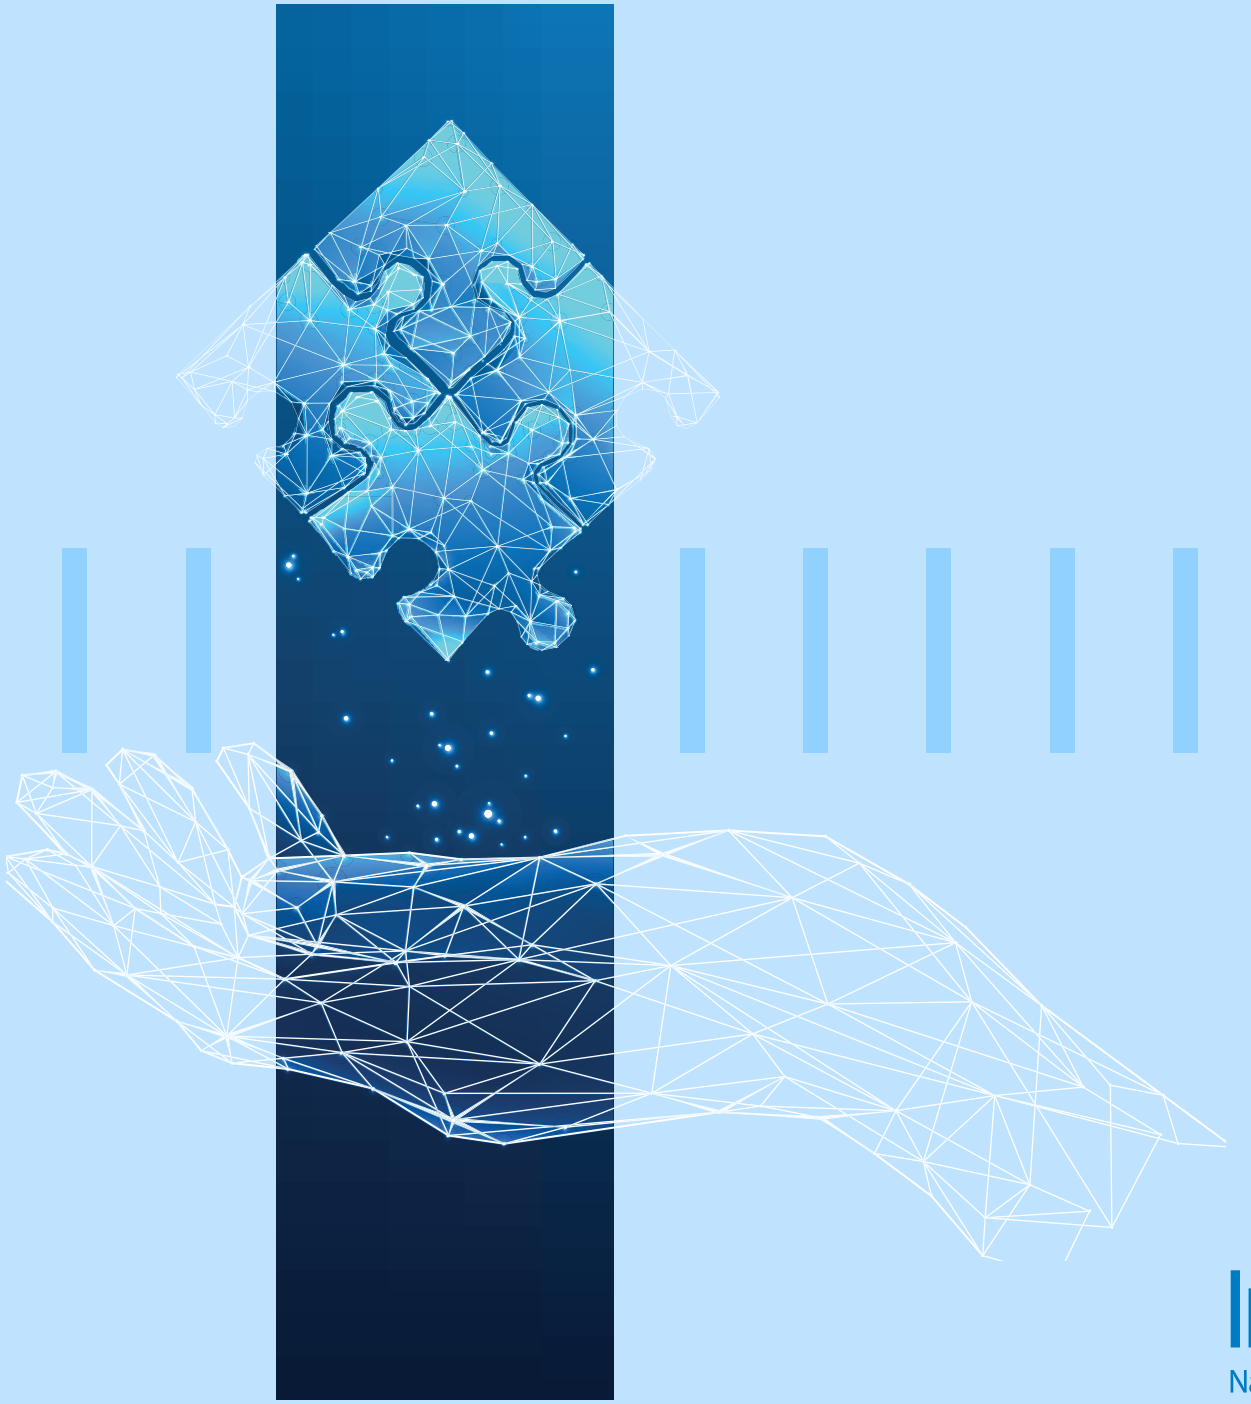
THE GAP BETWEEN KNOWING AND DOING

A PRACTICAL AGENDA FOR CYBERSECURITY IN THE LLM ERA



Infosys
Navigate your next

Executive Summary

Fifty-one professionals in the cybersecurity space, drawn from a variety of backgrounds in both academia and industry, answered a structured questionnaire at the ACM Asia Conference on Computer and Communications Security. Their responses reveal a community that is technically fluent, and yet operationally under-equipped. The meteoric rise of AI and large language model adoption has outpaced the governance and defensive tooling required to secure it. This report conveys what these professionals told us: where they are at now, what they will fear in the future, and what they are building to prevent those fears from being realized.



External Document © 2026 Infosys Limited

Who Answered, and Why It Matters

Background of Respondents:

PhD/Graduate students
51.9%

Academic Researchers
35.3%

Industry Practitioners
15.7%

The majority of respondents are in academia, either as researchers or students. These are people who spend their time carefully studying threats, reading papers, building proof-of-concept systems, and thinking about what is next. Industry practitioners make up a smaller but distinct section of the cohort. This cohort reflects a sophisticated, forward-looking group that are not just IT generalists filing out a survey.

