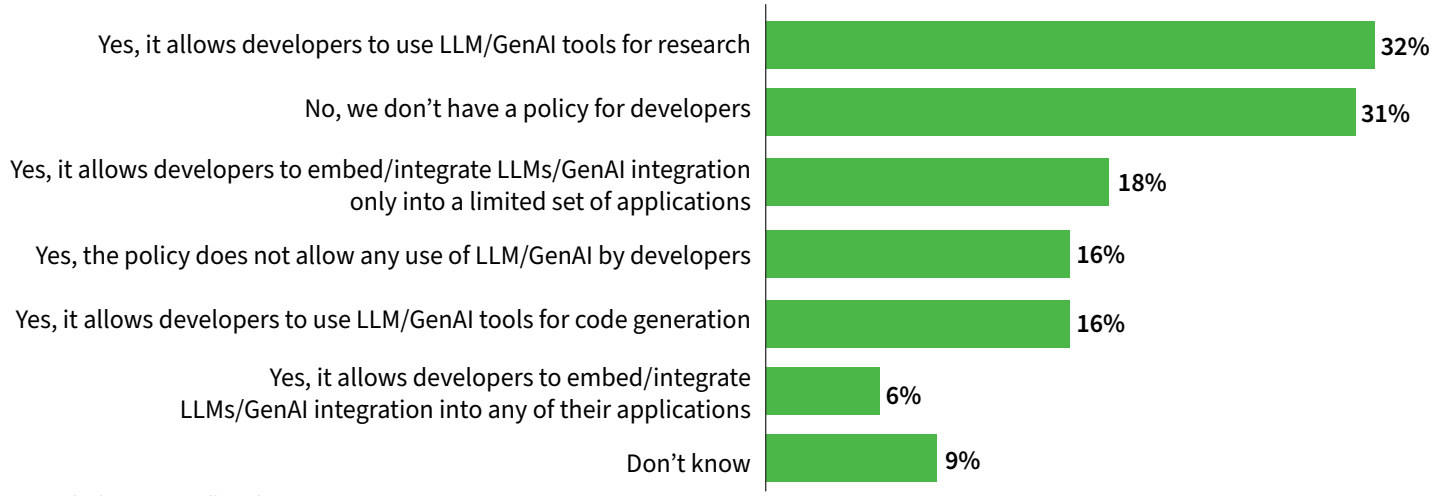

Note: Multiple responses allowed
Data: InformationWeek survey of 210 cybersecurity and IT management, May 2024

Figure 3

Top AI Security Concerns

What are the top AI security and compliance concerns at your organization?

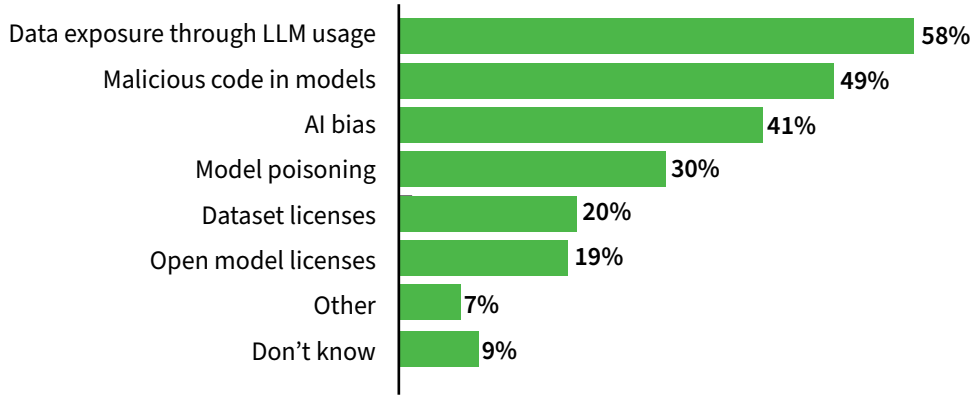

Note: Maximum of three responses allowed
Data: InformationWeek survey of 210 cybersecurity and IT management, May 2024

The top three AI security and compliance concerns at organizations are:

- Data exposure through LLM usage (58%)
- Malicious code in models (49%)
- AI bias (41%)

Other concerns weigh heavily as well, with 30% of organizations naming model poisoning as a top issue, and about one in five organizations reporting that licensing for datasets and ML models stands as a challenge **(Figure 3)**.

