FROM HUMAN-LED DEFENSE TO AGENTIC AI

REIMAGINING ENTERPRISE CYBERSECURITY WITH TOPAZ FABRIC

Infosys®
Navigate your next

Abstract

As enterprise digital environments grow more distributed and dynamic, traditional human-led cybersecurity models are struggling to keep pace with the volume, velocity, and complexity of modern threats. This paper explores the transition from fragmented, tool-centric defense to an AI-augmented, platform driven approach enabled by **Infosys Topaz Fabric** and the **Cyber Next platform**. By unifying alerts across identities, endpoints, networks, applications, and data, Topaz Fabric enables security operations to shift from reactive alert handling to anticipatory, context-aware defense, thereby preserving human judgment while extending it through intelligent automation and agent-based collaboration.

This paper explores:

- The shift from fragmented, tool-centric cybersecurity to AI-augmented, platform-driven defense
- How Topaz Fabric unifies alerts across identities, endpoints, networks, applications, and data
- The move from reactive alert handling to anticipatory, context-aware defense
- How agentic AI extends human judgment without replacing defenders
- A responsible blueprint for adopting AI grounded in resilience, governance, and transparency

Transformation of Cyber Services through Topaz Fabric

For many enterprise leaders, cybersecurity now sits at the center of digital resilience. As organizations expand their digital ecosystems, the systems designed to defend them are under pressure to evolve just as quickly.

Cybersecurity rarely changes in one dramatic moment. More often, the shift happens gradually until one day organizations realize the old ways of working are no longer enough. Over the past several years, enterprises have expanded their digital environments at a pace few security teams could have predicted.

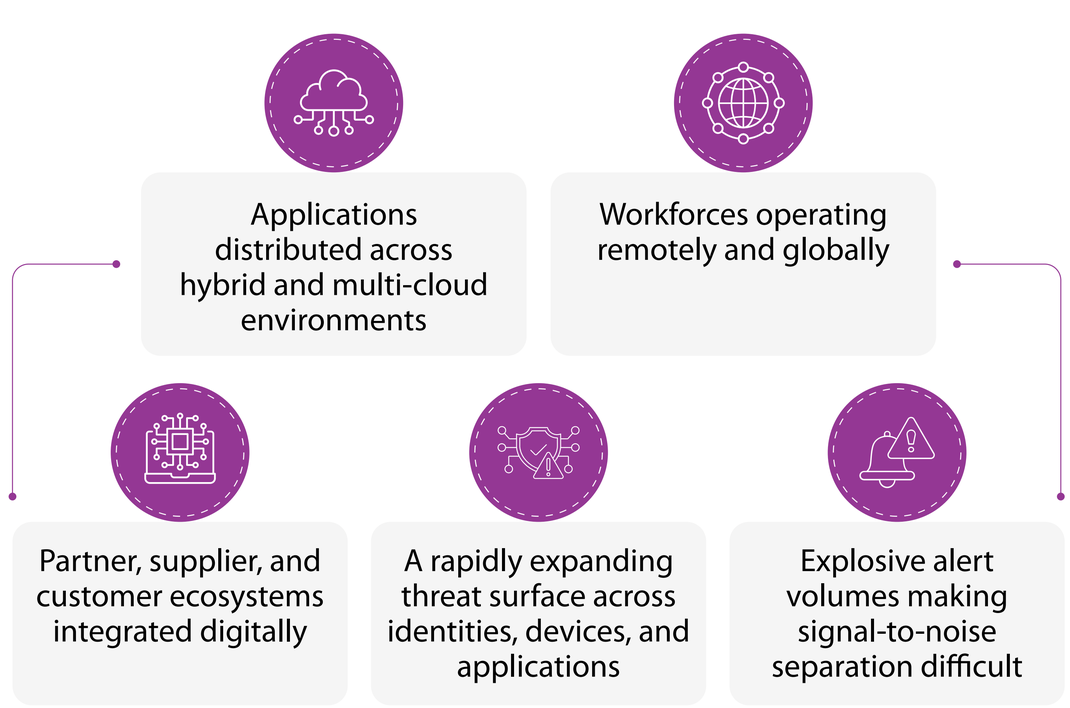

Inside many security operations centers, the reality is easy to see. Alerts stream in from different systems, each designed to monitor its own corner of the environment. Analysts move from one console to another trying to understand what those alerts mean when viewed together. Sometimes the answer is straightforward. Sometimes it takes time to piece together the story behind the activity.

At the same time, attackers have not stood still. Automation allows them to test large numbers of systems quickly, and increasingly they are experimenting with artificial intelligence to improve how they discover weaknesses. The imbalance this creates is subtle but important. Defenders are responsible for protecting environments that grow more complex every year, while the tools they rely on often operate in isolation from one another.