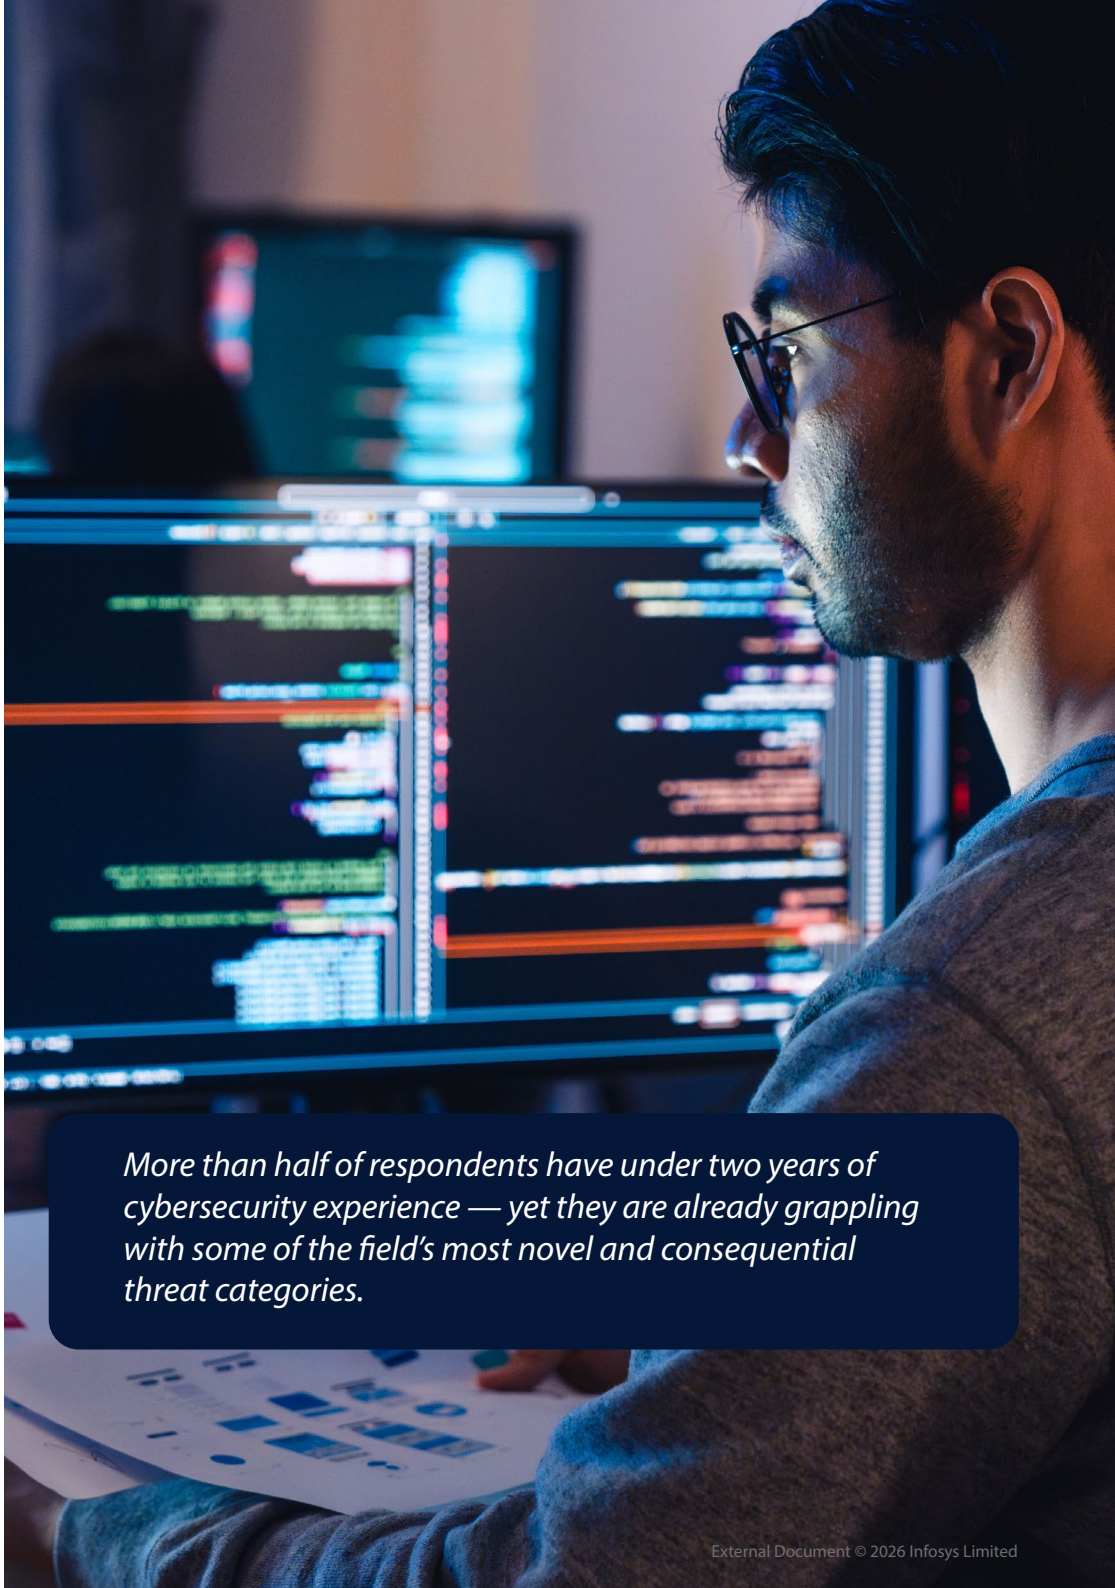
Experience in Cybersecurity:

