CISOFORUM

Outlook 2026

How Cybersecurity Leaders in India Are Navigating Digital Transformation, Regulatory Complexity, and Emerging Threats

Contents

MANAGEMENT
Managing Director: **Dr Pramath Raj Sinha**
Printer & Publisher / CEO & Editorial Director (B2B Tech): **Vikas Gupta**
COO & Associate Publisher (B2B Tech): **Sachin Nandkishor Mhashilkar**

EDITORIAL
Group Editor: **R Giridhar**
Executive Editor: **Jatinder Singh**
Senior Correspondent & Editorial Coordinator–CISO Forum: **Jagrati Rakheja**
Principal Correspondent: **Musharrat Shahin**

DESIGN
Creative Director: **Shokeen Saifi**
Assistant Manager - Graphic Designer: **Manish Kumar**

SALES & MARKETING
Senior Director - B2B Tech: **Vandana Chauhan**
Head - Brand & Strategy: **Rajiv Pathak**
National Sales Head - B2B Tech: **Hafeez Shaikh**
Regional Sales Head - South: **Sourabh Dixit**
Senior Sales Manager - South: **Aanchal Gupta**

COMMUNITY ENGAGEMENT & DEVELOPMENT
Head - Databases: **Neelam Adhangale**
Head - Community Engagement: **Ekta Srivastava**
Senior Community Manager: **Vaishali Banerjee**
Senior Community Manager: **Reetu Pande**
Senior Community Manager: **Snehal Thosar**

OPERATIONS
General Manager - Events & Conferences: **Himanshu Kumar**
Senior Manager - Digital Operations: **Jagdish Bhainsora**
Manager - Events & Conferences: **Sampath Kumar**
Senior Producer: **Sunil Kumar**

PRODUCTION & LOGISTICS
Senior Manager - Operations: **Mahendra Kumar Singh**

Preface

The Security Reset

In 2026, cybersecurity is no longer a technical discipline focused on protecting systems—it is a leadership mandate focused on protecting enterprise trust, decision integrity, and regulatory credibility. As organizations accelerate AI adoption, expand across hybrid and multi-cloud environments, and digitize core operations, the attack surface has shifted from networks to identities, data flows, third-party dependencies, and algorithmic decision layers.

This CISO Priority Survey captures that inflection point through the views of senior infosec leaders operating at the sharp end of digital transformation. Their responses reveal a striking pattern: while boards are increasingly aware of cyber risk, many organizations remain structurally unprepared for the speed, scale, and systemic nature of today's threats. The challenge is no longer about detecting more attacks—it is about governing risk in environments where humans, machines, and AI agents act autonomously, often beyond traditional control boundaries.

For CISOs, this report is a mirror. It reflects how peers are reframing security from a cost center to a strategic capability that protects business continuity, brand trust, and regulatory standing. It surfaces where maturity is lagging—especially in recovery readiness, identity lifecycle governance, and AI risk enforcement—and where investment is moving toward foundational trust architecture rather than incremental tools. Most importantly, it arms CISOs with evidence to move conversations with the board from "security spend" to enterprise resilience, compliance defensibility, and decision velocity.

For boards and CxOs, this report is a governance wake-up call. Cyber and AI risks are now inseparable from business risk, regulatory exposure, and reputational resilience. The survey findings underscore that leadership misalignment—not technology gaps—is the biggest barrier to stronger security outcomes. Oversight must therefore evolve from periodic risk reviews to shared ownership of cyber and AI governance, with clear accountability for resilience, recovery readiness, and third-party exposure.

Read this report not as a snapshot of threats, but as a guide to governing trust at digital speed. The organizations that internalize these signals will not merely defend against disruption—they will build the structural resilience required to compete, comply, and lead in an AI-native economy.

As always, we welcome your feedback and comments.

Jatinder Singh
Editor
CISO Forum
Jatinder.singh@9dot9.in

R. Giridhar
Group Editor
9.9 Group
r.giridhar@9dot9.in

Introduction

The Security Priority Reset

CISOs reveal the real risk agenda for 2026—and what boards must do differently to govern cyber risk, AI risk, and business resilience

The CISO Priorities Survey 2026, based on insights from senior information security leaders across large and mid-sized Indian enterprises, reveals a profound shift in how cybersecurity is being defined, governed, and operationalized. This research captures a pivotal moment where traditional security paradigms are giving way to strategic imperatives driven by emerging technologies, regulatory evolution, and the fundamental reimagining of digital trust. Security is no longer anchored to infrastructure protection alone; it has become a core pillar of enterprise trust, regulatory defensibility, and decision integrity in an AI-native operating environment.

CISOs identify generative AI, identity sprawl, third-party dependencies, and multi-cloud complexity as the most consequential security exposures—signaling a move away from perimeter-centric risk models. The framing of cyber risk to business leadership has evolved in parallel: security is increasingly articulated in terms of trust erosion, regulatory consequences, operational continuity, and reputational risk, rather than technical vulnerability counts.

Investment priorities point to a strategic rebalancing—from accumulating tools to building foundational trust architecture, including identity governance, zero trust enforcement, cyber resilience, and third-party risk management. Threat perceptions further reinforce that systemic, supply-chain, and geopolitical risks now rival conventional cybercrime in board-level impact potential. Yet organizational barriers—notably leadership misalignment, budget prioritization conflicts, talent burnout, and tool sprawl—continue to constrain execution.

The maturity assessment highlights a persistent gap between detection and recovery readiness, exposing enterprises to prolonged business disruption when incidents occur. Meanwhile, the rapid expansion of generative AI has outpaced governance and enforcement capabilities, creating new decision-layer risks that most organizations are managing through policy more than operational controls. Regulatory pressure, particularly DPDP compliance, is reshaping security architecture choices, elevating data governance and access control from hygiene factors to strategic imperatives.

The survey data reveals both progress and persistent challenges. While 49% of respondents report that their boards now view cybersecurity as critical to enterprise resilience and brand trust, nearly half of organizations still struggle with fundamental capabilities like enterprise-wide cyber resilience planning and supply chain risk governance. This maturity gap, combined with ongoing challenges around talent shortages and executive support, underscores that the journey from tactical security operations to strategic risk management remains incomplete for many organizations.

The respondent profile—dominated by CISO-level leaders operating across regulated, asset-intensive, and large-scale enterprises—grounds these insights in real-world operational accountability. This is not theoretical security strategy; it is the

agenda of practitioners responsible for protecting enterprise continuity, credibility, and compliance at scale.

Taken together, the findings point to a clear mandate for 2026: cybersecurity must evolve from a control function to a governance discipline. Organizations that succeed will be those where boards co-own cyber and AI risk, invest in structural resilience, empower CISOs with architectural authority, and measure security not by the absence of incidents, but by the enterprise's ability to absorb shocks, recover trust, and sustain decision velocity in a high-risk digital economy.

The Convergence of Legacy and Emerging Threats

The survey shows a clear top-tier concern: Generative AI and Shadow AI use has overtaken traditional cloud or perimeter risks as the most pressing security domain for CISOs. This reflects the democratization of powerful AI capabilities and the resulting explosion of unsanctioned AI tool usage across enterprises.

information (hallucination) that could corrupt business processes or decision-making. So, there is a clear shift from infrastructure anxiety to decision-layer anxiety—where ungoverned AI usage can expose data, create regulatory violations, and embed bias into business workflows.