For many organizations this has prompted a broader question. Rather than continuing to add more security technologies, how should cyber defense function in a world where systems, identities, and data are constantly in motion?

At Infosys, one response to that question is taking shape through **Topaz Fabric**. The thinking behind it is not complicated. Security teams already rely on many sources of information: identity systems, infrastructure logs, application behavior, threat intelligence, and data alerts. What has often been missing is a unifying mechanism that allows these alerts to be correlated and interpreted collectively rather than in isolation.

Topaz Fabric acts as the connective layer that allows cyber defense to operate as a coordinated system rather than a collection of separate tools. It brings together data alerts, analytical models, governance policies, and response mechanisms into a single operating environment. In doing so, it allows security teams to move from reactive monitoring toward anticipatory defense, where threats can be understood earlier and responses coordinated across identities, infrastructure, applications, and data.

In practice, this means alerts from across the enterprise no longer remain trapped inside individual security tools. Identity activity, endpoint behavior, network telemetry, and application alerts can be analyzed together inside the same operating environment. Analytical models help interpret those alerts, while automated workflows assist defenders in coordinating investigation and response. Instead of moving between multiple consoles to assemble context, security teams begin with a unified view of risk across the environment.

Once those alerts start coming together, the way security teams work begins to shift. Analysts no longer spend most of their time moving between tools trying to assemble context. Instead, they can focus on interpreting what the activity actually means. Alerts that once looked isolated often turn out to be connected when viewed in the wider environment, revealing patterns that were easy to miss before.

Platforms such as Cyber Next Topaz Fabric help support this kind of environment. Cyber Next Topaz Fabric unifies Posture, Protect, and Prevent, thereby establishing continuous security posture and context, deploying AI agents for active protection, and enabling prevention through content engineering and digital trust. This work is supported by the scale of **Infosys CyberSecurity operations**.

The organization brings together:

7,000+
cybersecurity professionals

500K+
workloads protected

140+
clients protected today

300+
global clients

10M+
identities secured

100+
AI agents removing repetitive analyst work

2.2M
events monitored per second

None of this removes the role of human defenders. Their experience still shapes the decisions that matter most. What begins to change instead is the scale at which they can operate. With intelligence flowing more easily across platforms and environments, teams gain a clearer understanding of risk and can respond more quickly when something unusual appears.

Seen in this context, the evolution of cybersecurity services is less about adding another tool and more about rethinking how defense operates as an integrated system. Enterprises need systems that can learn continuously, interpret alerts in context, and support responses that move as quickly as the environments they protect.

The shift is already underway. The challenge now is learning how to build cyber defense around architectures that allow intelligence, platforms, and human expertise to work together rather than separately.

Emulating Human Defenders in the Path to AI-driven Cyber Defense

Despite widespread discussion of artificial intelligence in cybersecurity, the foundational principle remains unchanged. The most effective defense systems continue to rely on human expertise. Analysts notice small details that automated systems might overlook. Threat hunters follow instincts that come from watching attackers over many years. Incident responders understand how a technical alert might affect real business operations.

Anyone who has spent time inside a security operations center will recognize how this works.