0-2 years
51%

2-5 years
39.2%

5+ years
9.8%

Experience levels skew junior but the threats addressed by this survey, including LLM prompt injection, agentic AI, and non-human identities, are new enough that more experienced practitioners may not actually be the most informed. In this area, graduate students and professors working on cutting edge research may have sharper threat models than decade-seasoned engineers.



More than half of respondents have under two years of cybersecurity experience — yet they are already grappling with some of the field's most novel and consequential threat categories.

External Document © 2026 Infosys Limited

Survey Findings

AI for cybersecurity has become operational: AI is used daily for vulnerability discovery, anomaly detection, malware classification, secure code generation, log analysis, and threat hunting.

Control lags behind: Defense against prompt-injection, observability tools built for AI, and identity governance lag behind usage.

Guardrails need to adapt: Agentic systems change the landscape of security threats in the cyber space; IAM platforms must account for what these systems can do, and platforms need to be prepared for new attack surfaces.

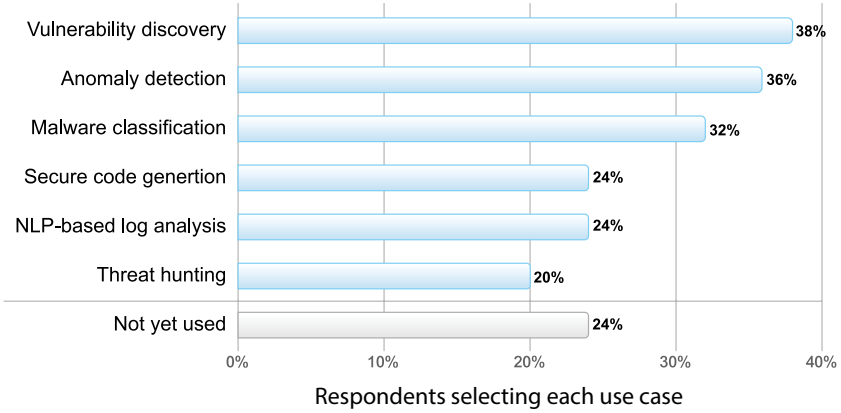
Post-Quantum planning needs to begin: Cryptographic migration needs to begin with the data and systems that need protection beyond the expected life of today's algorithms.

Finding 1: AI for cybersecurity has become operational

The use of AI tools in cybersecurity is no longer experimental; these tools are regularly used by the survey respondents. 38% use them for vulnerability discovery, 36% for anomaly detection, and 32% for malware classification. Over a quarter of respondents use AI for secure code generation or analysis of log data, and one in five use it for threat hunting. Only 24% report no AI or machine learning use in their security work.

AI and machine learning are already embedded in security work

Share of respondents selecting each use case (n=50)



Source: Asia CCS survey Narrative Repost (n=50). Multiple selections were allowed.

Among survey respondents, there is a wide range of deployments in use. 62% of respondents use proprietary models such as Claude or GPT, 27.7% use open-source models like LLaMA, and 23.4% use fine-tuned or domain specific models.