Identity & Access Management at scale

AI Overtakes Cloud as the Biggest Security Challenge

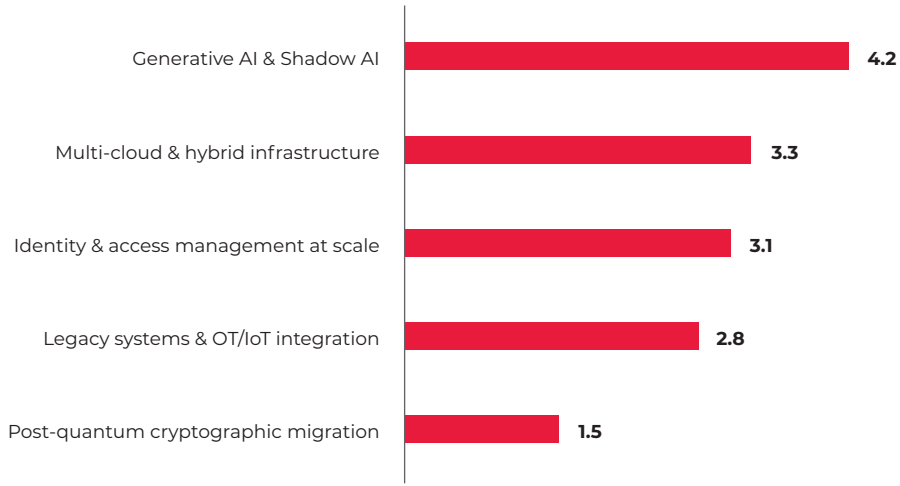

AI and identity sprawl are emerging as the new, ungoverned attack surface for enterprises. Leaders will need to consider AI and identity as core enterprise risk domains, not IT sub-projects

Unlike previous shadow IT challenges—where unauthorized cloud services were the primary concern—shadow AI introduces novel risks around data exfiltration, intellectual property leakage, and the generation of false but convincing

ranks consistently high, reinforcing the reality that identity—not networks—is now the security control plane of the enterprise. This underscores the persistent difficulty of managing digital identities in an era of remote work, third-party integrations, and API-driven

architectures. The proliferation of service accounts, machine identities, and temporary credentials has created identity sprawl that exceeds the capacity of many organizations to effectively monitor and govern.

Multi-cloud and hybrid infrastructure continue to rank high but are no longer novel risks; instead, CISOs are worried about control fragmentation, inconsistent policy enforcement, and operational blind spots across platforms. The abstraction layers that make cloud computing powerful—APIs, containers, serverless functions—also create new attack surfaces and configuration pitfalls. The challenge intensifies in hybrid environments where traditional perimeter-based controls coexist uneasily with cloud-native architectures, forcing security teams to master multiple control planes while maintaining consistent policies across heterogeneous environments.

Legacy systems and OT/IoT integration remain structurally risky—especially in manufacturing, BFSI, and utilities—where modernization is slow but exposure is real thanks to the growing imperative to connect operational technology to corporate networks.

While practical quantum computers capable of breaking current encryption remain years in the future, the harvest-now-decrypt-later threat means that adversaries are already collecting encrypted data for future decryption. For organizations handling sensitive intellectual property, financial records, or personally identifiable information with long-term value, the cryptographic clock has already started ticking. The challenge extends beyond technology—it encompasses inventory management, risk assessment, and the gradual replacement of cryptographic libraries across sprawling digital estates.

For now, post-quantum cryptography appears as a lower-ranked near-term priority, indicating that CISOs acknowledge the threat but see AI misuse and identity sprawl as more immediate business risks than cryptographic obsolescence. This prioritization reflects pragmatic risk management: CISOs are triaging what can damage enterprises in 12 to 24 months, not just what might break security architectures in five years.

> "Cybersecurity must evolve from a control function to a governance discipline."

What This Means?

- Treat AI governance as a board-level risk category, not an IT experiment
- Mandate identity governance as part of enterprise risk posture
- Implement unified security architecture across clouds—not tool sprawl

Boards Recognize Strategic Value of Cybersecurity

How boards perceive cybersecurity fundamentally shapes resource allocation, risk governance, and the organizational positioning of security leaders. Survey data suggests that board-level understanding has evolved significantly, though not uniformly, across Indian enterprises.

paradigms where security was primarily framed through compliance lenses, or treated as an operational cost center. Organizations in this category typically feature regular board-level security briefings, integrate cyber risk into enterprise risk management frameworks, and link security investments to business outcomes

Cyber Risk Is Now a Trust and Compliance Issue

The shift toward resilience-focused board engagement marks maturation of cyber risk governance, though one-quarter of organizations still frame security primarily through compliance or cost lenses.

Nearly half of respondents (49%) report that their boards view cybersecurity as a critical component of enterprise resilience and brand trust—the most sophisticated and strategically aligned perspective available in the survey. This represents a meaningful shift from earlier

rather than mere threat mitigation.

The second-largest cohort (21%) reports boards that view security primarily through a compliance and legal risk lens. While this perspective ensures regulatory obligations receive appropriate attention, it

can inadvertently limit security's strategic influence. Organizations in this category often excel at audit readiness and documentation but may underinvest in proactive threat hunting, advanced analytics, or security architecture innovation that lies outside regulatory mandates. The compliance-first mindset also tends to produce checkbox behaviors that satisfy auditors without necessarily reducing actual risk.

About a fifth (18%) of the survey respondents report boards that recognize security as a strategic enabler of digital transformation—a perspective that aligns security with business velocity, rather than viewing it as an impediment. These organizations typically involve security leaders in early-stage product development, M&A due diligence, and market expansion planning. Security enables rather than merely protects, facilitating faster cloud migrations, API-driven business models, and partner ecosystem integration through thoughtful control design.

However, a concerning 7% of respondents report boards that primarily view cybersecurity as an operational cost center with limited ROI visibility. In these organizations, security budgets compete directly with other cost centers, measured primarily on efficiency metrics rather than risk reduction or business enablement. This perspective typically correlates with reactive security programs, minimal executive engagement, and persistent underinvestment.

An additional 6% report boards that frame security primarily as a reputational and customer trust issue—not inherently negative, but potentially narrow in scope if it excludes operational resilience and strategic considerations.

CISOs are increasingly positioning cybersecurity not as a technical problem but as a business risk with reputational, customer trust, and regulatory consequences. This marks a maturation of the CISO narrative: risk is no longer about malware; it is about brand credibility, regulatory survival, and customer confidence. CISOs also increasingly see AI not just as a tool risk, but as a business logic risk—where flawed models can cause financial, legal, and ethical damage at scale. The idea of cyber risk as operational resilience (business continuity, uptime, recovery) is now mainstream, not aspirational.

What This Means?

- Cyber risk should be reported in terms of business impact (revenue, trust, compliance)
- AI governance should be tied to enterprise risk oversight
- Vendors should be held to shared security accountability

> "Boards view cybersecurity as a critical component of enterprise resilience and brand trust.

Investing in Enterprise-Wide Resilience

When asked to rank their organization's top cybersecurity priorities for 2026, respondents revealed a strategic agenda focused squarely on resilience, risk quantification, and regulatory adaptation. The most striking finding: building enterprise-wide cyber resilience emerged as the overwhelming top priority, with 44% ranking it first, followed by regulatory compliance.

This emphasis on resilience represents a philosophical evolution in security leadership thinking. Rather than pursuing the elusive goal of perfect prevention, organizations are embracing the reality that breaches will occur, and are focusing instead on minimizing impact through rapid detection, coordinated response, and swift recovery. Resilience encompasses business continuity planning, incident response capabilities, supply chain redundancy, and the organizational muscle preparation required to navigate crisis scenarios without catastrophic business disruption.

Investment priorities reveal a strong pivot away from “buying more tools” toward architectural resilience and governance maturity. Identity security, zero-trust enforcement, and third-party risk platforms dominate funding intent,

CISO Priorities Have Evolved from Protection to Enterprise Resilience

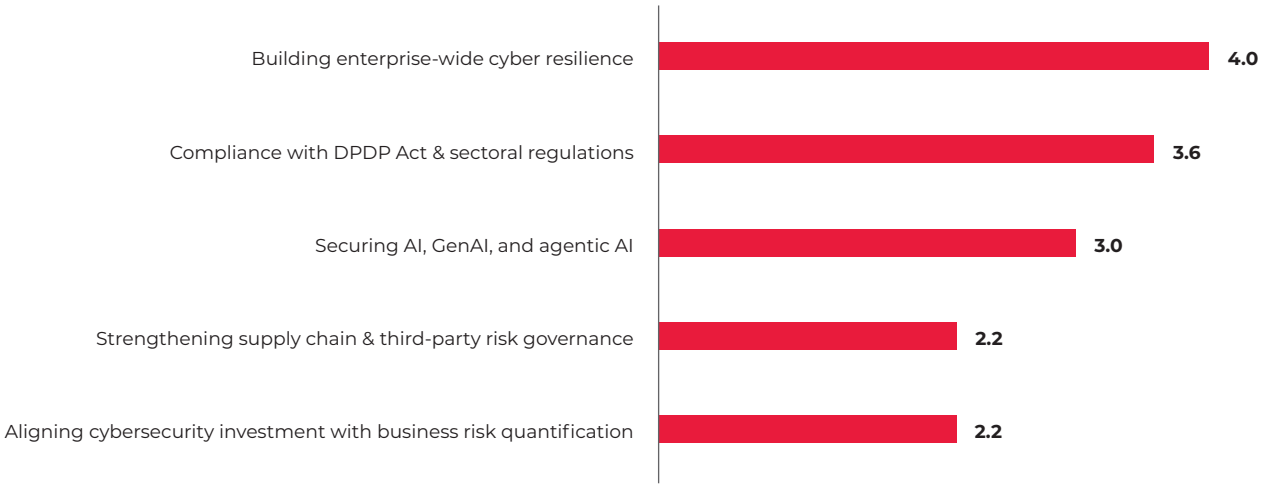

The focus is shifting from preventing incidents to building enterprise resilience and recovery readiness, with compliance and AI security forming the strategic triad for 2026.

indicating CISOs want to reduce structural risk rather than chase tactical alerts. The prioritization of enterprise-wide resilience over point-solution security represents a strategic shift that boards must actively support with appropriate resources and cross-functional mandates. Boards should also verify that resilience programs include regular testing through tabletop exercises and simulations, not just documentation.

