HELP AG
an e& company

STATE OF THE
MARKET REPORT
2026

AGENTIC AI & SECURITY | SOVEREIGN PLATFORMS | THREAT LANDSCAPE | POST QUANTUM READINESS | MARKET INTELLIGENCE

TABLE OF CONTENTS

01

EXECUTIVE LENS

SHAPING POLICY AND REGIONAL DIRECTION

So What Defines Cybersecurity Leadership Today?

By Dr Aleksandar Valjarevic, Acting CEO

> **AI becomes truly transformational when it stops being an application layer and starts becoming part of the organisation's decision architecture.**

A start to 2026 like this has reinforced one reality clearly: there has never been a more demanding, or more consequential, time to be in cybersecurity.

This is not a phase defined by increased activity alone. It is an operating environment defined by sustained pressure, where cybersecurity is embedded directly into how organisations function, sustaining continuity and maintaining trust in systems that do not pause or stabilise.

At this scale and speed, security can no longer rely on human effort alone. Automation and AI are becoming foundational to modern defence. Cybersecurity is no longer a supporting discipline in digital transformation.

Sustained Pressure and Execution Reality

Cybersecurity is now operating in a sustained state of structural pressure that has become embedded in normal business conditions.

Organisations now face close to 2,000 cyberattack attempts per week on average, a 70% increase over two years[1]. The more important shift is not scale, it is precision of targeting. Activity is increasingly concentrated where modern organisations actually operate: identity systems, cloud environments, and interconnected digital services.

Enterprise environments have evolved faster than most security operating models have adapted. Infrastructure is now hybrid by default, workloads span multiple clouds, and third-party integration is embedded directly into core business processes. Availability is no longer a technical metric — it is a direct dependency of business continuity.

Most organisations have already scaled capability. The gap is no longer investment, it is consistency of execution.

Security is now judged by whether it continues to perform as environments change, integrate, and operate under this pressure.

The volume and velocity of modern threat activity are now exceeding the limits of manual security operations. Security teams are expected to interpret thousands of signals daily across identity, cloud, and infrastructure layers. This is accelerating the shift toward automation, orchestration, and AI-assisted analysis as core operating requirements rather than advanced capabilities.

1. Check Point Software Technologies Cyber Security Report 2026.

From Coverage to Continuity

Security effectiveness is not defined by the presence of controls, but by sustained operational performance across cloud, identity, and infrastructure layers. In tightly connected environments, inconsistency in one domain does not remain isolated, it propagates across the system.

Security must therefore operate as a continuous capability embedded into how environments are designed and operated, not as a layer applied after deployment.

Maintaining continuity at this scale requires security operations that can function continuously, not episodically. AI-driven analytics, automated investigation, and orchestrated response are becoming essential to achieving the consistency and speed required across distributed environments.

A Regional Model Shaped by Design

In the Middle East, cybersecurity has followed a different trajectory. It has been embedded into national transformation from the outset, supported by regulatory clarity and long-term digital ambition.

This creates a stronger baseline of maturity but also higher expectations for execution. The question is no longer whether capability exists, but whether it performs reliably at scale.

Sovereignty is now a structural constraint and design principle. It defines how systems are built, how data is governed, and how control is exercised across distributed environments. While this strengthens accountability and clarity, it also increases architectural complexity and raises the requirement for alignment across infrastructure, identity, and operations.

Technology Acceleration and Human Layer

The threat landscape has shifted from opportunistic activity to deliberate, structured, and large-scale execution.

Attackers are using AI-assisted techniques to scale reconnaissance, refine targeting, and accelerate exploitation. Identity remains the primary access vector. Supply chain exposure and configuration weaknesses continue to represent persistent and repeatable pathways into enterprise environments.

Post-quantum cryptography is now moving into active planning cycles. Artificial intelligence is simultaneously reshaping both attack and defence capability. These are no longer future considerations, it is already shaping current architecture decisions.

But technology does not replace execution, it amplifies it.

Security teams are now operating in environments where they must interpret AI-driven outputs, validate system behaviour, and make decisions under uncertainty. Across the region, continued investment in talent development and nationalisation is strengthening this human layer — embedding expertise within operational environments and improving long-term resilience.

> **Cyber sovereignty has shifted from a compliance requirement to strategic power shaping how nations build and control their digital future.**

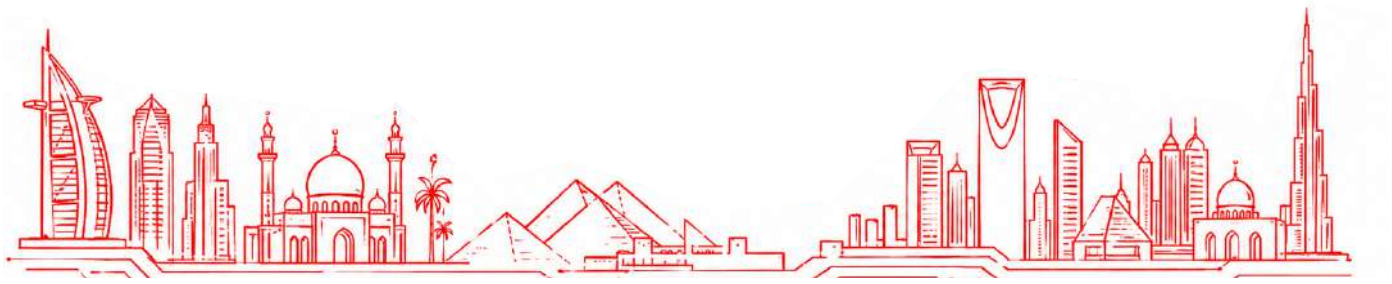

Security as Operational Architecture

Cybersecurity is now embedded in system architecture and enterprise decision-making.

The role of the CISO has already expanded from security oversight to operational resilience and business continuity accountability. Security is now embedded in how the business operates. This is reinforced by financial exposure, with the average cost of a data breach in the region now exceeding $7 million[2], with impact extending beyond IT into operational and regulatory domains.

Across complex enterprise and government environments, the challenge is maintaining consistent performance across distributed systems operating as a single environment. The focus is operational coherence, ensuring security performs as an integrated system rather than a collection of controls.

This shift is accelerating the adoption of AI-driven security operations platforms that unify visibility, automate investigation, and support real-time decision-making across distributed environments.