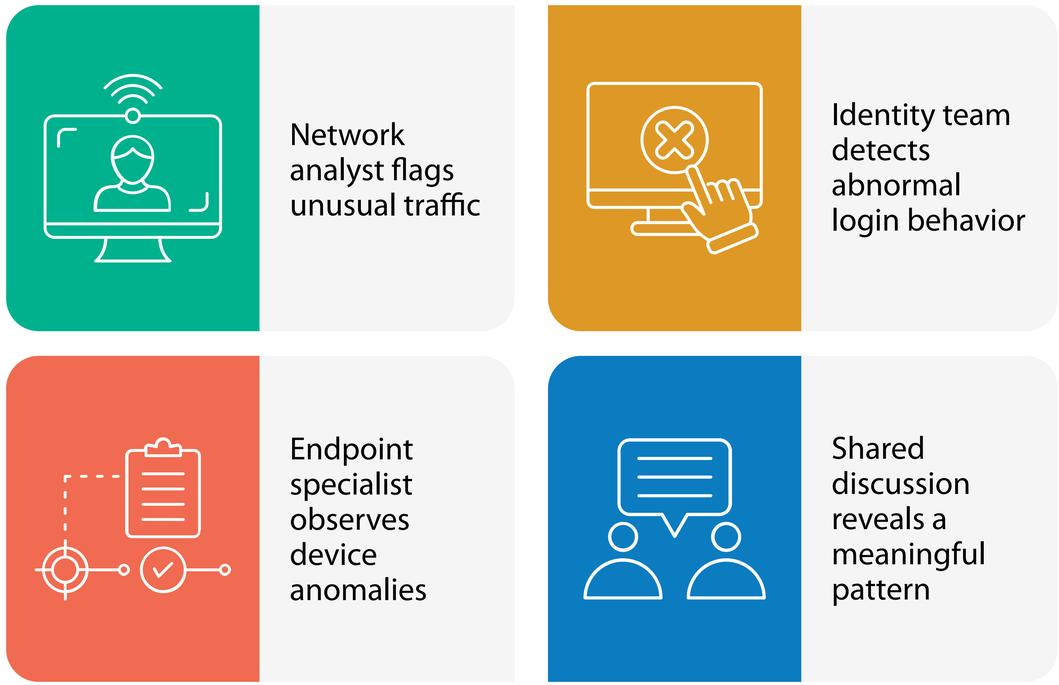

That ability to interpret alerts collectively is one reason human defenders remain so valuable. Security teams do not simply react to alerts. They question what they see, compare notes with colleagues, and use experience to decide what matters and what does not. Over time these habits become a kind of shared understanding about how the organization normally behaves.

As enterprises begin to explore agent-based AI in cyber defense, many leaders are asking how this same way of working can be reflected in intelligent systems. The intention is not to recreate people in software or remove human defenders from the process. The aim is to design systems that support the way effective teams already operate.

Another question that usually comes up is how much freedom automated systems should have when they operate inside a security environment. In practice, defenders already work within boundaries set by company policy, regulatory obligations, and the organization's appetite for risk. Those expectations do not disappear when artificial intelligence is introduced. If anything, they become more important. Systems that analyze activity or recommend responses still need to work within the same limits that guide the people responsible for protecting the enterprise.

These boundaries also help address a concern that often appears when automation is discussed. Security teams want to understand why something happened. If a system flags suspicious behavior or recommends isolating a device, defenders need to see how that conclusion was reached. When the reasoning behind a decision is visible, the system becomes easier to trust.

Another issue that security leaders frequently raise is control. Incidents rarely unfold in neat or predictable ways. New information appears. An alert that looked harmless may turn out to be important. Sometimes an automated response may need to be paused while analysts look more closely. Because of this, many organizations design their cyber defense systems so that people can intervene whenever necessary.

Different organizations describe this idea in different ways. Some speak about human-in-the-loop models. Others describe override mechanisms or escalation thresholds. The meaning is similar in every case. Automation should support defenders, not replace their authority to make decisions.

Responsible AI practices also play a role in this evolution. Cybersecurity systems process sensitive information and interact with critical infrastructure. As artificial intelligence becomes more involved in these environments, organizations must ensure that systems operate within ethical, legal, and privacy frameworks that are already part of enterprise governance.

In many respects this reflects how security teams already work today. Analysts document investigations and explain how they reached their conclusions. Logs are kept so that actions can be reviewed later if necessary. AI systems must follow the same discipline. Actions need to be recorded, explanations available, and outcomes open to review.

Another interesting development appears when several AI agents operate inside the same environment. Human defenders rarely work alone. They exchange observations constantly and build a shared view of what is happening across the enterprise. Agent-based cybersecurity systems are beginning to show similar patterns.

This way of working is not very different from what happens inside a security operations center today. A data protection analyst may notice unusual activity around sensitive information while, somewhere else, a network specialist is looking at traffic that does not quite match normal patterns. At the same time, an endpoint engineer might be investigating behavior on a device that feels slightly off. None of these alerts tells the whole story on its own. But when the observations are shared, a clearer picture begins to emerge.

Agent-based cybersecurity systems try to recreate that same exchange of context. Different agents can focus on different parts of the environment, such as identity behavior, endpoint activity, network telemetry, or application events. When those alerts are brought together, the system begins to resemble the way a small team of analysts might approach the same investigation, each contributing a piece of the overall understanding.

This type of interaction becomes useful in large enterprise environments where threats rarely appear in a single place. A login anomaly might connect to unusual activity on a device. That activity could relate to unexpected data movement across a network. When alerts from these different areas are interpreted together, the overall picture becomes clearer.

Human defenders still sit at the center of cyber defense. Tools can process alerts and surface patterns, but the responsibility for understanding those alerts remains with the people protecting the organization. Security teams bring context that software does not have. They understand which systems support critical operations, how activity normally behaves across the network, and what the business consequences might be if a response is triggered at the wrong moment.