Compliance with the Digital Personal Data Protection Act and sectoral → regulations claimed 39% of first-place rankings, underscoring the profound impact of India's evolving regulatory landscape. The DPDP Act has fundamentally altered the risk calculus for Indian enterprises, introducing potential penalties that can reach up to ₹250 crore per instance and creating personal liability for data fiduciaries. Unlike previous compliance regimes that could be satisfied through documentation and policy frameworks, DPDP Act compliance requires operational changes in data collection, processing, storage, and cross-border transfer practices.

Securing AI, GenAI, and agentic AI ranked third with 11% selecting it as their top priority. While this may seem modest compared to resilience and compliance, it represents remarkable prioritization for an emerging technology domain. The AI security challenge extends beyond traditional application security to encompass prompt injection attacks, model poisoning, training data provenance, and the novel risks introduced by autonomous AI agents that can take actions without human oversight. AI security platforms—especially for GenAI monitoring, data leakage prevention, and model governance—are emerging as a new budget category. This reflects a realization that AI introduces continuous risk, not episodic vulnerability.

Strengthening supply chain and third-party risk governance came at fourth place, acknowledging that enterprise security perimeters now extend deep into partner ecosystems and vendor relationships. However, the relatively low prioritization of supply chain risk governance is concerning given the demonstrated impact of recent third-party breaches. Boards should question whether their organizations truly understand the security posture of critical vendors, and if contracts include adequate security requirements, audit rights, and liability provisions.

Aligning cybersecurity investment with business risk quantification rounded out the top five, reflecting the growing demand for security leaders to articulate cyber risk in business terms rather than technical metrics. This priority signals recognition that security's credibility and influence depend on demonstrating clear connections between security investments and quantifiable risk reduction. The message is clear: security architecture is now a business efficiency lever, not just a defense mechanism.

What This Means?

- Fund architectural modernization, not tactical patchwork
- Implement AI risk controls alongside AI investments
- Treat cyber resilience as business continuity insurance

Ransomware and Human Risk Dominate Concerns

Security leaders rated nine distinct threat vectors on a five-point severity scale, revealing a threat landscape where ransomware operations and human risk factors command overwhelming executive attention. The data underscores a sobering reality: despite years of awareness and investment, organizations continue to face existential threats from professionalized ransomware ecosystems, while the human element remains the weakest link in enterprise security architectures.

Ransomware with multi-layered extortion emerged as the most severe threat by a commanding margin, earning the highest rating with 79% of respondents rating it as high or extremely severe—nearly ten percentage points above the second-ranked threat. This finding reflects ransomware's evolution from opportunistic malware to sophisticated, industrialized operations that combine data exfiltration with encryption in devastating double-extortion attacks. Modern ransomware gangs operate as professional criminal enterprises with customer service departments, negotiation protocols, and affiliate programs that enable attackers of varying skill levels to deploy destructive payloads.

Extortion, Human Risk, and Emerging Technologies Lead Concerns

The threat landscape shows clear stratification with ransomware and human-centric risks representing the most immediate and severe organizational risks.

The multi-layered extortion model has fundamentally changed the ransomware calculus for organizations. Beyond traditional encryption that disrupts operations, adversaries now threaten public disclosure of exfiltrated data, leading to impossible choices: pay ransoms that fund criminal ecosystems and provide no guarantee of data deletion, suffer extended downtime that threatens business viability, or accept public disclosure of sensitive data that damages customer trust and regulatory standing. The average ransom demand has escalated into millions of dollars, while total recovery costs—including forensic investigation, system rebuilding, legal fees, regulatory fines, and business disruption—often exceed the ransom itself by factors of ten or more. For many organizations, a successful ransomware attack represents an extinction-level event that forces fundamental questions about business continuity and organizational viability.

Insider threats and human risk earned the second-highest severity rating, with 71% of respondents rating them as high or extremely severe. This persistent vulnerability reflects that humans remain the most exploitable component of enterprise security despite substantial investment in awareness programs and technical controls. The insider threat encompasses both malicious actors—disgruntled employees seeking revenge or financial gain, corporate espionage, fraud—and well-intentioned users whose actions inadvertently create risk through configuration errors, phishing susceptibility, or policy violations. The challenge has intensified with remote work arrangements, BYOD policies, and cloud-based collaboration tools that blur traditional workplace boundaries and make monitoring more complex. The rise of sophisticated social engineering—including AI-generated deepfakes and emotionally manipulative pretexting—has made even security-aware users vulnerable to manipulation.

AI-driven attacks ranked third, with 70% rating them high or extremely severe. This concern reflects the democratization of offensive capabilities—sophisticated techniques once limited to nation-state actors are now accessible through automated tools and AI-assisted exploitation frameworks. Adversaries leverage AI for enhanced spear-phishing campaigns that adapt messaging based on target behavior, deepfake-enabled social engineering that impersonates executives with disturbing accuracy, automated vulnerability discovery that finds zero-day exploits faster than defenders can patch, and adaptive malware that evolves to evade signature-based detection.

Identity-centric breaches secured fourth place with 68% severity ratings, underscoring that credentials have become the primary attack vector in modern breaches. Rather than breaking through firewalls, adversaries increasingly steal or abuse legitimate credentials to walk through the front door. The proliferation of cloud services, API keys, service accounts, and machine identities has created vast credential sprawl that exceeds many organizations' ability to effectively monitor and govern. The traditional perimeter has dissolved into thousands of individual access decisions, each representing potential compromise.

Cloud and API misconfigurations ranked fifth, reflecting the security debt accumulated through rapid cloud adoption. The flexibility and abstraction that make cloud computing powerful also create countless opportunities

for misconfiguration—overly permissive S3 buckets, exposed management interfaces, unsecured APIs, and excessive IAM permissions.

Supply chain attacks rounded out the top concerns, with 58% high/extreme severity ratings acknowledging that trusted vendor relationships now represent critical attack vectors, as demonstrated by incidents like SolarWinds and the 3CX supply chain compromise. Cyber insurance uncertainty also features as a growing worry, with CISOs concerned about shrinking coverage, claim exclusions, and rising premiums—making cyber risk increasingly uninsurable without strong governance maturity.

The ransomware finding demands immediate board-level attention to organizational resilience and crisis preparedness. Boards should verify that their organizations maintain offline, immutable backups that can enable recovery without ransom payment, conduct quarterly ransomware simulation exercises that test cross-functional response coordination, and pre-negotiate relationships with forensic firms, legal counsel, and crisis communications specialists who can mobilize immediately following an incident. Boards must also scrutinize cyber insurance policies to ensure they provide adequate coverage for ransomware incidents, including business interruption, data restoration costs, and regulatory penalties. Many policies contain significant exclusions or sub-limits that leave organizations exposed.

The high severity of insider threats and human risk requires board focus on organizational culture, not just technical controls. Security awareness must evolve beyond annual compliance training to include regular phishing simulations, role-based education, and clear consequences for policy violations. The convergence of AI-driven attacks with traditional threats underscores the need for continuous investment in threat detection and response capabilities that can identify and contain rapidly evolving attack techniques.

What This Means?