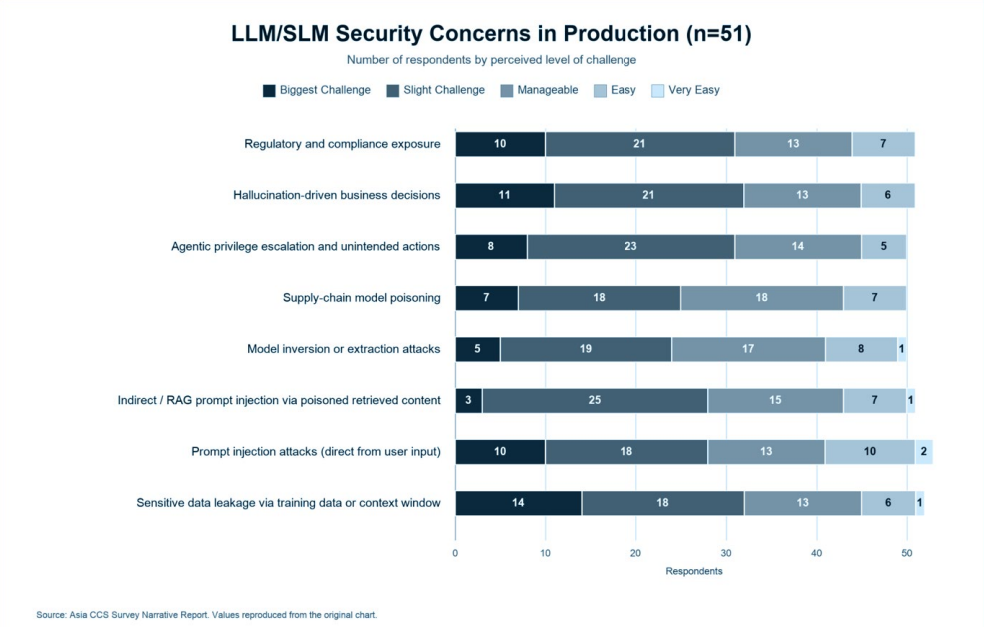
The data above highlights a growing concern among data privacy experts. Nearly two in three respondents use models like GPT or Claude that they do not host themselves, meaning third-party data exposure is an active risk. The gap between the use of these tools and the governance to regulate them begins to emerge.

Nearly two in three respondents are sending prompts to public LLM APIs — making third-party data exposure not a future risk, but a present one.

External Document © 2026 Infosys Limited

Finding 2: Control lags behind

The survey asked respondents to rank eight threat categories for LLM deployment in production. What emerged was a picture clustered around two distinct fears: what AI knows and what AI does.



Source: Asia CCS Survey Narrative Report. Values reproduced from the original chart.

The leakage of data through AI models training or context window was highlighted as the biggest challenge by fourteen respondents. What if the next AI is trained on your company's proprietary data? What if one of your employees uploads confidential client files to an AI model? These are the fears that respondents chose to highlight as one of the most pressing concerns of deploying LLMs in production.

Other concerns brought up included prompt injection attacks (when an attacker inserts malicious requests in a prompt to an AI model) and hallucination-driven decisions (AI models inventing facts or figures that are then used to make critical business decisions).

The growing capabilities of agentic AI was noted as the third biggest challenge. This is the fear of AI agents taking actions such as downloading files, calling APIs and entering secure systems without a human authorizing each action. This is the newest threat which has emerged in the last three years and is only growing in importance.



GROWING THREAT OF AGENTIC AI

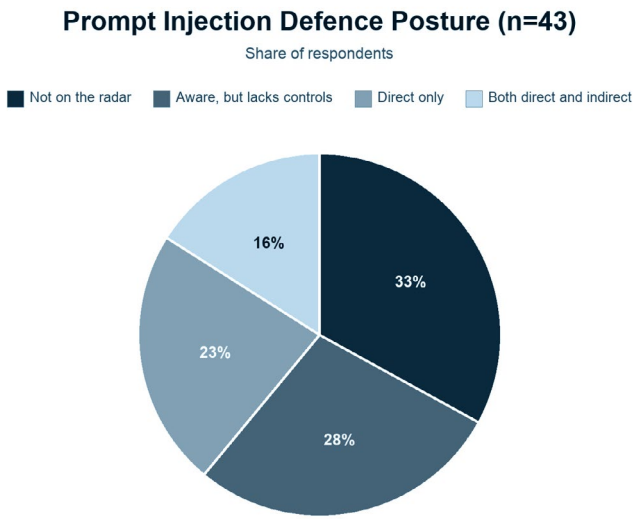
Agentic privilege escalation has the highest aggregate concern rating of any LLM threat — a category that barely existed three years ago is now the community's most anticipated risk.