Artificial intelligence changes how information is examined rather than who is responsible for the outcome. Systems can sift through large volumes of telemetry and highlight activity that deserves attention. Analysts then interpret those alerts, compare them with what they know about the environment, and decide how the organization should respond.

Over time this changes the rhythm of work inside a security operations center. Instead of chasing individual alerts across multiple tools, analysts spend more time looking at patterns and asking what those patterns mean. AI agents contribute by analyzing activity and bringing related alerts together, which helps defenders see connections that might otherwise remain hidden.

What emerges is not a machine-led model but a collaboration. Human defenders provide context and accountability. Intelligent systems provide speed and analytical reach. Together they create a form of cyber defense that is better suited to the scale and complexity of modern enterprise environments.

Agentic AI in Cybersecurity: A CISO-led Perspective

When a new technology appears in cybersecurity conversations, most Chief Information Security Officers react the same way. Interest comes first, but it is quickly followed by caution. CISOs spend their days thinking about **risk**, **reliability**, and **accountability**. Anything that changes how cyber defense operates will naturally be examined through that lens before it is allowed anywhere near production systems.

CISO Core Questions

How does the system behave when things go wrong?

Can decisions be reviewed and explained?

Does it introduce new risks?

Artificial intelligence has entered the discussion in exactly that way. Security leaders see the potential. They also see the pressure their teams are under. Security operations centers deal with a constant stream of alerts. Environments are expanding across cloud platforms, identities, and connected systems. Skilled analysts remain in short supply. From that perspective, the idea that intelligent systems could help absorb some of the operational load is appealing.

But curiosity alone is never enough to drive adoption in cybersecurity. CISOs also ask a different set of questions. They want to understand how a technology behaves when something goes wrong. They want to know how decisions are made inside the system and how those decisions can be reviewed later. Above all, they want to be certain that introducing a new capability does not quietly introduce a new vulnerability.

These concerns shape how many security leaders think about agent-based AI. Rather than seeing it as a replacement for human defenders, they tend to view it as an extension of the existing security organization. In that role the technology can assist with analysis, highlight unusual patterns, and help teams interpret alerts that would otherwise take hours to review manually.

This approach allows security teams to focus their attention where experience matters most.

Analysts still investigate incidents. Incident commanders still decide how to respond. The difference is that they begin their work with clearer information about what is happening inside the environment.

For CISOs, the real test of any new capability is whether it improves the security program in measurable ways. One of the areas they watch closely is **detection speed**. If alerts can be interpreted faster, teams can identify suspicious activity earlier. Another area is response time. When incidents are understood quickly, defenders have a better chance of containing them before they spread through the environment.

Improvements in these areas show up in metrics that security teams already track. **MTTD (Mean Time to Detect) and MTTR (Mean Time to Respond)** remain two of the most widely used indicators in security operations. When AI helps reduce the time required to interpret alerts or investigate behavior, those improvements become visible almost immediately in operational performance.

Governance also remains part of the conversation. CISOs remain accountable for the actions taken by their security programs, whether those actions originate from a person or from an automated system. Because of this, organizations often introduce clear operating frameworks when AI becomes part of cyber defense. These frameworks describe how data is used, which decisions can be automated, and where human approval is required.

Transparency plays an equally important role. Security teams must be able to explain what happened during an incident and why particular decisions were made. Automated systems therefore need to provide reasoning that analysts can understand and document. When explanations are available, defenders can review actions after the fact and ensure the system behaved as expected.

Over time, many CISOs come to see agent-based AI less as a new risk and more as a structured tool that helps manage complexity. Modern enterprises generate enormous volumes of security data. No human team can examine every alert that appears across endpoints, networks, and identities. Intelligent systems help narrow that field so that defenders can concentrate on the activity that matters.

That does not remove the human role from cybersecurity. In many ways it strengthens it. Security professionals continue to guide investigations, interpret unusual behavior, and decide how incidents should be handled. AI contributes speed and scale, while people provide context and judgment.

Seen from the CISO's perspective, that balance is what determines whether a technology becomes part of everyday security operations. If a system improves visibility, shortens response times, and operates within the organization's governance framework, it earns its place in the security architecture. If it cannot meet those expectations, it remains an experiment rather than an operational capability.

For most security leaders the conclusion is straightforward. Agentic AI is not about surrendering control of cyber defense to machines. It is about giving the people responsible for security better tools to understand what is happening across increasingly complex digital environments.

Aligning with Infosys CyberSecurity Strategy

Whenever organizations rethink their cyber defense model, the conversation eventually moves beyond tools. Technologies matter, of course, but the bigger question usually becomes how everything fits together. Many enterprises have learned the hard way that adding security technologies one by one does not automatically create a stronger defense. In fact, it can have the opposite effect over time.