- Include geopolitical cyber risk in enterprise risk registers
- Demand vendor risk audits and breach transparency
- Treat misconfiguration risk as an operational governance failure

> "Ransomware with multi-layered extortion emerged as the most severe threat."

The Human Capital Crisis in Cybersecurity

Beyond external threats, security leaders face a constellation of internal challenges that constrain their ability to build effective security programs. When asked to rate internal obstacles to improving cyber resilience, respondents revealed that organizational and human factors pose greater barriers than technology limitations.

The shortage of skilled talent has emerged as the most severe internal challenge, with 50% rating it high or extremely severe. This finding reflects a cybersecurity workforce gap that the industry has struggled to close despite a decade of awareness and education initiatives. The challenge extends beyond raw headcount—organizations need specialists in emerging domains like cloud security architecture, AI governance, OT/IoT security, and threat intelligence analysis. Traditional hiring approaches are proving inadequate when the labor pool cannot meet demand, forcing organizations toward alternative strategies like automation, managed security services, and cross-training programs.

End-user education and training ranked second, with 49% severity ratings acknowledging that employee security awareness remains a persistent weakness.

Talent and Training Are the Real Security Bottlenecks

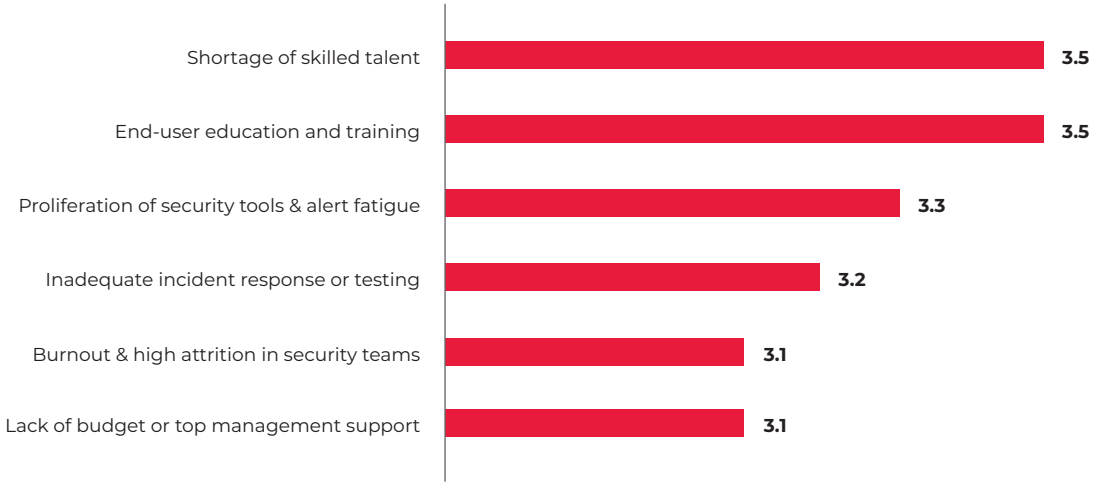

Talent shortage surpassing budget constraints as the top challenge marks a fundamental shift—organizations can't hire their way out of the cybersecurity workforce gap.

Despite substantial investment in security awareness programs, the data suggests these initiatives have not achieved their intended impact. The challenge may lie in delivery methodology—annual compliance training modules fail to create lasting behavioral change, or to prepare users for sophisticated social engineering attacks that exploit current events, organizational relationships, and emotional triggers.

Tool proliferation and alert fatigue claimed third position, reflecting the unintended consequences of point-solution security strategies. Many organizations have accumulated dozens of security tools over years of incremental acquisition, creating operational complexity, integration challenges, and alert volumes that exceed analyst capacity. The resulting alert fatigue leads to missed critical warnings buried in noise, analyst burnout, and inefficient workflows where teams spend more time managing tools than investigating threats.

Inadequate incident response testing ranked fourth, with 48% severity ratings exposing a troubling gap between documented incident response plans and operational readiness. Many organizations maintain comprehensive playbooks that have never been tested under realistic conditions, leaving them unprepared for the chaos, time pressure, and cross-functional coordination required during actual incidents.

Burnout and attrition in security teams rounded out the top tier at 39% severity, reflecting the human cost of understaffing, alert fatigue, and the perpetual defender's disadvantage where adversaries need only succeed once while defenders must succeed constantly.

The talent crisis requires board-level workforce strategy beyond incremental hiring. Boards should look at whether their organizations have developed non-traditional talent pipelines—career conversion programs, partnerships with technical training providers, and creative compensation structures that acknowledge market realities. The relatively low severity rating for budget constraints (36%) compared to talent shortages (50%) suggests that money alone cannot solve the workforce challenge. Boards should focus on retention through career development, manageable workloads, and automation that reduces repetitive tasks.

Inadequate incident simulation and testing indicates that resilience is still aspirational in many organizations. Boards should mandate quarterly tabletop exercises that include executive participation, not just security team drills. These exercises should simulate realistic scenarios—ransomware with data exfiltration, supply chain compromise, insider threat—and test not just technical response but crisis communications, legal notifications, and business continuity activation. Organizations that discover response plan inadequacies during exercises are far better positioned than those who discover them during actual breaches.

What This Means?

- Treat cybersecurity as a leadership problem, not a technology gap
- Fund technology simplification and talent retention
- Mandate resilience testing as a board KPI

The Operational Readiness Gap

Security leaders rated their organization's maturity across seven critical capabilities, using a five-point scale from immature/ad-hoc to optimized/ continuously improving. The survey reveals a clear maturity gap between strategic intent and operational reality. While most respondents rate their governance frameworks and policy coverage as "defined" or "improving," far fewer report mature capabilities in continuous monitoring, identity lifecycle governance, and automated response. This gap reflects a common enterprise pattern: organizations invest heavily in policy and compliance signaling, but under-invest in the operational muscle required to enforce security at scale.

Detection capabilities appear stronger than response and recovery capabilities—suggesting that many organizations can identify incidents faster than they can contain and recover from them. However, this imbalance is dangerous in an era of AI-enabled attacks, where speed of exploitation is accelerating. Resilience maturity—incident simulations, crisis playbooks, and recovery orchestration—lags behind perimeter and detection investments, indicating that many enterprises remain prevention-biased rather than resilience-ready.

The most telling signal is the uneven maturity

Excelling at Today's Challenges, Unprepared for Tomorrow's Threats

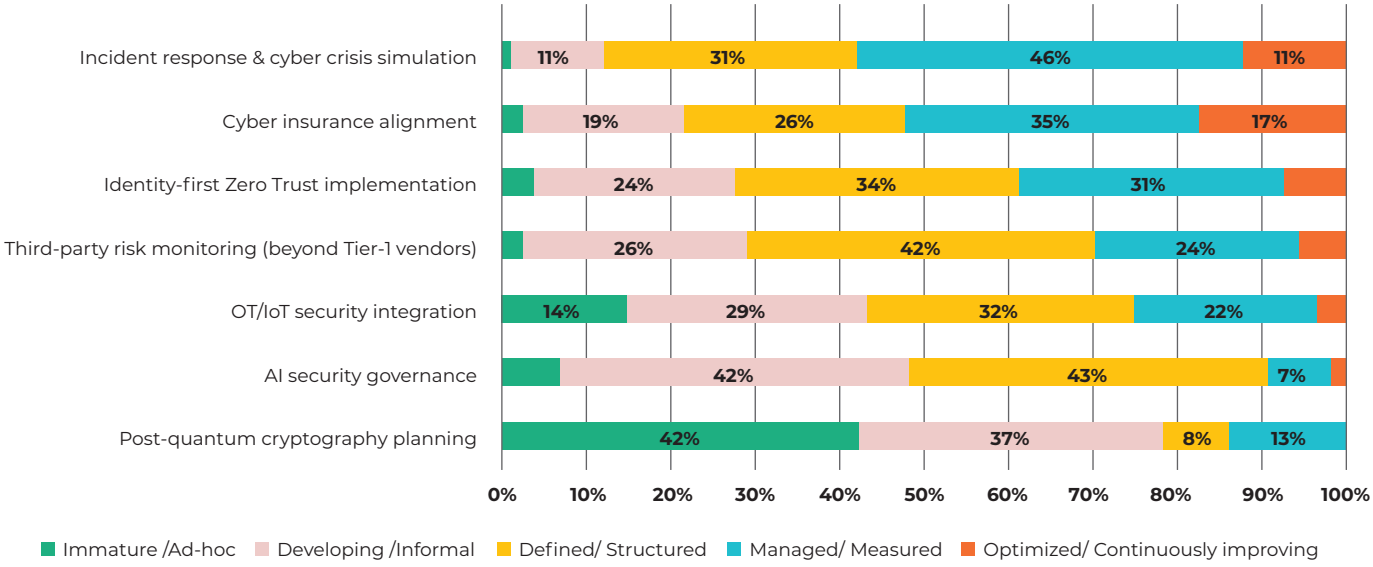

Organizations demonstrate they can build sophisticated capabilities when mandated, yet remain critically unprepared for emerging quantum and AI-driven threats they already recognize as significant risks.