Looking Ahead

This year has reinforced a defining shift: cybersecurity is no longer about expanding capability — it is about sustaining performance under continuous pressure.

Digital dependency is deepening, AI is reshaping both attack and defence at scale, and post-quantum readiness is moving from theory into implementation. The pace and volume of modern threats now exceed what human effort alone can manage.

Complexity is no longer temporary; it is the operating baseline. In this environment, automation and AI are becoming fundamental to how organisations detect, investigate, and respond in real time.

So what defines cybersecurity leadership today?

It is the ability to maintain consistent security performance under real-world pressure by combining human expertise with intelligent automation to operate at the speed, scale, and persistence of modern threats.

2. IBM Cost of a Data Breach Report 2025.

Abdulla Ebrahim Al Ahmed
Chief Government & VVIP Relations Officer Company - e& UAE

The UAE's digital ambitions are built on trust, resilience and national capability. As AI becomes embedded across government services and sovereign digital ecosystems continue to expand, cybersecurity must operate at the same speed and scale.

Organisations today need security that is continuously adaptive, locally aligned and designed to protect critical infrastructure and citizen data in an AI-driven environment.

Strengthening sovereign cybersecurity capabilities is key to enabling innovation while maintaining the trust that forms the foundation of the UAE's digital future.

Saudi Arabia's Next Cybersecurity Chapter

Execution at Scale

By Safwan Akram, Country Manager - KSA

A Structured National Foundation

Saudi Arabia has established one of the most structured cybersecurity environments in the region. The differentiator is no longer alignment to frameworks, but the consistency with which organisations operate under real-world conditions.

This maturity is being reinforced by sustained national investment. Saudi Arabia's cybersecurity market reached approximately SAR 15.2 billion in 2024, reflecting 14% year-on-year growth, while the sector's contribution to GDP rose to SAR 18.5 billion, underscoring its expanding role in national economic value creation.[1]

As national transformation accelerates, cybersecurity is increasingly embedded into economic continuity, digital trust, and system-wide resilience. It is shaping how organisations scale, interconnect, and maintain stability within highly complex digital environments.

At a leadership level, cybersecurity is structured, enforced, and integrated into national development priorities. Regulatory bodies have established a consistent baseline, elevating cybersecurity into a board-level discipline directly linked to organisational performance.

Frameworks such as the Essential Cybersecurity Controls (ECC) and Critical Systems Cybersecurity Controls (CSCC) define national expectations, while sector oversight continues to expand.

What matters now is how cybersecurity performs in practice, not how it is defined on paper.

Operational Resilience as a Measurable Capability

Regulatory alignment establishes intent. Operational resilience demonstrates effectiveness.

Across the Kingdom, organisations are strengthening Security Operations Centres (SOCs), threat intelligence functions, and incident response capabilities to ensure consistent detection, response, and recovery.

Resilience is increasingly defined by performance under pressure. This requires integrated technologies, structured processes, and teams capable of acting decisively in time-critical scenarios.

Organisations are now assessed not only on preparedness, but on the reliability and consistency of execution.

Architecture-Led Security in a Converged Environment

Saudi Arabia's giga-projects, smart cities, and industrial platforms are creating highly interconnected environments spanning IT, OT, IoT, and cloud infrastructure. These ecosystems operate without traditional boundaries, increasing both exposure and interdependence.

At the same time, the Kingdom remains among the most frequently targeted countries in the region, driven by the scale of its digital expansion and its strategic importance.

In response, security is shifting upstream into architecture. Organisations are embedding security into design decisions from the outset. Large-scale initiatives are increasingly developed without legacy constraints, enabling integrated, zero trust-aligned models from inception.

In this context, architecture is becoming the primary control point for consistency, visibility, and governance across complex environments.

1. Saudi National Cybersecurity Authority Key Economic Indicators in Cybersecurity Sector in the Kingdom Report 2025

Cybersecurity in the Kingdom is shifting from regulatory alignment to measurable operational performance.

Local Capability and Sovereign Control

As cybersecurity investment matures, the focus is moving toward capability ownership and operational control within the Kingdom.

The cybersecurity workforce now exceeds 21,000 professionals[1], with growing demand in specialised domains such as threat hunting, digital forensics, and OT security. Meeting this demand requires structured development pathways, practical experience, and sustained investment in national skills.

Cybersecurity is also increasingly linked to sovereignty. Data residency, in-country operations, and control over critical security functions are becoming central considerations.

This local content expectations are reshaping delivery models. In-country SOCs, Saudi-led operational teams, and partnerships with national institutions are emerging as key differentiators.

Resilience is being built not only through technology, but through ownership.

Cyber Risk as a Leadership Discipline

Accountability across the cybersecurity landscape continues to evolve.

Regulatory expectations now extend beyond organisational compliance to defined ownership at leadership level. Cyber risk is treated as a core business variable, integrated into governance, risk, and audit structures.

Leadership teams are expected to demonstrate visibility, control, and readiness including clear decision-making structures, defined response ownership, and confidence in operational effectiveness.

For CISOs, the role continues to expand beyond protection toward enabling continuity, supporting transformation, and maintaining stability across increasingly complex environments.

Cybersecurity is no longer a supporting function. It is part of how organisations are governed and operated.

Sector-Specific Risk in a Diversifying Economy

As Saudi Arabia's non-oil economy expands, cybersecurity risk is becoming more closely aligned to sector dynamics.

Exposure is no longer uniform. It is shaped by how each sector operates, scales, and interacts with digital ecosystems.

Logistics environments prioritise supply chain integrity and operational coordination. Tourism focuses on safety, privacy, and customer trust. Industrial sectors emphasise uptime, safety, and operational continuity.

This shift is driving more tailored security strategies, where controls and capabilities are aligned to operational realities rather than applied uniformly.

The sectors experiencing the fastest growth are also where cyber exposure is scaling most rapidly, creating new complexity in how risk is managed.