Security environments often grow in layers. A company introduces a tool to manage identities. Later it adds endpoint protection, network monitoring, vulnerability scanners, and threat intelligence feeds. Each solves a specific problem. But as the environment grows, the connections between them become harder to see. Analysts begin their day surrounded by dashboards that show pieces of the picture rather than the whole.

That experience has shaped the way Infosys approaches cybersecurity. The strategy has never been about building the largest collection of tools. Instead, the focus has been on how defense operates as a system. When alerts across the enterprise are interpreted collectively, security teams spend less time reconciling fragmented alerts and more time understanding actual security conditions.

Infosys structures its cybersecurity strategy around three connected ideas that guide how defense is designed and delivered across enterprise environments.

Infosys CyberSecurity Strategy

The first is **Cyber Resilience at the Core**, recognizing that attacks are inevitable and systems must be prepared to withstand and recover from them.

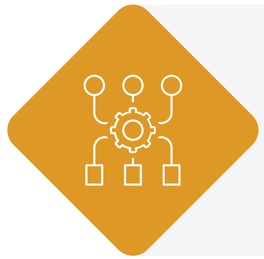

The second is the use of **Platform-Enabled Defense**, where alerts from identities, endpoints, networks, and applications are interpreted together rather than through

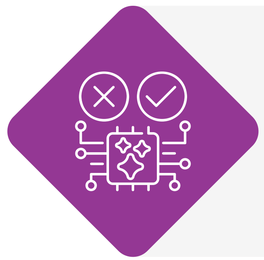

The third is **AI-Powered Services,** where intelligent systems assist defenders by analyzing activity at scale and supporting faster investigation and response.

Resilience is where this thinking usually begins. Many organizations now accept that incidents will occur at some point. Systems expand, new services appear, and attackers continue to experiment with new techniques. Because of this, cyber defense has gradually shifted from a mindset of absolute prevention to one that also considers how the business continues operating when something goes wrong.

A resilient environment behaves differently during an incident. Systems are designed in a way that disruption stays contained. Critical services recover quickly. Teams understand which assets matter most and can act before a situation spreads further than it should. Over time, that ability becomes just as important as blocking attacks in the first place.

The idea of platforms emerged from similar experiences. Security teams repeatedly discovered that investigations slowed down whenever alerts were scattered across different tools. An alert might appear in one system, but the explanation lived somewhere else. Analysts spent time collecting context rather than interpreting it. Platforms help close this gap by bringing alerts together. Identity behavior, endpoint activity, network patterns, and application alerts begin to appear in the same operational view. When that happens, the meaning behind an alert often becomes clearer much earlier in the investigation.

Cyber Next plays this role inside the Infosys CyberSecurity environment. It links enterprise systems, third-party security technologies, and AI-driven capabilities through a platform that allows alerts to be interpreted in context rather than in isolation.