External Document © 2026 Infosys Limited

Finding 3: Guardrails need to adapt

The previous finding focused on what practitioners fear about AI use, but how defended are modern systems against those fears? Results from the survey shows that respondents are afraid that security and guardrails systems are lagging behind AI advancement, and stress on the urgent need for new ways of managing these attack surfaces.

For example, prompt injection was the second-most-feared LLM production risk. When respondents were asked about deploying defence mechanisms to such attacks, only 16.3% reported doing so.



Source: Asia CCS Survey Narrative Report. Values reproduced from the original chart.

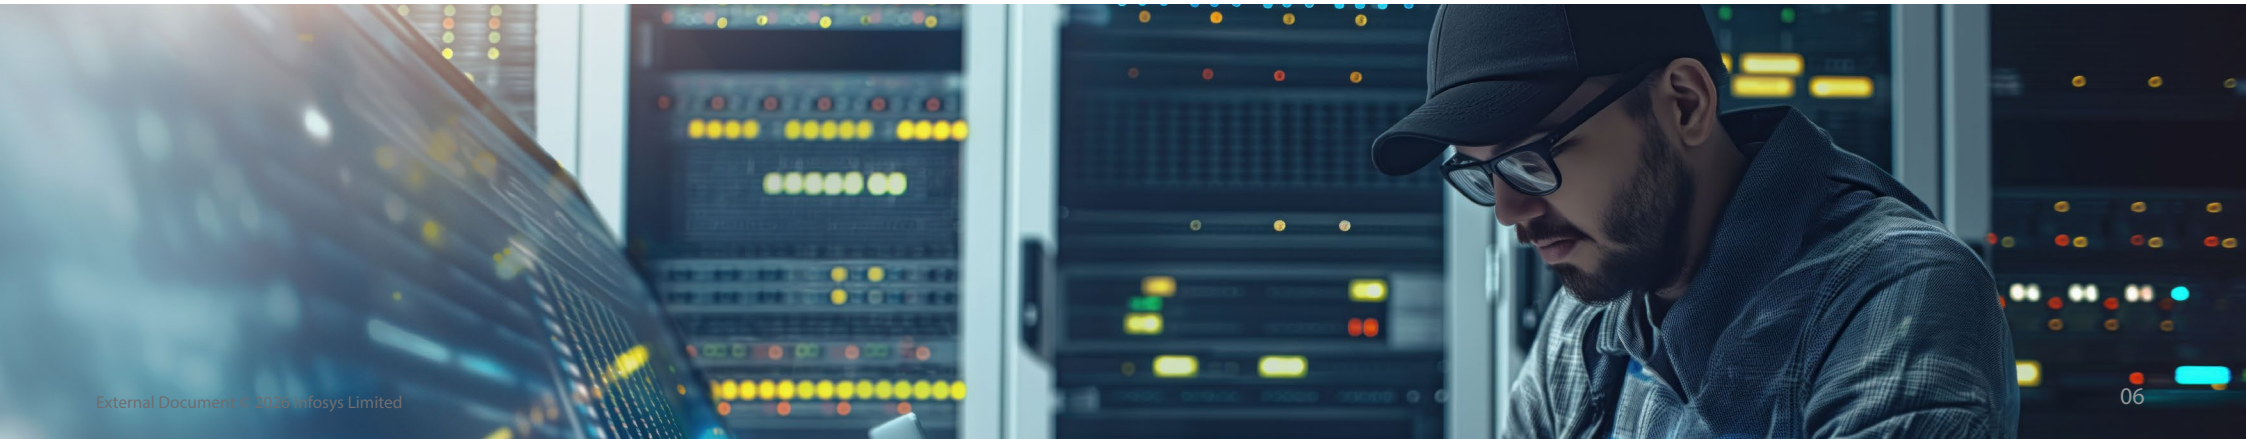
The controls that practitioners are deploying (content filtering at 47.7%, audit trails and LLM-specific logging at 36.4%) are good starts but do not hold up to new agentic use cases. The ones that do are not widely deployed. Only 4.5% of respondents reported separating model permissions from human permissions. This is a fundamental control for agentic systems that ensures agentic systems are only allowed to undertake actions that they have been allowed to.