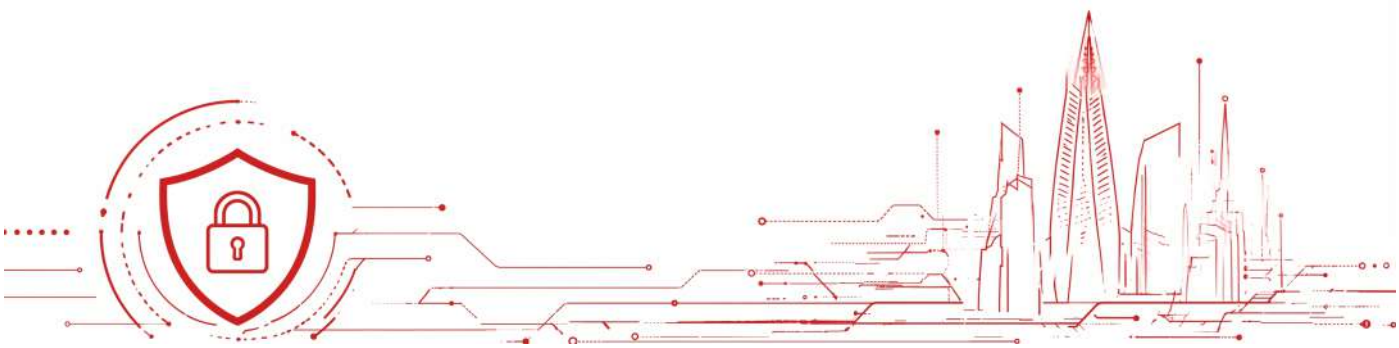

> **The next cybersecurity challenge sits within Saudi Arabia's non-oil economy, where growth, exposure, and risk are scaling simultaneously.**

Execution as the Defining Differentiator

Saudi Arabia has established a strong cybersecurity foundation supported by regulatory clarity and sustained investment.

The defining variable is execution at scale.

For CISOs and technology leaders, the priority is translating strategy into consistently performing operations supported by strong architecture, skilled teams, and clear governance.

Cybersecurity is increasingly positioned as an enabler of growth, continuity, and trust.

Final Perspective

Saudi Arabia's cybersecurity journey is entering a new phase. The focus is shifting from establishing frameworks to demonstrating sustained operational performance across increasingly complex digital environments.

The organisations that will lead the next stage of progress are those able to operate consistently, adapt continuously, and maintain control as digital interdependence expands.

Cybersecurity in the Kingdom is no longer defined by what has been implemented, but by how effectively it performs in practice.


2025
CYBER ECONOMY
DRIVING CYBER INVESTMENT AND STRATEGIC GROWTH

The GCC Cyber Market at a Glance

$7.29 million

Average data breach cost in the Middle East, 61% higher than the global average.

(IBM Cost of a Data Breach Report 2025)

MENA end-user security spending to reach

$4.07 billion

in 2026. 10.1% YoY growth – with managed security services growing at 16.6%, the fastest segment.

(Gartner)

GCC cybersecurity investment to surpass

$9.6 billion

by 2032 sustaining a 7.2% CAGR.

(P&S Intelligence)

92% of UAE firms and **91%** of Saudi firms intend to deploy AI agents and autonomous systems capable of executing complex security tasks without human intervention.

(Cisco AI Readiness Index 2025)

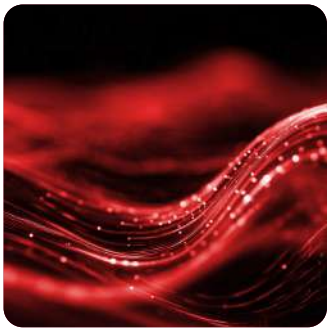

88% of organisations now operate in hybrid or multi-cloud environments.

66% of cybersecurity experts surveyed say they lack strong confidence in their ability to detect and respond to cloud threats in real time.

(Fortinet, 2026 State of Cloud Security Report)

Middle East & Africa (MEA) OT security market is projected to grow from USD 4.36 billion in 2025 to

$9.65 billion by 2030,

at a CAGR of 17.2%.

(MarketsandMarkets)

Sovereign cloud IaaS spending in the MENA region to hit

$250 million in 2026,

one of the highest growth rates globally (89%).

(Gartner)

Market Economics & Demand Signals

Reframing Cybersecurity Investment Across the Region

Cybersecurity investment across the Middle East continues to accelerate.

But growth is no longer the defining story.

The market is entering a more consequential phase where the real pressure is no longer just building capability – it is making that capability work at scale, consistently, across fast-moving and increasingly complex environments.

Cybersecurity spending across the region is projected to surpass $9.6 billion by 2032[1], while information security spending across MENA is expected to reach $4.07 billion in 2026[2], reflecting sustained double-digit growth.

At the same time, cyber risk economics continue to intensify. The average cost of a data breach in the Middle East has reached $7.2 million[3], significantly above the global average, while cybercrime is forecast to reach $23 trillion[4] annually by 2027.

The implication is increasingly clear: cybersecurity is no longer evaluated by the scale of investment but by whether that investment delivers measurable operational resilience.

This distinction is now reshaping the market.

A Market Defined by Expansion - But Not Yet by Integration

Investment across the region is rising across cloud, sovereign infrastructure, managed services, and enterprise modernisation at the same time.

Managed security services alone are growing at 16.6%[2] year-on-year, making them the fastest-growing segment of the cybersecurity market. This reflects a clear shift toward continuous security operations and managed detection and response.

However, increased investment is not yet translating into proportional reductions in operational risk.

Most organisations now operate with significantly expanded security capability but still face challenges in maintaining visibility, integration, and consistent enforcement across hybrid environments.

The issue is no longer whether controls exist. It is whether they work together cohesively.