Ultimately, only a few companies have confidence in their ability to keep up with the fast-emerging AI security standards and regulations. Approximately two-thirds of firms say they're not at all confident or only somewhat confident in their ability to comply with new and emerging regulation around AI usage in software.

AI Policies Are Still Primitive

Managing the risk associated with AI and LLMs starts with establishing policies that govern the use of this technology within an organization. Most companies lack a comprehensive policy for the application of LLMs and GenAI by employees in various departments, including non-tech areas like marketing, sales, and legal.

Less than half (48%) of respondents reported that they have this kind of enterprise-wide policy guiding employee use of LLM/GenAI technologies. The rate of policy leadership dips considerably at smaller companies. Only 37% of companies with fewer than 500 employees have a policy that guides LLM use across the business.

Digging into how companies govern LLM/GenAI usage among developers, almost a third of firms have no policies at all regarding how developers code with LLMs or even what kind of LLM technologies they're allowed to embed in enterprise software. Another 9% don't know whether they have a policy, which indicates these firms may not have one at all—or at least may not have one with enough visibility to be impactful.

Further examination of established policies reveals that two in five firms lack any significant kind of governance framework or policies for developers as they navigate rapid proliferation of LLMs.

With regard to LLM use within day-to-day development activity, only 16% of companies openly permit devs to use GenAI for code generation. An equal proportion of firms completely ban use of LLMs by developers. Many more—32%—take a middle-of-the-road approach, condoning developers to use LLMs for research and ideas as they code but prohibiting code generation since the resulting code is not always optimal and has been shown to introduce vulnerabilities.

Read more about analyzing common vulnerabilities introduced by Code-Generative AI.

But even with solid policy guidance many organizations will likely struggle to enforce the rules unless they lock down the developer desktop. In most organizations, this is not yet a practical option.

Delving into the types of AI that developers are embedding into their software, as mentioned earlier, 24% of firms say they allow some level of LLM/GenAI integration. As for broader adoption of AI components and ML models, many companies are still in exploratory phases and don't have a policy to govern that. Incredibly, 58% either have no policy in place or don't know if they have a policy that sets rules for how developers use open-source AI models and components. Even more companies lack guidance for how developers find and use training data. Some 60% of respondents admitted that their firms don't have a defined policy for training data sourcing and licensing.

Meanwhile, when companies that do have AI/ML model and component use policies, the content is inconsistent. Sixteen percent of firms do not allow developers to use any kind of AI/ML models or components in software they currently create, whereas 15% do allow use, but only with private AI/ML models and components provided by internal data science or ML teams **(Figure 4)**. The survey showed that 32% of organizations do officially allow open-source components, but 22% say only when models and components are approved by the organization.

Figure 4

Security Policy for Integrating ML Models

Do you have a company security policy on how AI components and ML models are integrated into software today?

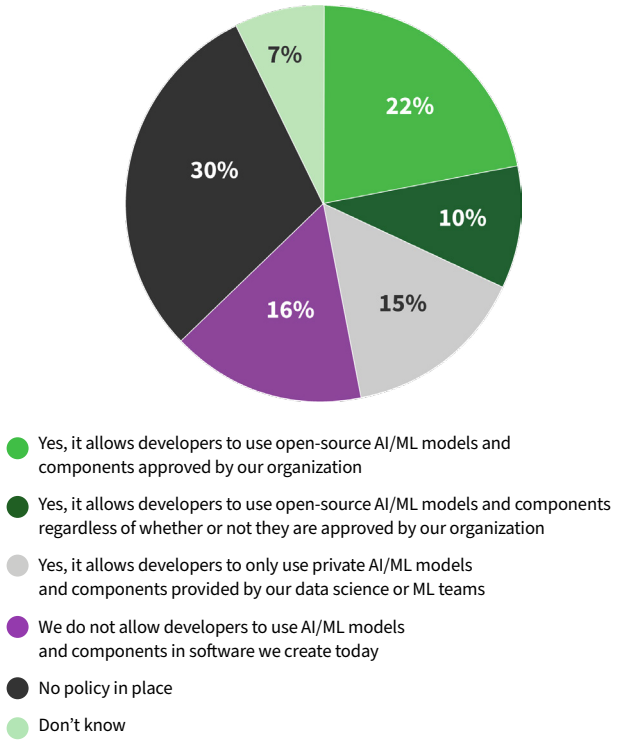

Data: InformationWeek survey of 210 cybersecurity and IT management, May 2024

Enforcement of Policies Remains Limited

Clearly, many organizations are open to significant risk due to a lack of policymaking (or perhaps the inability to make or enforce policies) from IT and business leadership. The survey shows an even deeper risk lingering under the very thin security blanket of policy oversight. Additional survey questions indicate that organizations with AI development policies often lack the means to enforce them effectively.