Cyber Next 2.0 – Resilience at Core | Platform Driven | AI-Powered

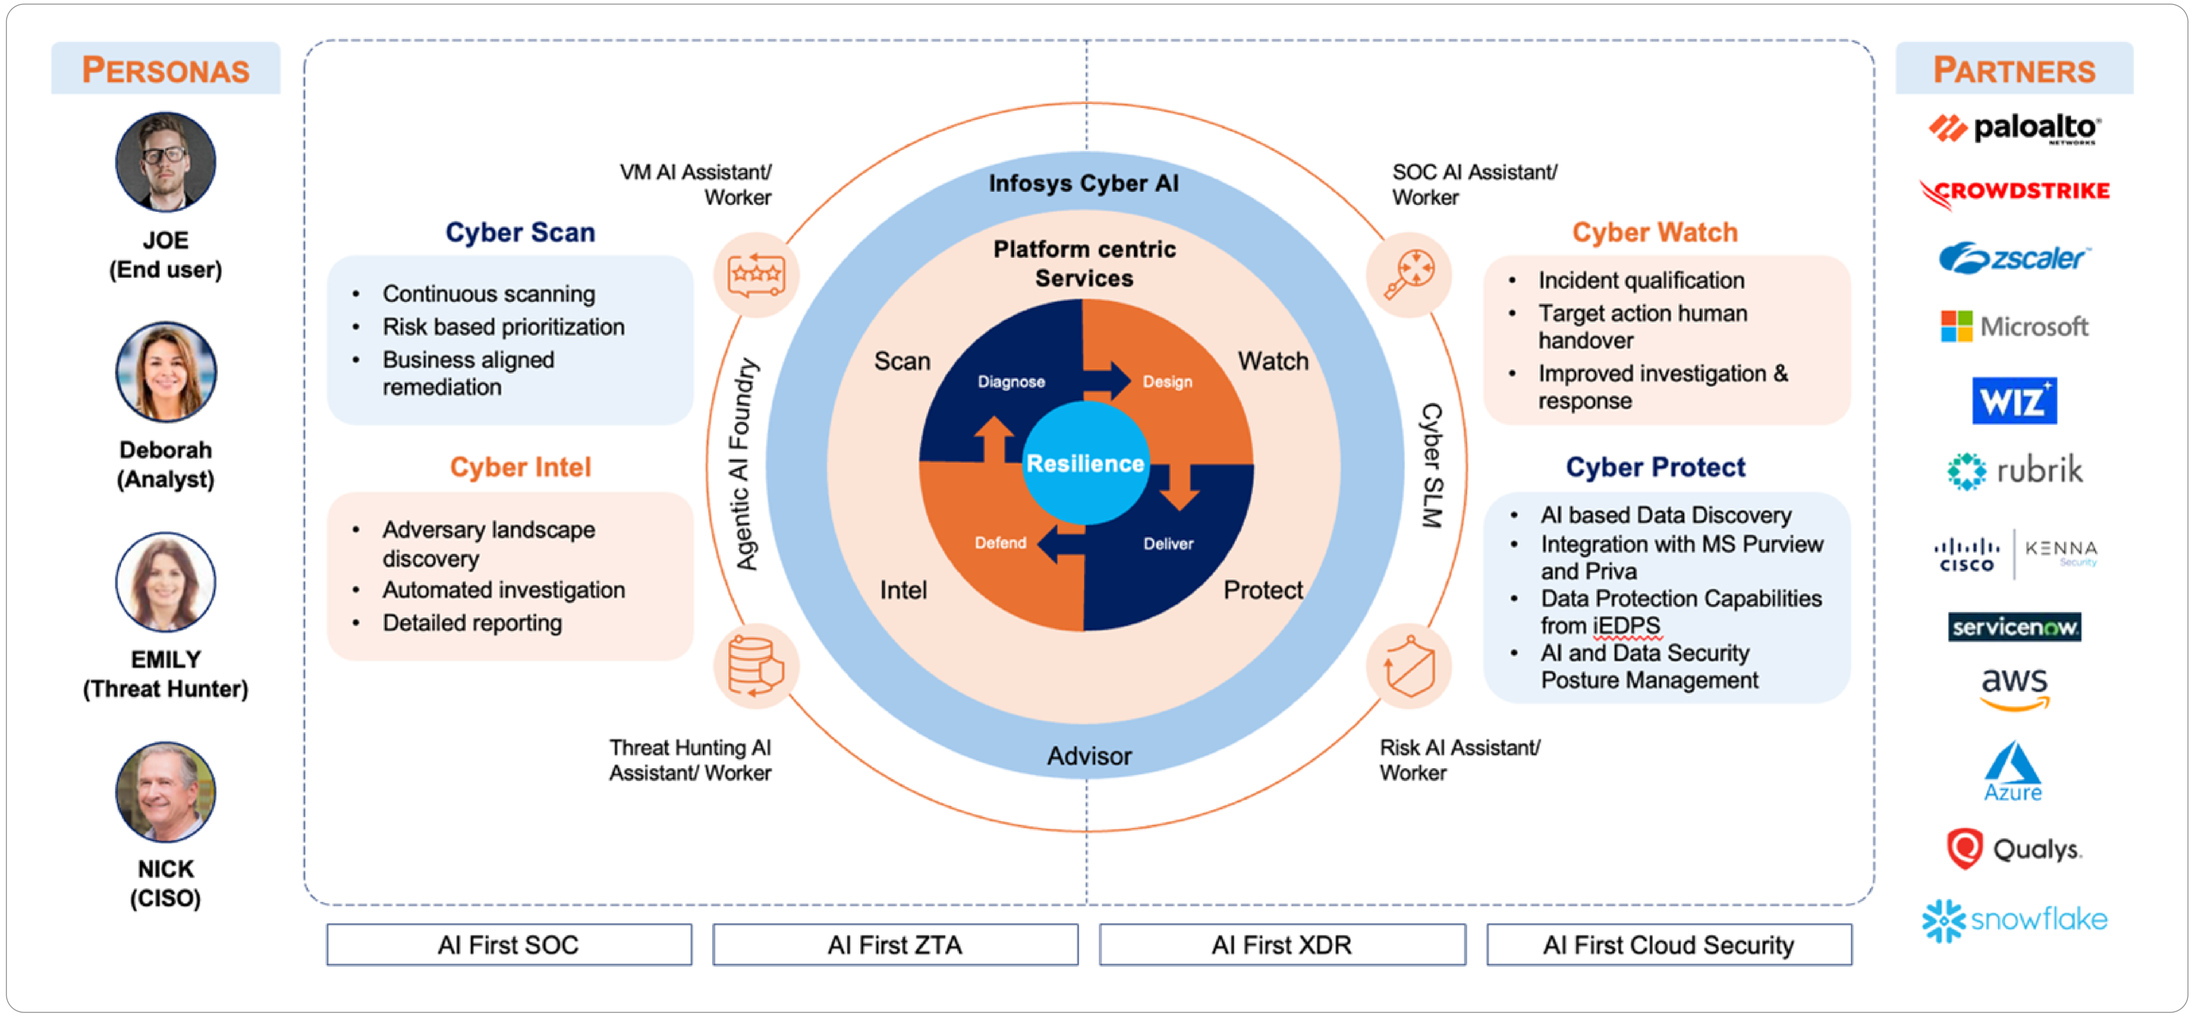

Figure: Cyber Next platform architecture connecting posture, protection, and prevention capabilities through an AI-first security operations environment.

Artificial intelligence enters the picture almost naturally at that point. This progression reflects a broader shift toward AI-first security operations, where automation supports defenders, AI augments decision-making, and intelligent systems increasingly assist with coordinating response. Once alerts are connected through a platform, analyzing them becomes the next challenge. Modern enterprises generate enormous amounts of security data. Patterns often exist inside that data long before defenders notice them. AI helps surface those patterns earlier and highlight behavior that deserves closer attention.

Topaz Fabric helps connect these elements. Instead of operating as a separate intelligence layer, the fabric allows AI capabilities to interact directly with the platforms and systems where security activity takes place. Alerts can be interpreted faster and responses coordinated across the environment without forcing analysts to move between multiple systems.

For security teams, the change is often subtle but meaningful. Less time goes into assembling information. More time goes into understanding risk and deciding how to respond. The technology fades into the background, and the defenders regain a clearer view of what is happening around them.

The impact of this approach can already be seen in production environments. In one large energy enterprise, the Cyber Next platform enabled:

300+
detection use cases

30,000
endpoints secured

20+
automated response playbooks

25%+
operational cost reduction

Seen from this perspective, the strategy is not built around a single technology. It is built around how resilience, platforms, and intelligent analysis reinforce one another. Each plays a role in helping enterprises defend environments that continue to grow more complex every year.

Over time that combination allows cybersecurity to operate less like a collection of independent tools and more like a coordinated system that supports the organization as it evolves.