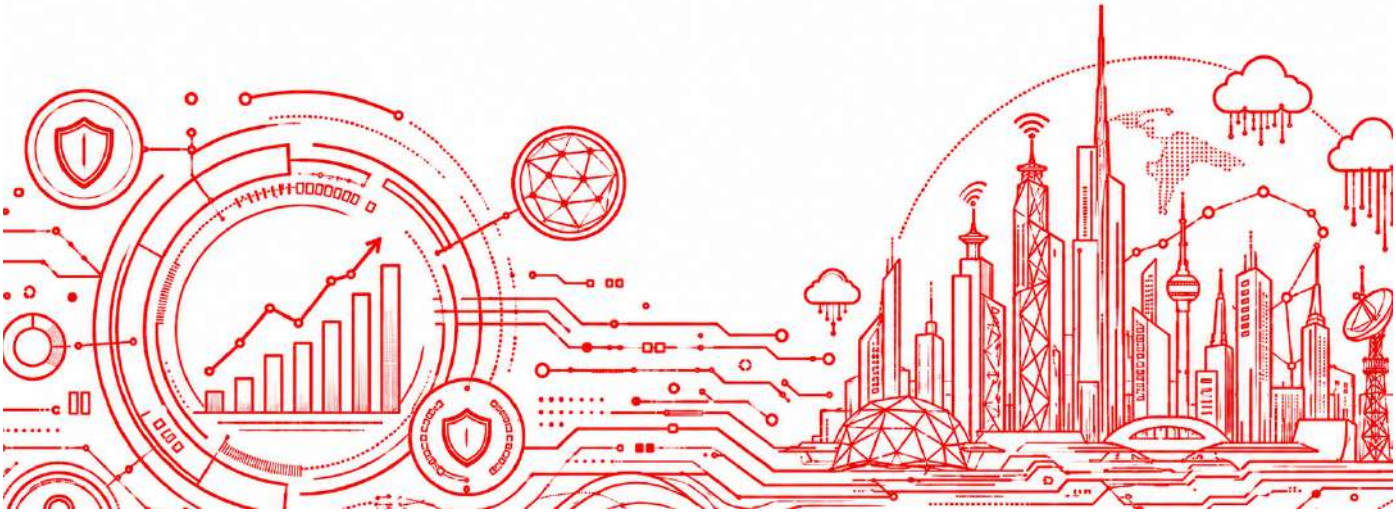

1. P&S Intelligence GCC Cybersecurity Market Size & Share Analysis
2. Gartner Forecasts MENA End-User Spending on Information Security (Press Release) 2025
3. IBM Cost of a Data Breach Report 2025
4. Atlantic Council Counting the Costs Report 2025

The Execution Gap is Becoming the Defining Challenge

The widening gap between cybersecurity investment and breach impact reflects a maturity and execution gap, not underinvestment.

The evidence is visible in breach data: even as detection capabilities improve, many organisations still measure response in hours or days rather than minutes – sustaining a gap between awareness and containment.

This is being driven by four structural forces:

- Attack complexity is increasing faster than defensive integration.
- Security tool adoption is outpacing operational maturity.
- Digital exposure is expanding across cloud, OT, and sovereign infrastructure.
- Threat actors are scaling through AI, automation, and identity exploitation.

This gap is a key reason breach costs remain elevated despite continued investment growth.

Security capability is expanding. Operational coherence is not keeping pace.

Global Context: The Same Problem, Different Pressure Points

This is not a regional pattern — it is a global operating reality.

Cybersecurity spending worldwide is expected to exceed $300 billion annually by 2026–2027[6], yet operational resilience is not improving at the same speed.

The signal across global research is consistent: attackers are compressing the attack lifecycle from drastically, encrypted traffic is limiting visibility, and identity compromise remains the primary entry point for breaches. The Middle East reflects the same pattern – but under higher speed of transformation and higher exposure intensity.

Investment has Outpaced Integration

Over the past decade, organisations across the region have rapidly expanded their security estates.

Cloud adoption accelerated.

Digital services multiplied.

OT and IT environments converged.

Third-party ecosystems expanded.

Security investment followed the same trajectory. But successive layers of tooling have created fragmented environments characterised by overlapping telemetry, inconsistent enforcement, and rising operational overhead.

The result is a structural paradox: organisations often have more security capability than ever before, while struggling to convert that capability into consistent, coordinated outcomes.

The market is no longer optimising for security coverage. It is optimising for operational coherence.

Why Security Operating Models are Being Rebuilt

This pressure is driving a redesign of cybersecurity operating models. Historically, strategies focused on:

- control deployment
- compliance alignment
- technology expansion

6. IDC Security Spending Guide 2023

Today, priorities are shifting toward:

- simplification
- interoperability
- platform efficiency
- operational endurance
- measurable resilience

This is accelerating consolidation toward platform-led security architectures.

The question is shifting from: "*How many capabilities do we have?*" to "*How consistently do they operate together?*"

Managed Security is Moving into an Outcome Economy

Managed security services are evolving beyond traditional monitoring and escalation.

Enterprise priorities are now centred on:

- Faster detection and response
- Measurable exposure reduction
- Continuous validation
- Agentic Ai led operation
- Lifecycle accountability

This is precisely why managed security services are growing at 16.6%[2] in the region – organisations are accelerating their security maturity by partnering for integration and continuous operations, not just monitoring.

Security delivery is shifting from activity-based services to outcome-based performance.

Sovereign Cloud is Reshaping Security Architecture

The global sovereign cloud market is projected to grow from $154.7 billion in 2025 to $1133.3 billion by 2034.[5]

Across the Middle East, sovereignty is now influencing core architecture decisions across:

- Cloud strategy
- SOC design
- Data governance
- Incident response
- Infrastructure localisation

Sovereignty is no longer a compliance layer.

It is a structural design principle for digital infrastructure.

AI Is Compressing the Security Timeline

Artificial intelligence is reshaping both attack and defence cycles.

Defenders are embedding AI into detection engineering, behavioural analytics, triage automation, exposure management, and response orchestration.

According to IBM, organisations using AI and automation extensively reduce breach-related losses by an average of $1.9 million[3] while improving detection accuracy and response times. As adversaries compress attack timelines, AI-driven defence turns speed into a strategic advantage.

5. Fortune Business Insights Sovereign Cloud Market Size

Yet the same technology in adversary hands is shrinking the gap between exposure and exploitation. Attackers are using AI to scale reconnaissance, phishing precision, credential attacks, vulnerability discovery, and social engineering.

A clear divide is emerging between organisations scaling through headcount and those scaling through intelligence.

Viewpoint from the Field

By Girish Dani, Vice President of Sales

Cybersecurity investment across the Middle East is becoming more structured, measurable, and closely aligned to operational outcomes. Organisations are prioritising capabilities that improve business continuity, simplify security operations, and sustain resilience across distributed environments, which means increased focus on AI automation and sovereign platforms.

Detection and response remains a core investment priority, particularly across hybrid IT, OT, and cloud environments. At the same time, identity-led security, Secure Service Edge (SSE), Distributed Denial-of-Service (DDoS) protection, and platform consolidation are becoming increasingly interconnected as organisations move toward integrated architectures.

Boards are also shifting their expectations. There is now greater focus on operational readiness, recovery capability, cyber-risk quantification, and long-term programme sustainability. Cybersecurity discussions are increasingly framed in terms of business impact and operational performance, not just technical control.

Resilient Organisations are those addressing complexity early, consolidating visibility, and aligning cybersecurity directly with business continuity objectives.