When examining how LLM/GenAI usage is enforced among developers, 37% of respondents cited security scanning tools, automated policies, and regulated developer environments as the most common methods.

The tooling cited is either extremely limited or nascent in its capability to perform this kind of enforcement work. In most cases, these cited tools will struggle to truly control the use of LLM in coding and to detect LLM integration in software. Even more troubling is that right behind these top enforcement answers, the next most cited method was manual review, which was named by 33% of firms **(Figure 5)**. As with many other areas of security oversight, the manual enforcement option could be ineffective, as there are just not enough people even in the largest organization to keep tabs on this activity.

What's more—in general—there aren't many foolproof traditional approaches for companies to detect if the source code is derived from human or LLM. A solid 68% of respondents admitted that they either can't or they're not sure if they can detect where source code comes from when developers use LLM/GenAI tools in their development workflows.

Enforcement capabilities are even more limited when looking at the ability to enforce ML model/AI component usage, with 63% of firms reporting that they either have no way to enforce component usage or depend on manual review. The numbers look similar for enforcing policies about training data sourcing and licensing, for which 59% of respondents say they have no enforcement mechanism or rely on manual review.

Figure 5

Enforcing Policies

How do you enforce those policies and manage usage of LLMs/GenAI by developers?

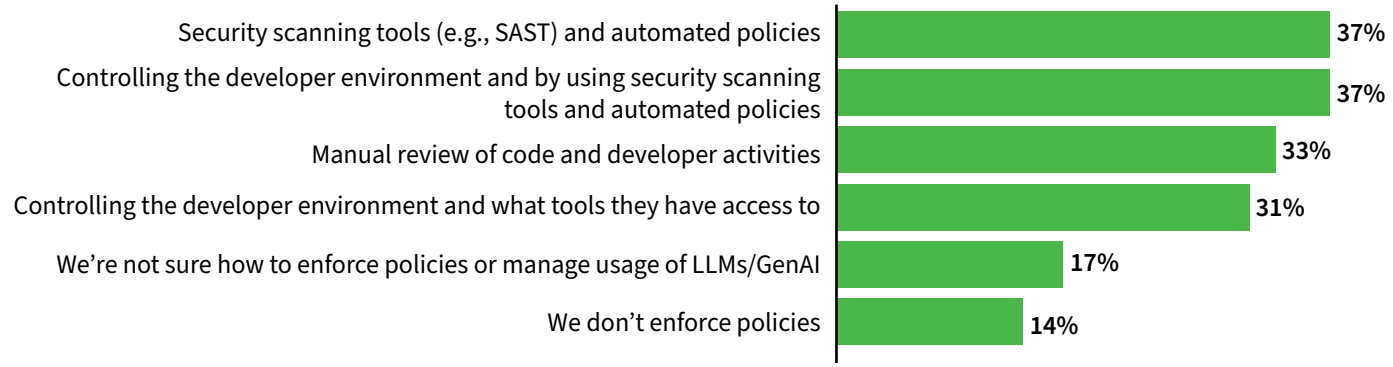

Note: Multiple responses allowed
Data: InformationWeek survey of 210 cybersecurity and IT management, May 2024

Again, in today's pace of development, manual reviews simply do not scale. No matter what kind of business a software engineering organization operates within, the prevailing trend is that organizations want developers to move fast. If AI security depends on manual steps to be taken, organizations will see these practices kill the productivity of developers, data scientists, and business stakeholders alike. The slowdowns will hurt competitiveness and time to market, and those stakeholders will look for ways around security best practices. In modern development, the most secure organizations strive to make security as automated as possible and use manual review as a backup mechanism. These survey results indicate that AI-related policy enforcement still has a way to go to catch up with traditional standards for security oversight.

Enterprises Still Seeking AI Supply Chain Visibility

Most application security teams are already struggling to build up software supply chain visibility across their traditional portfolio of applications. Survey results indicate that AI will add new layers of obscurity to the supply chain as AI software components, models, and training data are added into the mix. The data indicates that organizations are not ready to make those AI supply chain risks visible, let alone manageable.