The Report as a Playbook for Enterprise AI Adoption in Cyber Defense

By the time organizations reach the point where artificial intelligence enters their cybersecurity conversations, the tone of the discussion usually changes. Early excitement about the technology tends to settle into a more practical question. Security leaders begin asking how these capabilities must be introduced inside environments that already carry significant operational responsibility.

Cyber defense programs rarely change overnight. Most enterprises build them gradually, layer by layer, over many years.

New technologies appear, threats evolve, and security teams adapt as the environment grows. Introducing AI follows the same pattern. The organizations that succeed with it usually begin by defining a few simple principles before thinking about scale.

The first principle is **clarity around how decisions are made**. Security teams need to understand the reasoning behind a system's behavior. When an AI capability identifies unusual activity or recommends a response, defenders must be able to see how that conclusion was reached. Transparency matters because security programs depend on trust. Analysts need to know that the systems assisting them are working within the same logic and discipline that guide human investigations.

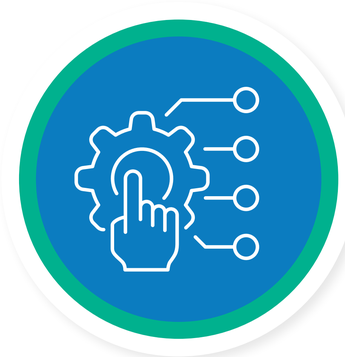

The second principle involves **maintaining control**. Incidents rarely unfold in predictable ways, and security leaders are cautious about introducing systems that might act faster than people can review them. Because of this, organizations normally ensure that analysts remain able to pause automated responses, adjust investigations, or step in when a situation requires judgment. Automation helps with speed, but defenders remain responsible for the decisions that affect the enterprise.

The third principle concerns **governance**. Cybersecurity systems operate across sensitive data, infrastructure, and business applications. AI capabilities must therefore respect the same regulatory and privacy frameworks that guide the rest of the organization's security program. When those boundaries are defined early, enterprises avoid introducing automation that later becomes difficult to manage.