Sector View: Financial Services

FORTINET

Financial services has become one of the clearest indicators of where cybersecurity across the Middle East is heading next.

The sector now operates at the intersection of sovereign infrastructure, AI acceleration, expanding regulation, and increasingly sophisticated threat activity — all while being expected to maintain uninterrupted trust and operational resilience.

Regulatory maturity across the region continues to accelerate. Bahrain, Jordan, Oman, Qatar, Saudi Arabia, and the UAE ranked at or near the highest possible score in the ITU Global Cybersecurity Index 2024, underscoring how cyber resilience is now treated as a pillar of national economic stability.

At the same time, the threat landscape is intensifying. IBM's Cost of a Data Breach 2025 reports that AI-driven attacks now account for 16% of cyberattacks globally, with financial institutions remaining a primary target. Post-quantum security is also emerging as a strategic concern, with many organisations across EMEA still unprepared for the transition.

From Fortinet's perspective, the most significant shift is architectural. As sovereign cloud adoption accelerates across the GCC, security is increasingly being designed into infrastructure from the outset rather than added later as a compliance layer. Data residency, localised control, and sovereign security operations are becoming embedded requirements.

At the same time, Banking-as-a-Service models, digital ecosystems, and third-party integrations are expanding faster than many traditional security architectures were designed to support. This is driving a broader industry shift toward more integrated and consistent security approaches across cloud, network, identity, and application environments.

The defining challenge is no longer the deployment of additional security controls. It is maintaining operational resilience across highly connected financial ecosystems, where speed, continuity, and trust have become inseparable from cybersecurity.

Outlook

Cybersecurity across the Middle East is entering a more mature operating phase — defined by integration, discipline, and operational accountability rather than expansion alone.

Over the next 12 to 18 months:

- Sovereign delivery models will become standard.
- Artificial intelligence will be embedded across security operations.
- Managed security services will shift toward outcome-based delivery.
- Regulatory enforcement will continue to strengthen.

The defining challenge is no longer expanding cybersecurity capability.

It is the ability to operationalise it consistently across fast-moving, increasingly complex environments.

Leadership will not be defined by the scale of investment but by the ability to reduce fragmentation, sustain operational coherence, and translate capability into measurable resilience at the speed at which risk now evolves.

2025

THREAT LANDSCAPE

UNDERSTANDING ESCALATION, DISRUPTION, AND RISK

Partner Insights: Global Threat Landscape 2025

FORTINET

26%

of global exploitation attempts were targeted at EMEA with the Middle East continuing to face persistent state-sponsored campaigns against government and technology sectors.

(2025 Global Threat Landscape Report)

zscaler

146%

increase in ransomware attempts blocked year-over-year in the Zscaler cloud marking the sharpest spike observed in the past three years.

(ThreatLabz 2025 Ransomware Report)

paloalto NETWORKS

99%

of organisations experienced attacks targeting AI applications or services while managing an average of 17 cloud security tools, creating significant visibility gaps.

(State of Cloud Security Report 2025)

f5

Only 2%

of organisations are considered "highly ready" to scale AI securely despite 96% already deploying AI models across their environments.

(2025 State of AI Application Strategy Report)

NETSCOUT

8 million+

DDoS attacks were recorded globally in H2 2025 with 42% involving multi-vector techniques spanning between two and five attack types.

(DDoS Threat Intelligence Report H2 2025)

rubrik

90%

of organisations experienced a cyberattack in the past year — with 86% paying ransom demands and 74% reporting compromised recovery systems.

(Zero Labs 2025 State of Data Security: A Distributed Crisis)

CLOUDFLARE

47.1 million

DDoS attacks mitigated by Cloudflare in 2025 representing a 121% increase year-over-year, including a record 31.4 Tbps attack.

(Q4 2025 DDoS Threat Report)

proofpoint

99% of organisations experienced account takeover attempts while **88%** faced post-compromise abuse, with attackers establishing persistent mailbox access within minutes.

(Proofpoint ATO Analysis, Nov 2024–Nov 2025)

HELP AG an e& company

Signals of Escalation

When Threats Outpace Innovation

Insights from Help AG – Managed Security Services

The 2025 threat landscape marks a structural shift in how cyber attacks are executed and scaled. Adversaries are no longer constrained by manual effort. They are leveraging automation, identity exploitation and artificial intelligence to operate with greater speed, precision and consistency.

Threat activity is increasingly persistent, identity driven and designed to bypass traditional control boundaries. Initial access is now dominated by phishing, compromised credentials and application - level exploitation, reflecting a move away from opportunistic vulnerability attacks towards the systematic abuse of trust.

Artificial intelligence has emerged as a force multiplier, enabling attackers to accelerate reconnaissance, scale social engineering and adapt techniques in real time. As a result, attack lifecycles are compressing, reducing the time between initial access and impact and widening the gap between offensive and defensive capabilities.

This shift is not incremental. It represents a change in the operating model. Security effectiveness now depends less on visibility alone and more on the ability to detect, decide and respond at speed.

Identity Becomes the Primary Attack Surface

One of the most significant shifts observed in 2025 is the move from exploiting systems to exploiting trust.

Across incidents handled by Help AG, phishing, identity compromise and application exploits accounted for initial access in 90% of confirmed compromises. Overall, 46.5% of true positive incidents were directly associated with these vectors.

This indicates a structural change in attacker behaviour. Identity is no longer one component of the attack surface. It has become the primary entry point.

Global market research pointing to a sharp rise in phishing and Business Email Compromise closely aligns with Help AG MSS observations. Attackers are prioritising scale, using identity based techniques to bypass traditional controls and gain rapid access to critical environments.

The New Reality of Cyber Attacks 2025

90%

of breaches begin with identity, phishing or application exploits

60%

of attacks end in ransomware.

32%

increase in major compromises.

65%

faster progression from initial access to compromise execution.

<40 hours

from breach to impact in advanced cases.