Respondents reported that developers and AI/data science pros are sourcing models, AI components, and AI training data from a range of sources. Some of the most common AI repositories used today are Microsoft AI GitHub, PyTorch Hub, TensorFlow Hub, and Hugging Face **(Figure 6)**.

Just a scant 23% of organizations have a single source of truth for software components that includes visibility into components that go into training and using AI models in software. In fact, 19% say they struggle to track regular software components, let alone ML components.

When it comes to monitoring and controlling how AI models are used in code, the methods are still crude and overlapping. An alarming 49% of firms say they either don't know or have no way to monitor and control usage of ML models in their apps. Meanwhile, 34% say they have a way to do it, but it doesn't include transitive dependencies. Just 17% say they can monitor and control ML models, including transitive dependencies.

Figure 6

AI Model Repositories

Which AI model repositories do your data science and development teams use?

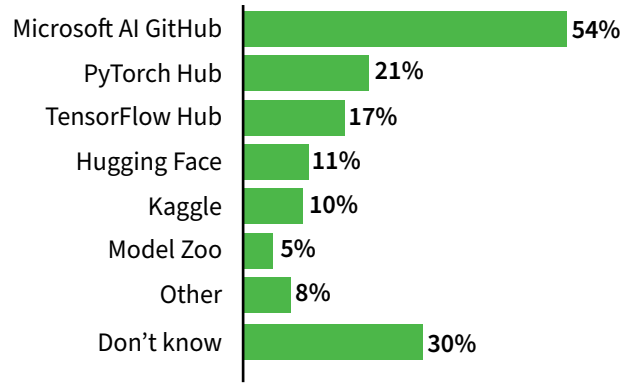

Note: Multiple responses allowed
Data: InformationWeek survey of 210 cybersecurity and IT management, May 2024

In many cases, development organizations can't even track the open-source packages used by developers that contain transitive dependencies to ML models. Over two-thirds of organizations either don't have or aren't sure if they have this type of visibility into their open-source usage. This is a huge visibility issue as ML use becomes increasingly embedded into open-source ecosystems.

Another big issue is model version control. While over half of organizations say they have version control for ML models in their development workflow, 29% say they're not confident in these mechanisms.

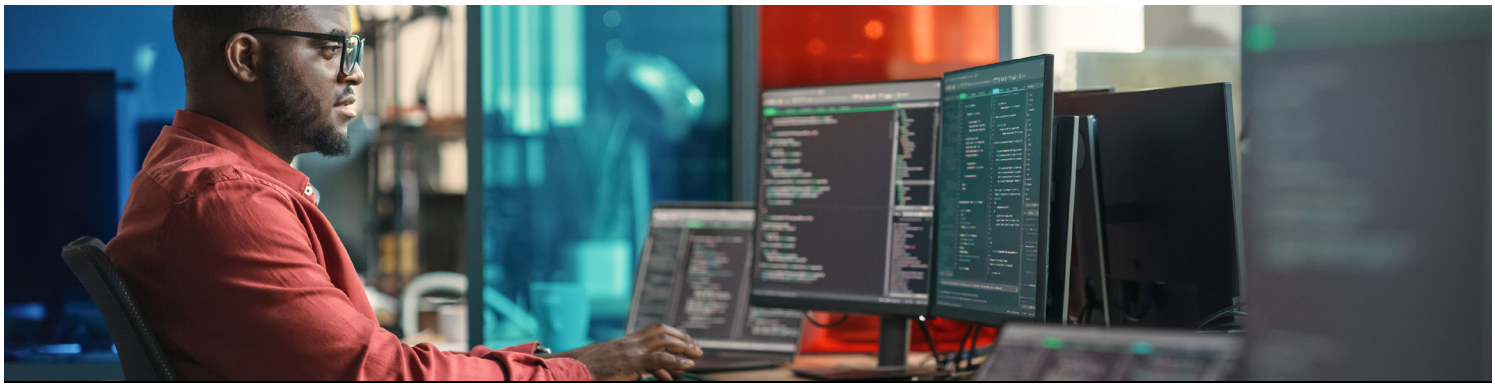

Meanwhile, on the data lineage and data governance side of things, only 21% of firms say they keep their training data use locked down to internal, first-party sources **(Figure 7)**. But 57% of respondents say they either use a combination of third-party and first-party data or they solely rely on third-party data to train their AI models. Another 22% don't even know where their training data comes from, indicating an extremely low level of maturity about AI data governance.