CISOFORUM

of identity governance. While IAM tooling is widespread, lifecycle automation, entitlement hygiene, and continuous access evaluation remain weak. In an AI-driven enterprise where agents, bots, and service identities now outnumber humans, this gap becomes a systemic risk.

AI security governance ranked second-lowest in maturity, with only 8% reporting optimized programs. This finding aligns with the nascent state of AI security frameworks and the rapid evolution of both AI capabilities and associated risks. Unlike mature security domains with established standards and proven practices, AI security governance requires organizations to develop novel controls for prompt injection, model poisoning, training data integrity, and autonomous agent oversight—challenges for which industry best practices remain emergent. The low maturity of AI security governance creates immediate risk given the proliferation of shadow AI usage across enterprises.

OT/IoT security integration scored marginally higher, reflecting progress in industrial cybersecurity but persistent challenges in unifying IT and OT security paradigms. Incident response and cyber crisis simulation achieved the highest maturity rating, with 57% reporting managed or optimized capabilities. This relative strength likely reflects the concrete, testable nature of incident response compared to more abstract capabilities like AI governance or quantum readiness. Organizations can readily measure response times, test procedures through drills, and iteratively improve based on lessons learned.

Cyber insurance alignment, Zero Trust implementation, and third-party risk monitoring occupy the middle tier of maturity, all scoring between 3.0 and 3.5. Third-party risk monitoring particularly suffers from the difficulty of obtaining security visibility beyond tier-one vendors into the extended supply chain.

Post-quantum cryptography planning received the lowest maturity score of any measured capability, with 42% of organizations describing their approach as immature or ad-hoc. Only 13% report managed or optimized programs, while the remaining 45% fall somewhere in the developing-to-defined range. This maturity deficit is particularly concerning given that standards bodies have published approved quantum-resistant algorithms, and migration timelines extend across years, not months. Organizations delaying quantum readiness planning risk discovering cryptographic dependencies at the worst possible moment—under time pressure when quantum threats materialize.

The relatively strong incident response maturity offers a silver lining—organizations have demonstrated they can build sophisticated capabilities when properly resourced and mandated. Boards should leverage this proof point to drive similar maturity improvements in quantum readiness and AI governance before these domains transition from emerging concerns to active threat vectors.

What This Means?

- Implement maturity benchmarks, not just tool counts
- Prioritize recovery readiness and simulation exercises
- Treat identity lifecycle governance as a business risk control

Securing AI: From Ad-Hoc to Strategic Integration

With artificial intelligence deployed across enterprises at unprecedented scale, the question of how organizations secure these deployments takes on critical importance. Survey responses reveal a spectrum of approaches ranging from ad-hoc security reviews to comprehensive AI-native security frameworks, with the majority of organizations clustering around periodic assessment models rather than continuous, integrated governance.

One-time security reviews before deployment emerged as the most common approach, adopted by 33% of respondents. This model treats AI systems similarly to traditional application deployments—conducting security assessments before production release but lacking ongoing monitoring or reassessment as models evolve, training data changes, or new attack vectors emerge. While better than no security review, this approach fails to account for the dynamic nature of AI systems, where model behavior can drift over time, adversarial techniques continuously evolve, and integration points multiply as systems mature.

Periodic AI risk assessments ranked second at 28%, representing a more mature approach that recognizes AI security as an ongoing concern rather than a one-time gate. However,

AI Governance Exists but Enforcement Lags

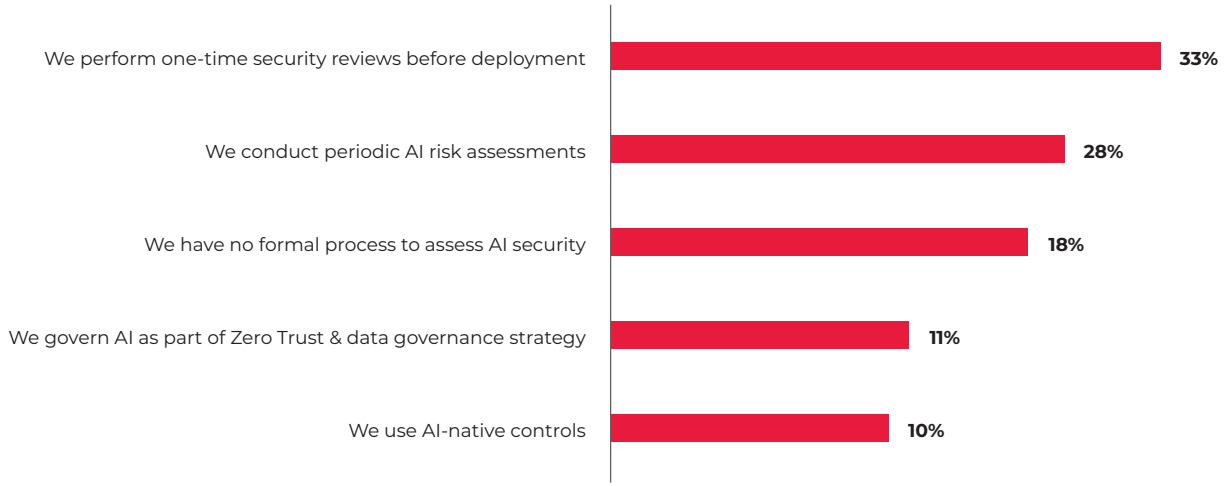

AI usage policies are widespread, but enforcement and auditability remain inconsistent. The prevalence of one-time security reviews suggests most organizations treat AI like traditional software, missing the unique risks of adaptive, probabilistic systems.

the periodic nature still leaves gaps between assessments where new vulnerabilities or usage patterns may emerge undetected. The approach also tends toward compliance-oriented checkbox exercises rather than continuous risk management integrated into development and operations workflows.

A concerning 18% acknowledge having no formal process to assess AI security—a significant governance gap given the pace of AI adoption and the novel risks these systems introduce. This cohort likely includes organizations where AI deployment has outpaced security program maturity, creating shadow AI implementations that operate outside formal governance frameworks. The absence of assessment processes leaves organizations blind to risks around data handling, and adversarial attacks.

Eleven percent report governing AI as part of Zero Trust and data governance strategies—an approach that integrates AI security into broader security architectures rather than treating it as a separate domain. This methodology applies Zero Trust principles of explicit verification, least privilege, and assume breach to AI systems while leveraging existing data classification, access controls, and monitoring capabilities.

Finally, 10% have adopted AI-native controls—specialized security mechanisms designed specifically for AI/ML systems, such as adversarial robustness testing, model explainability frameworks, and techniques to detect training data poisoning or model extraction attempts.

While policies for acceptable GenAI usage exist in many organizations, enforcement mechanisms—such as prompt monitoring, data leakage prevention, model governance, and audit trails—are far less mature. This creates a governance illusion: AI appears "controlled" on paper but remains operationally porous. Given that employees across organizations are already using AI tools—both sanctioned and unsanctioned—the absence of governance creates uncontrolled risk around data handling, intellectual property protection, and regulatory compliance.

The prevalence of one-time security reviews reflects a fundamental misunderstanding of AI systems. Unlike traditional software, where code remains static between releases, AI models exhibit emergent behaviors, require ongoing retraining, and face continuously evolving adversarial techniques.While AI security features prominently in investment and risk discussions, the maturity of AI governance frameworks significantly lags adoption—revealing a growing gap between how strategically important AI is perceived and how rigorously it is governed in practice.

The relatively small adoption of AI-native controls (10%) and Zero Trust integration (11%) suggests that most organizations are applying traditional security frameworks to fundamentally different technology. Boards should evaluate whether their security teams have developed specialized expertise in AI security, or need to partner with vendors offering AI-native security solutions.