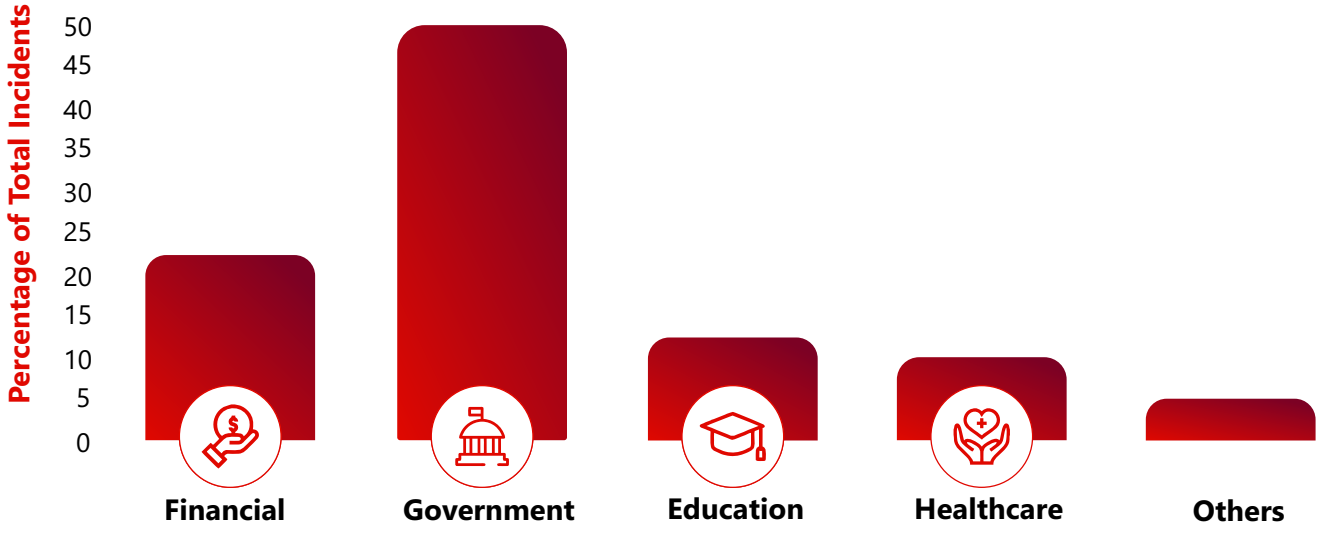

Sector

Distribution of Confirmed Threat Activity Across Sectors

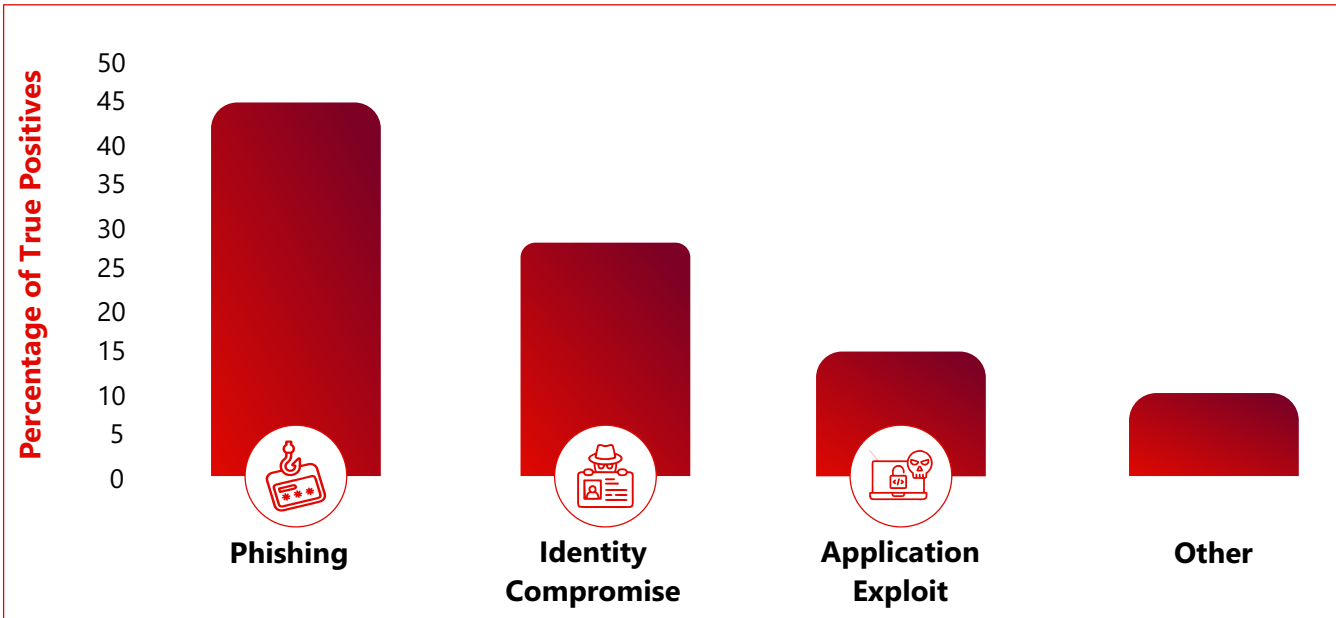

Attack Vector

Distribution of Initial Access Vectors in Confirmed Breaches

Financial Motivation and Targeting Trends

Financial gain remained the dominant driver of threat activity throughout 2025.

The financial sector recorded the highest number of true positive detections, accounting for approximately 22% of all incidents handled. However, financially motivated activity is no longer confined to traditionally lucrative sectors. Government entities accounted for 50% of known compromises, with Education and Healthcare also heavily targeted.

Once access is established, attacker behaviour is primarily focused on monetisation:

- Ransomware deployment: 60%
- Data theft and resale: 24%
- Cryptocurrency mining: 4%
- Additional activities including SMS pumping, toll fraud and payment re-routing

These patterns indicate consistent targeting of high value environments, particularly those containing sensitive financial data, mission critical systems and privileged access paths.

Observed Post-Compromise Activities 2025

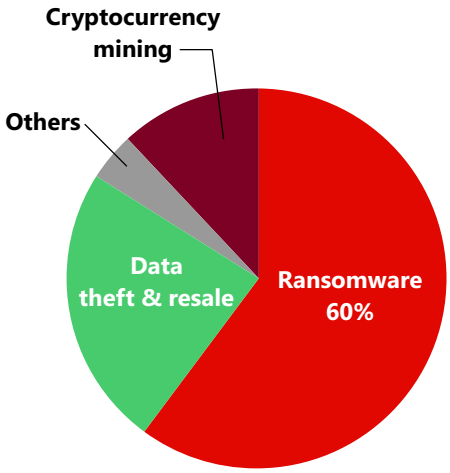

Ransomware Evolves into Control First Campaigns

Ransomware has evolved beyond immediate encryption. Attackers are now adopting multi-stage, control-first approaches, prioritising access, persistence and privilege escalation before any visible impact.