The most common model training dataset sources include Microsoft Azure Open Datasets, Google's Dataset Search, and datasets on GitHub. But again, only 25% have a way to track and control who trains or interacts with model training data throughout its lifecycle with a high degree of confidence.

Just as with model version control, over half of organizations have some mechanism in place to govern ML dataset versioning, but 31% are not confident in it.

Finally, there's the issue of harmonization of version controls. Most organizations have found a way to integrate version controls for models and data, but it tends to be hodgepodge. Around a quarter of organizations have cobbled together a homegrown solution, and another 30% say they have controls on the same platform but that it still has kinks. Just 37% say it's on the same platform and works well.

Figure 7

Source of Training Data

Where do you most commonly source model training data from today?

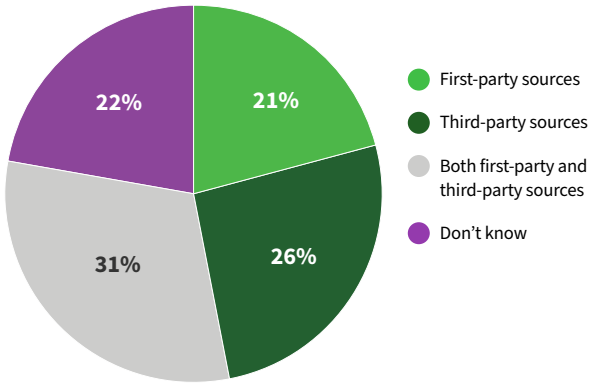

Data: InformationWeek survey of 210 cybersecurity and IT management, May 2024

Nascent MLSecOps practices

Organizations need a holistic program to roll out comprehensive ML, security, and operations practices that work well together. To support that program, they also need the right tooling that bridges visibility and control over ML components across data science and developer environments. As things stand, though, only 21% of orgs have a formal MLSecOps team or program to carry out this mission.

Many security teams are examining embedded AI risks through threat modeling and/or red teaming exercises, though these practices are still emerging. The survey found that 21% of organizations today solely do threat modeling for AI, and 11% only do red teaming. Just 9% do both.

Security awareness is also still low among data science teams. Today, only 21% of firms regularly train their data scientists on security principles. The good news is that 36% of companies at least give their data science staff some initial basic training, and 28% do semi-regular or ad hoc training **(Figure 8)**.

Figure 8

Training on App Security

How well trained are your data science teams in application security principles?

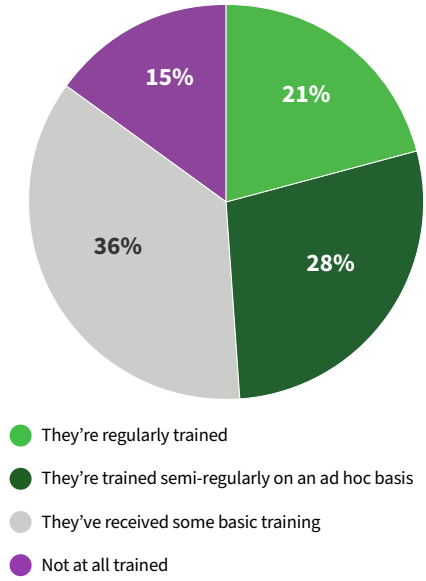

Data: InformationWeek survey of 210 cybersecurity and IT management, May 2024

The No. 1 security technology that firms today use to manage models and datasets, protect their integrity, and prevent model theft is role-based access control. This was named by 44% of respondents **(Figure 9)**. Other common mechanisms are data encryption and digital rights management.

The reliance on role-based access control is a nice glimmer of hope for the future. While additional infrastructure and practices are still definitely necessary to appropriately control how and when developers, data scientists, and business users utilize AI features and components, access control will be foundational to all of them. Most mature security control is based on roles, and AI will be no different. It's good to see that this table-stakes architecture is in place at a large number of organizations today, as it will be essential to further maturation of MLSecOps practices in the future.

Conclusion

The adoption of AI components and integration within enterprise software presents tremendous business opportunities for organizations in the very near future. But as organizations gear up for this AI-led digital transformation they'll need to do that in line with their security and compliance obligations. The prevailing winds of today's risk environment—including regulatory and attack trends—demand that technology and business leaders start collaborating to design a software pipeline that delivers responsible and trustworthy AI.