What This Means?

- Mandate AI governance structures with named accountability
- Require auditability for AI decisions affecting customers or compliance
- Treat AI misuse as a material enterprise risk

DPDP Enforcement Drives Strategic Planning

Regulatory frameworks are increasingly shaping cybersecurity strategies, forcing organizations to balance compliance obligations with operational flexibility and innovation velocity. When security leaders rated how various regulatory factors would impact their 2026 strategies, a clear hierarchy emerged with domestic data protection regulation commanding greatest attention while global frameworks receive more modest focus.

Digital Personal Data Protection Act enforcement earned the highest impact rating, with 60% of respondents rating it as high or extremely high impact on their 2026 planning. This finding reflects the DPDP Act's potential to fundamentally reshape how Indian organizations handle personal data. The law introduces consent requirements that may force redesign of customer interaction flows, data minimization principles that challenge existing retention practices, and breach notification obligations that necessitate enhanced detection and response capabilities. The phased enforcement approach means organizations face ongoing uncertainty about timing and interpretation, complicating planning cycles.

Evolving cyber insurance requirements ranked second at 50% high/extreme impact,

Compliance Is Shaping Security Architecture Decisions

Regulatory exposure is now an operational risk. Leaders need to focus on the modernization of compliance-driven architecture.

reflecting how insurers are increasingly driving security improvements through policy requirements, premium adjustments, and coverage exclusions. Organizations seeking cyber insurance coverage face detailed security questionnaires, must demonstrate specific controls like multi-factor authentication and endpoint protection, and increasingly confront exclusions for ransomware payments or nation-state attacks. The insurance market's hardening has transformed coverage from a passive risk transfer mechanism into an active driver of security investment decisions.

India AI Governance Guidelines claimed third position, though with more modest 29% high/extreme ratings. These guidelines establish principles for responsible AI development and deployment but lack the enforcement mechanisms and specific requirements of formal regulations. Organizations view them as directionally important but less immediately constraining than DPDP Act obligations.

Sector-specific mandates from RBI, IRDAI, and other regulators scored 40% high/extreme impact, reflecting the reality that regulated industries face layered compliance obligations extending beyond general data protection requirements.

Regulatory pressure—especially DPDP compliance—has become a primary driver of security investment prioritization. CISOs report that regulatory readiness now shapes funding approvals, architecture decisions, and even vendor selection. This marks a shift from reactive compliance to compliance-led architecture design, where data classification, access controls, and breach response are increasingly mapped to regulatory obligations.

However, the survey suggests that many organizations still view compliance as a reporting exercise rather than an operational discipline. Documentation readiness outpaces technical readiness, creating exposure in the event of audits or breaches. Sectoral regulations in BFSI, healthcare, and critical infrastructure further amplify this pressure, forcing CISOs to balance innovation with regulatory defensibility.

The convergence of AI governance with DPDP compliance is emerging as a new complexity frontier. Enterprises deploying GenAI over sensitive datasets face regulatory uncertainty around consent, explainability, and accountability. However, many CISOs see regulation not as constraint, but as leverage to modernize data governance and identity controls.

What This Means?

- Treat DPDP readiness as operational risk, not documentation hygiene
- Align AI programs with regulatory accountability frameworks
- Implement regulatory stress tests for breach scenarios

> "DPDP compliance has become a primary driver of security investment prioritization."

Investment Intent in Automation and Identity Security Surge

Where organizations choose to invest reveals strategic priorities more directly than stated intentions. When asked to rank six cybersecurity capabilities by investment focus for 2026, security leaders painted a clear picture: AI-powered automation and Zero Trust identity architectures dominate spending plans while quantum readiness receives minimal allocation despite being recognized as a critical long-term threat.

AI-powered SOC automation and threat detection emerged as the overwhelming investment priority, with 42% ranking it as their top focus. This finding reflects multiple converging pressures: the talent shortage that makes analyst productivity critical, the alert fatigue created by tool proliferation, and the proven efficacy of machine learning for threat hunting and anomaly detection. The investment encompasses SOAR platforms, advanced analytics, threat intelligence integration, and automated response orchestration.

Identity-first Zero Trust architecture claimed second place with 31% first-rank selections, reflecting the strategic shift toward identity as the new security perimeter. Zero Trust

Security Investments Are Flowing to Automation and Trust Architecture

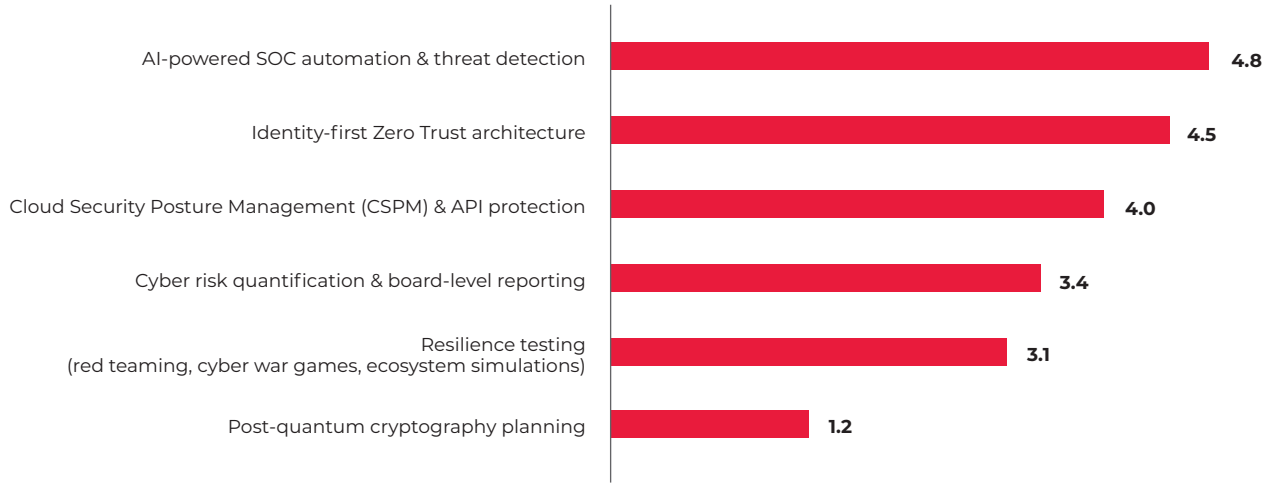

Investment focus is shifting from detection tools to trust architecture and recovery foundations that deliver more resilience.

principles—verify explicitly, use least privilege access, assume breach—require comprehensive identity and access management infrastructure including multi-factor authentication, privileged access management, identity governance, and continuous authentication mechanisms. The investment priority aligns with threat landscape findings where identity-centric breaches ranked among the most severe concerns. Organizations recognize that traditional network security models provide insufficient protection in cloud-centric, API-driven architectures.

Cloud Security Posture Management and API protection ranked third at 10%, representing continued investment in securing multi-cloud environments. CSPM tools provide visibility into cloud configurations, detect misconfigurations, enforce policy compliance, and integrate with DevOps workflows to prevent security drift. The API security component reflects recognition that APIs have become primary attack surfaces in digital business models.

Cyber risk quantification and board reporting secured fourth position, indicating growing investment in translating technical security metrics into business risk language that enables better resource allocation decisions. Resilience testing—red teaming, cyber war games, and ecosystem simulations—ranked fifth with 8% top rankings. While recognized as valuable for validating response capabilities, these activities receive lower investment priority compared to tool-based security capabilities.

Post-quantum cryptography planning received the lowest investment ranking, with respondents betting that quantum threats remain sufficiently distant to defer investment, or that they lack understanding of how to translate quantum readiness into actionable initiatives.

The capability investment pattern reflects a strategic pivot toward foundational controls. Identity governance, zero-trust enforcement, third-party risk platforms, and cyber resilience tooling dominate prioritization—signaling a move away from point solutions toward architectural cohesion. The investment concentration in AI-powered automation and Zero Trust makes strategic sense given demonstrated efficacy and clear ROI.

The emphasis on automation investments aligns well with talent shortage realities, but leaders should verify that automation initiatives include change management, analyst retraining, and process redesign—not just tool procurement. Many automation projects fail to deliver expected productivity gains because organizations automate broken processes rather than redesigning workflows.