The discussion around AI control and guardrails must also deal with integrating traditional IAM tools. Identity and Access Management (IAM) is the security framework of ensuring that secure resources can be accessed only by users with the permission to do so. In the pre-AI era, this involved API keys, service tokens, and OAuth credentials that enabled software systems to authenticate humans and other systems.

The survey reveals that only 21% of respondents have full governance over AI-generated identities at runtime. The combination of this finding with the agentic adoption survey is the single clearest risk signal. One in four respondents already use Agentic AI systems, but the systems to track their credentials at runtime aren't in widespread use.

LACK OF AGENTIC GUARDRAILS

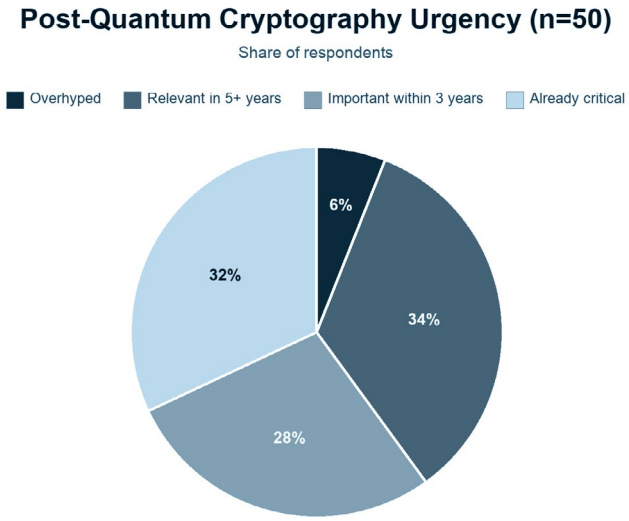
Only 4.5% of respondents have architecturally separated model permissions from human user permissions — the most fundamental control for containing agentic AI systems.



External Document © 2026 Infosys Limited

Finding 4: Post-Quantum planning needs to begin

The discussion on post-quantum cryptography provoked more disagreement among the survey's respondents than any other topic. When asked how urgent the transition to post-quantum cryptography is, the community was split evenly across three positions, reflecting a divided community with different threat models.



Source: Asia CCS Survey Narrative Report. Values reproduced from the original chart.

Even within the same community, practitioners shared very different timelines on when to act - 32% percent say it is already critical and that organisations need to act now, while 34% say it will only become relevant in five or more years. It is important to note that only 6% dismissed it as overhyped, reflecting the importance of quantum risk in near-unanimous terms even if assessments on urgency differ.

The differing opinions between the group who believes we need to act now versus waiting is not simply a case of being right or wrong. Organizations who handle sensitive data with a long shelf life - patient records, financial documents, national security information - worry about data that is stolen today and hacked in the future. For such organizations, the risk of a data breach in a pre-quantum world that is later decrypted in a post-quantum world is a real threat.

Data that is sensitive today and will still be sensitive in the future needs quantum-resistant protection now. For organizations handling data with short sensitivity windows, the urgency may not be as important, but the threat still warrants a valuable conversation about readiness and post-quantum preparation.

IS MY DATA SAFE?