Based on the current survey results, it's clear that organizations lack the visibility, governance, traceability, integration, and trust in AI/ML components that they have spent the last decade incorporating into their traditional software development pipeline.

Survey results also reflect that we're in the early days of AI adoption and enterprises have the opportunity to bridge the gap between AI/LLM development and DevSecOps before their practices solidify. One of the best places to start is by treating AI models like packages and governing models and AI components directly in the same workstream of artifacts that make up the rest of their software supply chain.

Figure 9

Managing Models and Datasets

What mechanisms do you leverage to manage models and datasets, protect their integrity, and prevent model theft or data exposure?

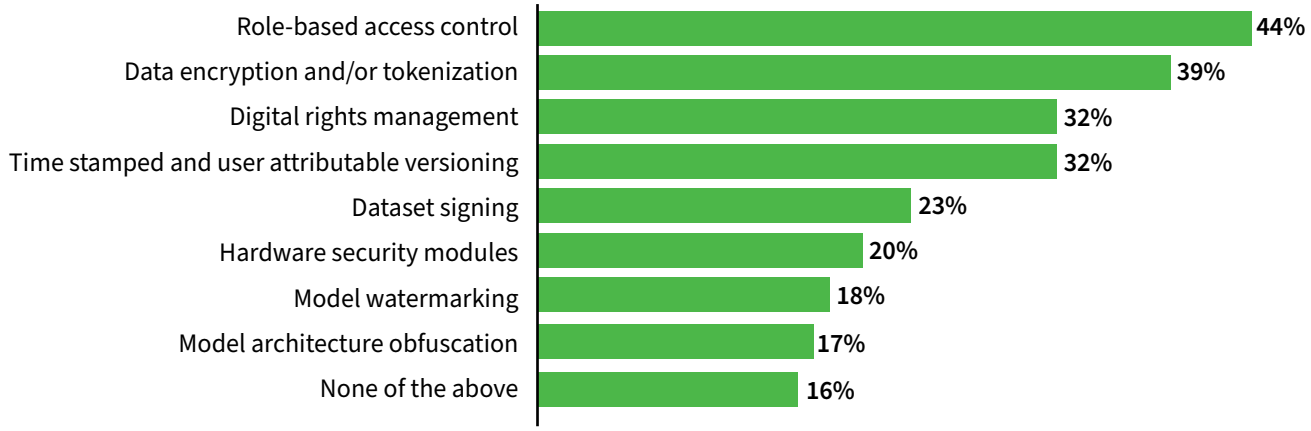

Note: Multiple responses allowed
Data: InformationWeek survey of 210 cybersecurity and IT management, May 2024

Survey Methodology

JFrog commissioned InformationWeek to conduct a survey to explore the rise of open-source and commercial artificial intelligence and machine learning solutions. The survey explored how well software developers and cybersecurity teams understand the importance of integrating application security into the software development lifecycle. Additionally, it looked at how cybersecurity and IT teams are protecting their organizations against malicious code and securing corporate data from improper use of AI technologies.

The survey was conducted online in May 2024 and collected responses from 210 IT and cybersecurity professionals predominantly located in North America. Respondents were recruited via email invitations containing an embedded link to the survey. The emails were sent to a select group of Informa Tech's qualified database. Informa Tech is a division of Informa, the parent company of InformationWeek. Informa Tech was responsible for all survey design, administration, data collection, and data analysis. These procedures were carried out in strict accordance with standard market research practices and existing U.S. privacy laws.

Respondents hailed from companies of all sizes. Twenty-nine percent of respondents were at large companies with 5,000 or more employees; 29% were at companies with 500 to 4,999 employees; 15% were at companies with 100 to 499 employees; and 27% were from organizations with fewer than 100 employees.

The final dataset includes job titles from executive level to staff. Thirteen percent held IT or cybersecurity executive job titles (CIO/CTO, CSO/CISO). Nineteen percent were titles including IT director/manager or networking director/manager, and 8% were cybersecurity director/manager titles. Other titles included IT and cybersecurity staff, corporate management, engineer/ML engineer, and consultants.

More than 21 vertical industries are represented including consulting, banking and financial services, education, government, technology, healthcare, and manufacturing.