The significant investment in cyber risk quantification deserves board support as it promises to improve decision-making quality by translating technical security metrics into business risk language. However, boards should ensure these initiatives produce actionable insights rather than just sophisticated dashboards. Effective risk quantification should inform resource allocation, enable risk-return tradeoff analysis, and facilitate meaningful comparison of security investments against other enterprise risk mitigation strategies.

What This Means?

- Fund architectural simplification over tool accumulation
- Treat identity and third-party risk as systemic controls
- Invest in automation to protect security team sustainability

From Defending Systems to Governing Trust

The 2026 CISO Priorities Survey captures Indian cybersecurity leadership at a critical juncture. Traditional security paradigms centered on perimeter defense and compliance checklists are giving way to strategic imperatives around resilience, identity-centric security, and the governance of transformative technologies like artificial intelligence. The data reveals both encouraging maturity in areas like incident response, and troubling gaps in emerging domains like quantum readiness and AI security governance. Across technology risk domains, business risk framing, investment priorities, threat perceptions, and organizational barriers, CISOs are signaling that cybersecurity has moved beyond perimeter defense into the realm of enterprise trust governance.

Several findings demand immediate executive and board attention. The maturity assessment exposes a persistent gap between security ambition and operational readiness—particularly in recovery, identity lifecycle governance, and continuous enforcement. With 18% of organizations lacking formal AI security assessment processes and employees across enterprises already using AI tools both sanctioned and unsanctioned, the governance vacuum creates uncontrolled risk. Organizations cannot afford to treat AI security as a future concern when AI deployment is a present reality. Leaders must establish governance frameworks, approved tool lists, and security review requirements before shadow AI implementations create irreversible data handling or intellectual property risks.

The talent shortage surpassing budget constraints as the top internal challenge marks a fundamental shift in security program challenges. Organizations can no longer rely primarily on hiring to build capability. Instead, they must embrace automation, develop non-traditional talent pipelines, prioritize retention through career development and workload management, and partner strategically with managed security service providers. The implication for boards: talent strategy deserves the same attention as technology investment, and compensation structures must acknowledge market realities rather than internal equity concerns.

The dominance of identity-centric threats and insider risk in the threat landscape validates the strategic shift toward Zero Trust architectures and identity governance. Organizations investing heavily in Zero Trust implementation and AI-powered automation align their spending with demonstrated threat vectors. However, this focus must not come at the expense of resilience capabilities. Building enterprise-wide cyber resilience—rated the top priority by 44% of respondents—requires more than technical controls. It demands cross-functional coordination, regular testing through realistic simulations, and organizational muscle memory that comes only through practice.

The regulatory landscape, particularly DPDP Act enforcement, continues to reshape security strategies in ways that extend far beyond compliance departments. The personal liability provisions and substantial penalty regime mean that data protection has become a matter of director duty, not just corporate

risk management. Boards must establish governance mechanisms appropriate to this reality—regular reporting on data flows, consent management practices, and breach preparedness, with clear escalation paths and decision authority.

Looking forward, successful security leaders will be those who can simultaneously manage present operational demands while preparing for strategic shifts in the threat landscape. This requires balancing investments between proven capabilities that deliver immediate risk reduction—AI-powered automation, Zero Trust, cloud security—and longer-term readiness initiatives around quantum cryptography and AI governance that may not show returns for years but cannot be deferred indefinitely.

The capability investment priorities reinforce this shift: CISOs are deliberately moving away from tool accumulation toward architectural coherence—identity-first security, third-party risk governance, zero trust enforcement, and cyber resilience.

Finally, the survey grounds these insights in the lived reality of large, complex Indian enterprises operating across regulated and high-risk sectors. This is not abstract security theory—it is the agenda of practitioners accountable for enterprise resilience at scale.

The overarching message to boards is unambiguous: cybersecurity in 2026 is no longer about asking whether the enterprise is secure, but whether it is governable, resilient, and trusted at digital speed. The organizations that thrive will be those where boards co-own cyber and AI risk, invest in foundational trust architecture, and empower CISOs not just as defenders of systems—but as governors of enterprise decision integrity.

Appendix A

Survey Methodology

The CISO Priorities survey reflects the views of senior cybersecurity leaders from India, primarily CISOs, Heads of Security, and leadership roles with enterprise-wide security responsibility. Most respondents operate at national or regional leadership scope, with representation across BFSI, manufacturing, IT services, healthcare, infrastructure, and digital-native enterprises. The respondent profile reflects decision-makers who directly influence enterprise security strategy, regulatory posture, and cyber investment allocation.

Leadership Profile: The majority of respondents (65%) identify as executive leadership with strategic oversight responsibilities, indicating that this research reflects board-level perspectives rather than purely operational viewpoints. Another 14% lead cross-functional technology initiatives that blend cybersecurity with broader IT and digital transformation mandates, highlighting the

Enterprise Strategy Is Being Shaped by Cybersecurity Decision Makers

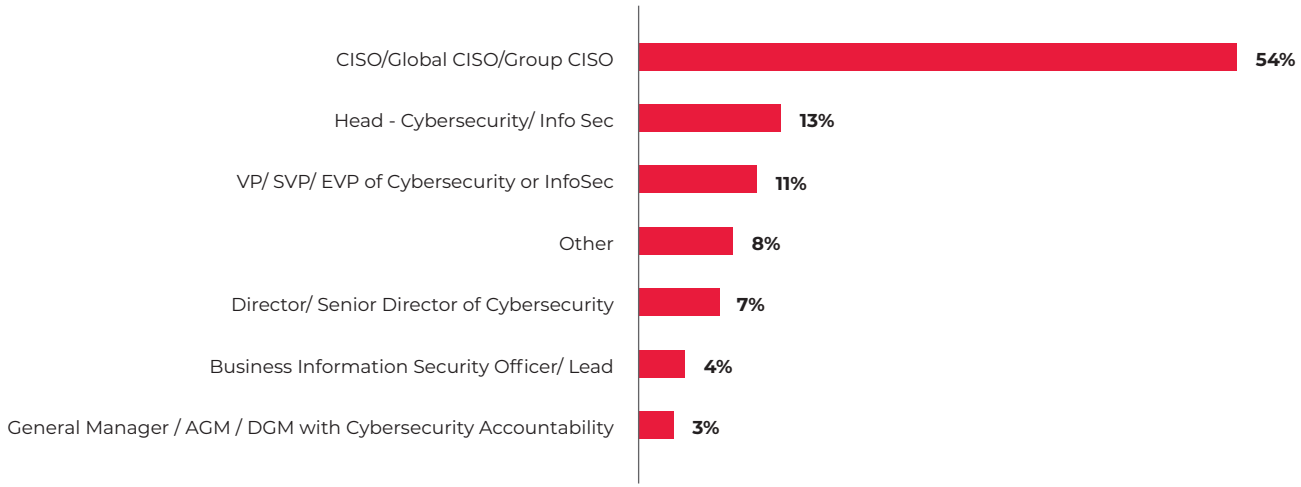

The survey reflects perspectives of leaders with direct ownership of enterprise security outcomes. The insights reflect CISO-level accountability, not theoretical viewpoints.

Most Security Leaders Operate with Enterprise-Wide Mandates

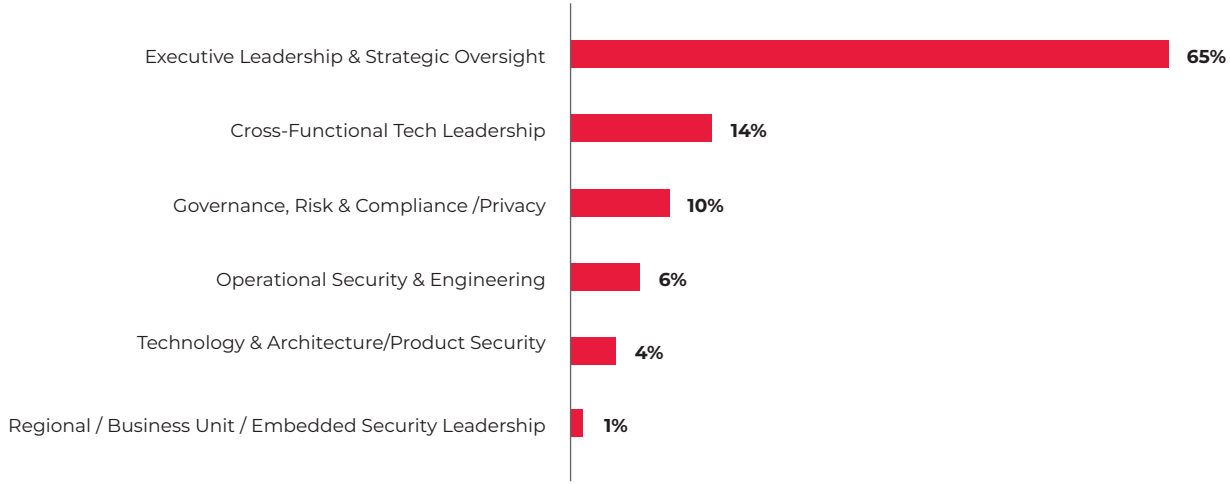

Security leadership is increasingly centralized at the enterprise level. Boards need to ensure that CISOs have authority commensurate with their risk accountability.

increasing integration of security into business strategy.