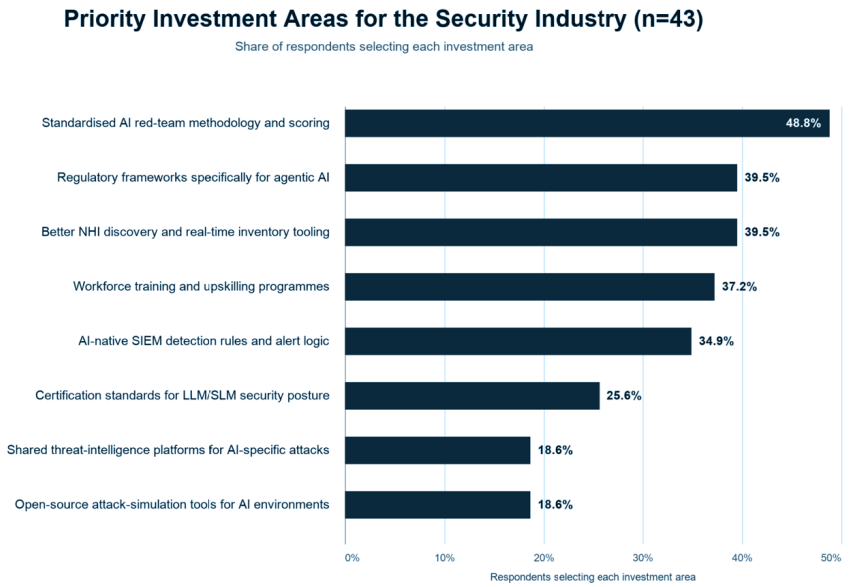
The 34% who view PQC as a 5-plus year concern should ask themselves: how long will the data I encrypt today remain sensitive? If the answer is longer than a decade, they are already exposed.



External Document © 2026 Infosys Limited

What the Community is Building Toward

The survey and the discussion it prompted makes it clear that investment is required to ensure that our cyber systems are secured and trustworthy. However, executives are often left wondering where this change needs to begin.



Source: Asia CCS Survey Narrative Report. Multiple selections were allowed.

The final part of the survey asked the respondents to reflect on which areas in the security industry they would want to prioritize. 48.8% of respondents named standardized AI red team methodology as the single largest investment priority. This is the framework for how to attack AI systems, score findings, and compare results across engagements.

Practitioners also wanted better investment in regulatory frameworks for agentic AI and non-human identity management (NHI). This reflects the findings around control and guardrails discussed earlier in this paper.



External Document © 2026 Infosys Limited



Summary

This survey highlights a field in transition. An industry that is working on absorbing novel threats that arrive each day with new innovations, while trying to operationalize defenses against them in an enterprise framework.

> *The gap between knowing and doing is not the product of complacency. It is the product of adoption velocity outrunning institutional capacity.*

The gap that exists between knowing about these threats and acting on them today is not the product of complacency, rather the result of product velocity that is outpacing the traditional institutional capacity for change.

Based on the survey, five points stand out as actionable conclusions from the data collected.

- Prompt injection defence cannot wait. 84% of respondents lacked control for such attacks in an environment where 62% are using private third-party models.
- Non-Human Identities (NHI) need to be introduced. Identity control and management systems need to track and control agentic AI at every stage of use.
- Human permissions and agentic permissions need to be separated. The pipeline to acquire credentials for humans and agents should be bifurcated and each system should be treated as its own independent security challenge.
- Post-quantum planning is critical for organizations with long data retention policies. Rather than waiting for quantum computers to emerge, organizations need to plan for the future and secure sensitive data today.
- Standardised AI red teaming needs to come online. Teams need to share a common framework on scoring approaches, producing findings, and benchmarking. Security needs standards to succeed.

External Document © 2026 Infosys Limited

Author's Bio



Vasudha Vasudev is a Consultant at Infosys CyberSecurity with five years of professional experience in the cybersecurity domain. Her work encompasses product strategy, stakeholder engagement, research, and innovation, alongside contributions to cybersecurity marketing and thought leadership. Her areas of interest include artificial intelligence, data privacy, and data protection.

For more information, contact askus@infosys.com

Infosys
Navigate your next

© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.