Once these foundations are in place, adoption usually unfolds gradually. Most organizations begin by applying AI in areas where security teams already face the greatest pressure. Security operations centers process enormous numbers of alerts and activity alerts every day. Intelligent systems can analyze these alerts quickly and highlight behavior that deserves attention. Analysts remain responsible for the investigation, but they begin their work with clearer information.

Over time the environment around these systems begins to evolve as well. Alerts from different security tools start to appear together rather than separately. Identity activity, endpoint behavior, and network patterns can be interpreted in context rather than in isolation. This is where platforms become particularly important.

Cyber Next supports this stage by linking enterprise systems, third-party technologies, and AI-driven capabilities through a shared operational environment. When alerts move through a common platform, both analysts and intelligent systems can see relationships that might otherwise remain hidden.

As organizations gain confidence in these capabilities, AI begins to play a broader role inside security operations. Investigations become faster to conduct, alerts are easier to interpret, and defenders are able to coordinate responses across complex environments with greater clarity. Human expertise continues to guide the process, but the scale of analysis increases significantly.

Throughout this transition, the balance between **people**, **platforms**, and **artificial intelligence** remains central. Security professionals provide judgment and an understanding of how technology decisions affect the business. Platforms connect alerts across environments so that defenders can see what is happening. AI contributes the analytical scale required to interpret the enormous volumes of activity moving through modern digital systems.

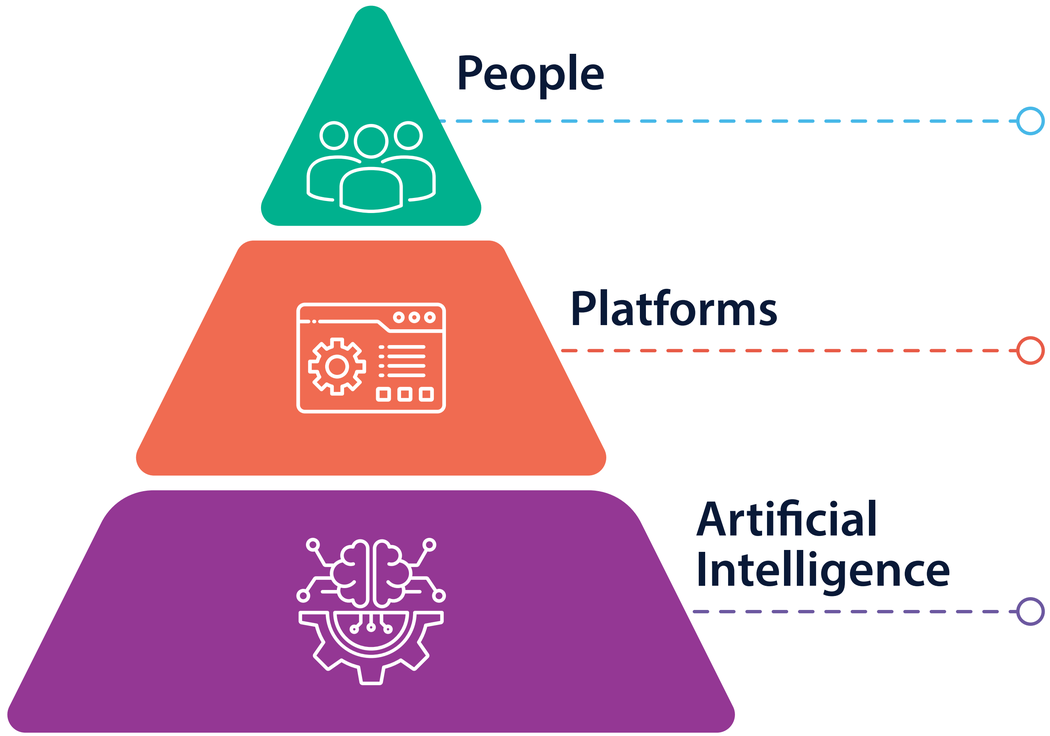

Topaz Fabric strengthens this relationship by connecting intelligence, governance, and execution across the cybersecurity ecosystem. Instead of existing as a separate capability, AI becomes part of the environment where defenders already work and make decisions.

When these elements align, cybersecurity begins to operate differently. Teams gain earlier insight into unusual activity. Investigations move faster. Responses remain coordinated across the enterprise.

Agentic AI, when introduced thoughtfully and supported by platforms such as Topaz Fabric, begins to reshape how cyber defense works across an enterprise environment.

For more information, contact askus@infosys.com

Navigate your next

Infosys.com | NYSE: INFY

Stay Connected