Help AG MSS investigations show consistent patterns across key detection categories:

- Endpoint malware detections: 25% of true positives
- Privilege escalation attempts: 23% of true positives
- Living off the land binaries and suspicious process creation: 8% of true positives

In nearly 90% of ransomware related major compromises, attackers established identity or domain dominance before initiating data exfiltration.

Ransomware is no longer the attack itself. It is the final stage of a successful intrusion.

Disruption-First (Legacy Ransomware)

- Initial Access (Phishing / Exploit)
- Rapid Payload Delivery Ransomware
- Immediate Encryption & Impact

Control-First (Modern Ransomware/APT)

- Initial Access (Identity / App)
- Privilege Escalation & Persistence
- Lateral Movement Data Exfiltration Extortion

Threat Activity at Scale: Acceleration and Compression

Help AG MSS telemetry throughout 2025 recorded a sustained increase in both the volume and speed of confirmed threat activity, alongside a growing divergence in how quickly attacks reach impact.

Key observations include:

- A 32% increase in major compromises, with government entities accounting for more than 50%.
- Advanced Persistent Threat (APT) activity from state sponsored threat actor groups like APT33 (Refined Kitten), APT34 (OilRig) and APT41 (Operator Panda) accounted for 17% of observed incidents, reflecting sustained targeted intrusion activity by state-aligned actors.
- An overall reduction in attacker dwell time of more than 20%.
- In two major DFIR cases, attackers progressed from initial access to impact in under 40 hours.
- Across other incidents, the average dwell time remained approximately 120 days.
- A 65% reduction in the time required for attackers to progress from initial access to compromise execution.

These findings point to a dual reality. While some intrusions remain prolonged, others are executed at speeds that outpace traditional detection and response models.

Reduction in Time-to-Impact Across Confirmed Attacks

This growing variance is being driven by automation and AI enabled tooling, allowing attackers to selectively compress timelines and accelerate impact where conditions permit.

Artificial Intelligence as a Force Multiplier

Artificial intelligence is no longer an emerging capability within the threat landscape. It is actively shaping how attacks are executed, scaled and refined.

Help AG MSS observations from 2025 highlight three defining shifts:

- AI is accelerating the attack lifecycle

  Automated reconnaissance, large-scale generation of targeted phishing campaigns and continuous testing of authentication mechanisms are reducing the time required to identify and exploit weaknesses.

- AI is amplifying dominant attack vectors

  With 90% of compromises originating from identity, phishing and application exploits, AI enables attackers to scale these techniques with greater efficiency and precision.

- AI is compressing time-to-impact

  A 65% increase in the speed of compromise execution has been observed, with some attacks progressing from initial access to impact in under 40 hours.

This represents a fundamental shift in attacker capability:

- Existing techniques are not being replaced, but optimised.
- Manual effort is being reduced across the attack lifecycle.
- Attack paths are being adapted in near real time.

The result is a threat landscape where attacks are more targeted, more efficient and increasingly difficult to disrupt using static or manual defence models.

Q1 2026: Validation of Emerging Threat Patterns

The 2025 outlook identified a clear trajectory: compressed attack lifecycles, operational use of artificial intelligence and closer alignment between cyber activity and geopolitical developments.

Within the first quarter of 2026, these projections materialised.

- Cyber activity scaled instantly in response to geopolitical events
  Following the escalation of tensions in late February 2026, reported attack volumes in the UAE increased from approximately 200,000 per day to between 500,000 and 700,000 daily attempts, as reported by national authorities.
- Artificial intelligence moved into operational use
  AI is no longer experimental in cyber attacks. In 2026 alone, Help AG supported three DFIR engagements involving AI-driven attacks where the full attack chain, from initial access to action on objectives, was completed within 48 hours with near-zero dwell time.
- Identity-based intrusion further intensified
  Identity continues to represent the primary attack surface, with early 2026 activity reinforcing the dominance of credential-based access and identity exploitation techniques.
- High-volume disruption activity increased across critical sectors
  Large-scale disruption campaigns accelerated sharply, with Help AG observing more than 114,000 DDoS attacks in Q1 2026 alone. Over 70% of attacks were multi-vector, highlighting a clear shift toward more coordinated, persistent, and harder-to-mitigate disruption activity.
- Wiper-based attacks emerged as a distinct destructive vector
  After a period of lower activity, destructive wiper malware re-emerged in several incidents, designed to irreversibly destroy data and cripple recovery efforts beyond the initial breach.

These developments show how rapidly threat patterns can escalate driven by external events. What was previously an emerging trend is now an operational reality.

The velocity gap identified in 2025 has not narrowed. It has widened.

Outlook

The patterns observed signal a transition towards deeper systemic risk.

Adversaries will continue to compress attack lifecycles, reduce reliance on human intervention and integrate cyber operations with broader geopolitical and informational objectives.

In 2026, the challenge will shift from recognising these patterns to operationalising them at scale. Organisations that translate these insights into measurable improvements in detection, response and resilience will be better positioned to manage both sustained and event driven threats.

The insights and statistics in this article are drawn from telemetry, DFIR investigations and confirmed incidents observed by Help AG's 24x7 AI-powered sovereign SOC and Managed Security Services operations.

The analysis reflects real-world threat activity across government, financial services, healthcare, education and critical infrastructure environments throughout 2025 and early 2026.

Beyond Volume

The New DDoS Reality

By Gaurav Srivastava, Director Managed Security Controls

In the UAE, digital services now underpin core economic and operational activity across government, finance, telecommunications, and logistics. As these services have scaled, continuous availability has shifted from a competitive advantage to a baseline expectation.

This environment leaves very limited tolerance for disruption, particularly where services operate in real time and at national scale.

Against this backdrop, Distributed Denial-of-Service (DDoS) activity has increased significantly, rising from 38,797 recorded incidents in 2019 to 371,303 in 2025, an increase of approximately **857%**.

The scale of growth is significant, but the more important shift is in how these attacks behave operationally.

From Volume to Operational Pressure

Historically, DDoS risk was defined by peak traffic volumes. That lens is now incomplete.

While high-intensity attacks continue to increase, with 135 incidents exceeding 100 Gbps in 2024, up from just 9 in 2020 duration is becoming the more operationally relevant factor.