Geographic Scope: While almost half (46%) of respondents lead security for India-only operations, 28% manage India plus select international markets,

Cyber Risk Leadership Is Expanding from Local to Multi-Region

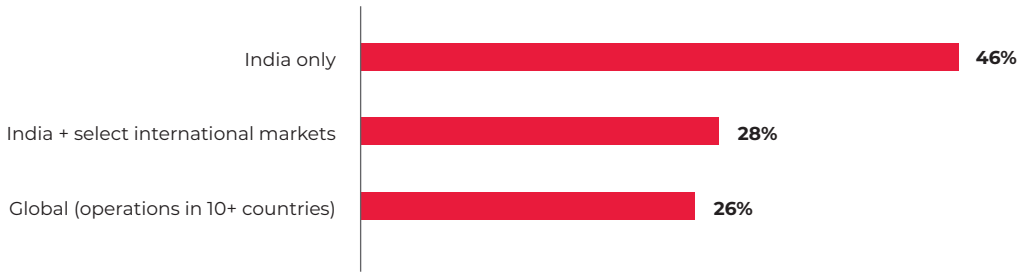

Cyber risk is now being managed at regional and multi-geography scale, not just locally. However, cross-border operations amplify regulatory and breach exposure.

and 26% oversee global security programs spanning 10 or more countries.

Industry: Banking and financial services represent the largest cohort (28%), followed by manufacturing (26%), and IT services (17%). This distribution aligns with India's economic composition and the sectors facing the most acute cybersecurity pressures from regulatory mandates and operational technology integration challenges.

Organizational Scale: Respondents represent organizations that span mid-market to large enterprises, with revenue bands skewed toward ₹2,500 crore+ annual turnover, ensuring that the findings reflect security priorities at scale. One quarter of respondents represent organizations with annual revenues exceeding ₹25,000 crore, while another 22% come from organizations with revenues below ₹1,000 crore, ensuring the findings reflect challenges across organizational scales. The diversity of sectors and enterprise sizes provides a realistic view of how cybersecurity priorities differ across regulated, asset-heavy, and digital-first organizations.

Industry Sector

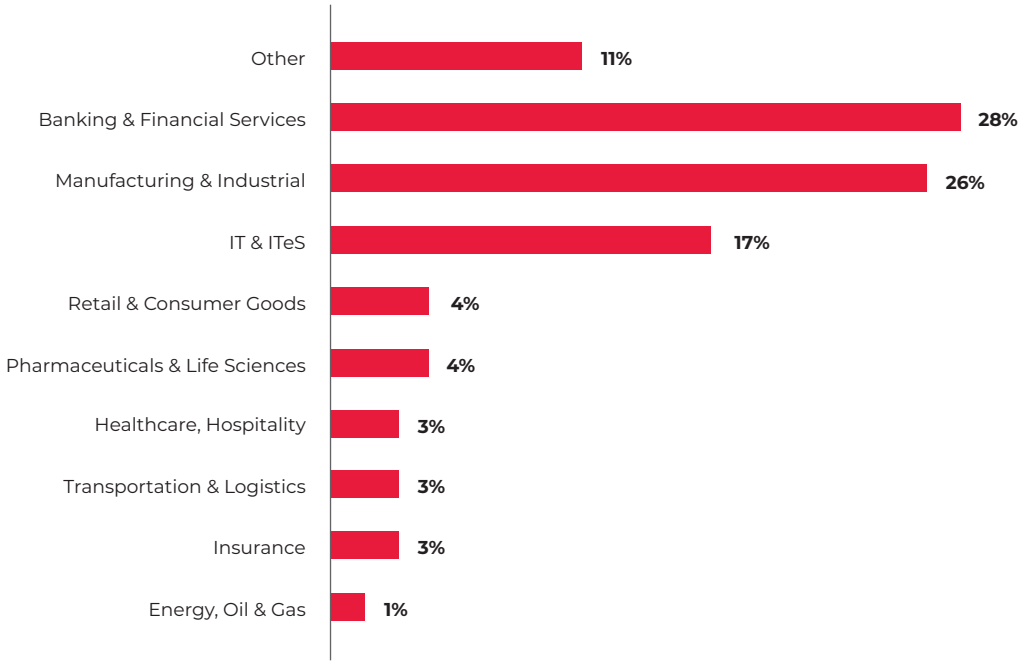

The distribution of respondents across industry sectors aligns with India's economic composition and the sectors facing the most acute cybersecurity pressures from regulatory mandates and operational technology integration challenges.

Large and High-Complexity Enterprises Are Driving Cybersecurity Trends

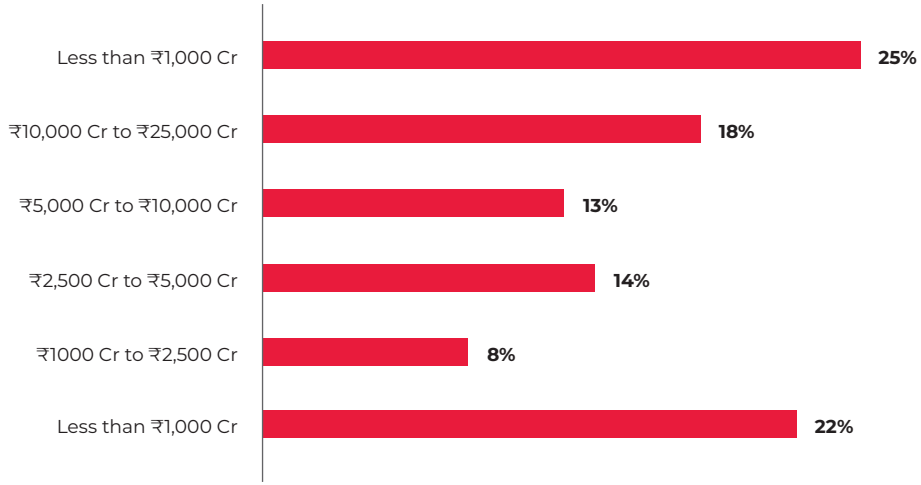

The survey insights reflect the realities of securing large, complex enterprises because scale increases attack surface and recovery complexity.

" The organizations that thrive will be those where boards co-own cyber and AI risk."

Key Contributors

Giridhar has more than 35 years of experience in areas spanning media, consulting and digital technology, working with leading B2B and B2C media organizations across the Asia-Pacific region. He has been actively involved with professional communities in developing content-driven engagements and platforms, and people recognition programs.

R. Giridhar
Group Editor
9.9 Group

With over 18 years of experience in research, consulting, media, and communication, Jatinder Singh currently serves as the Executive Editor at CIO&Leader. He is responsible for shaping the editorial strategy and direction of the publication. He specializes in writing about cutting-edge topics such as analytics, artificial intelligence, cloud computing, the Metaverse, and cybersecurity.

Jatinder Singh
Executive Editor - CIO&Leader
9.9 Group

CISOFORUM

The CISO Forum is a 17-year old community-led platform and content brand for CISOs, information security and risk leaders, run by the 9.9 Group. It offers insights, peer engagement, and best practices on enterprise security, risk, and governance to cybersecurity decision-makers in large and mid-size enterprises. Through its channels, the CISO Forum facilitates discussions on cybersecurity strategy, technology choices and leadership competencies, while positioning itself as an ecosystem hub for India's CISO and risk management community.