In 2025:

- **177,578 attacks lasted between 10 and 60 minutes**
- **51,767 extended beyond one hour**
- The longest recorded attack persisted for **over 85 days**

This changes the nature of disruption.

Short bursts test capacity. Persistent attacks test endurance.

Mitigation is no longer a temporary response activity. It is increasingly a sustained operational condition, requiring organisations to remain stable while defensive controls are continuously engaged.

The Shift in Defensive Assumptions

This evolution exposes a gap between traditional defence models and modern attack behaviour.

Many organisations still rely on threshold-based protection designed to detect sudden spikes in traffic and trigger mitigation responses.

However, a growing proportion of attacks operate below extreme volumetric thresholds while maintaining enough consistency to degrade services over time.

DDoS Activity Profile – 2025

Scale

371,303
Total DDoS Attacks

Intensity

539.1 Gbps
Largest Attack Bandwidth

550.5 Mpps
Highest Packet Rate

Persistence

85 days
Longest Attack Duration

13.94%
Attacks lasting >1 hour

Complexity

12
Max Attack Vectors in a Single Attack

Behaviour Shift

>60%
Multi-vector Attacks

91%
Sustained Low-Volume Attacks

This creates three operational pressures:

- Detection gaps: Gradual traffic patterns may not trigger spike-based alerting models.
- Sustained resource consumption: Continuous mitigation places ongoing load on infrastructure and security teams.
- Operational fatigue: Extended incident duration increases cognitive and operational strain.

The implication is clear: resilience is no longer about absorbing peaks, it is about sustaining operations under continuous pressure.

How Attack Focus Is Changing

The nature of targeting is also becoming more precise.

Rather than attempting to overwhelm entire organisations, attackers are increasingly focusing on dependencies that underpin availability.

Common targets include:

- DNS and NTP services, often used in amplification-based attacks.
- VPN gateways, particularly during peak usage windows.
- Application-layer components, including APIs and session management systems.

Attack patterns observed in early 2026 — including UDP floods, DNS amplification, and IP fragmentation — reinforce this shift.

The objective is no longer immediate disruption at scale, but sustained pressure on critical system dependencies.

This requires a shift in defensive thinking: resilience must extend beyond front-end protection to include the full service architecture.

Implications for Organisations

The shift toward persistent DDoS activity introduces implications that extend beyond downtime.

- Operational continuity becomes harder to maintain, as systems operate under prolonged defensive load.
- Decision-making becomes more complex, particularly in distinguishing between transient and sustained campaigns.
- Planning assumptions are disrupted, including capacity models, cost forecasting, and response readiness.

In this context, availability is no longer a static engineering outcome. It becomes an actively managed condition.

DDoS Activity – Q1 2026

114,566
Total DDoS Attacks
Volume

1464
Average Attacks per Day
Frequency

352.8 Gbps
Largest Attack Bandwidth
Intensity

16.12%
Attacks lasting >1 hour
Persistence

70%
Multi-vector Attacks
Complexity

Building Resilience for Persistent Threats

Addressing this shift requires evolution in approach rather than reinvention of technology. Three priorities stand out:

Continuous visibility - End-to-end monitoring across network, application, and access layers is required to identify both high-volume events and sustained low-intensity attacks over time.

Layered protection models - Combining on-premise controls with cloud-based mitigation ensures flexibility to respond to both short-lived spikes and prolonged activity.

Operational readiness for duration - Response structures must be designed to support extended incidents, ensuring teams, processes, and escalation paths remain effective over time.

> **Sustained low-volume attacks accounted for 91% of all DDoS activity in 2025, demonstrating how attackers increasingly prioritise persistence to degrade services over time rather than high-volume short term disruption.**

Final Perspective

DDoS activity is no longer defined solely by scale. It is defined by persistence.

As digital services become foundational to national and economic systems, availability can no longer be treated as an outcome achieved through capacity alone. It must be maintained continuously under evolving conditions of pressure.

The defining challenge is no longer whether organisations can withstand a spike but whether they can remain stable when disruption becomes sustained.

04

OFFENSIVE SECURITY

VALIDATING RESILIENCE IN PRACTICE

Inside the Mind of a Pen Tester

Seeing Beyond Controls

By Help AG Offensive Cybersecurity Team

A penetration test is often misunderstood as a technical exercise focused on identifying vulnerabilities within systems and applications. In reality, it is increasingly an exercise in understanding how security controls behave under pressure and where they fail to reflect real-world attacker behaviour.

Modern enterprise environments are complex, distributed, and heavily governed by layered security controls. But complexity does not automatically translate into resilience.

From a tester's perspective, the objective is not simply to identify gaps, but to understand how far existing controls can be pushed before they stop reflecting operational reality.

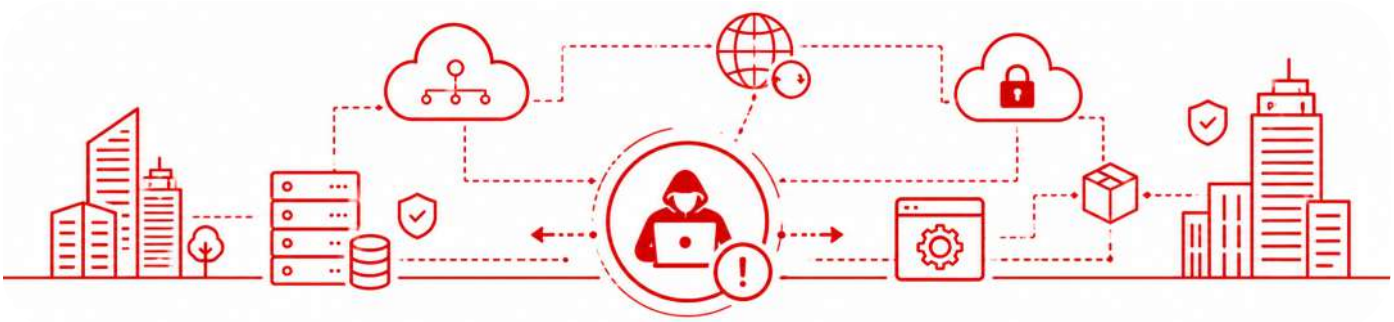

Beyond the Surface of Security Controls

On paper, most environments appear well protected. Controls are in place, policies are defined, and monitoring frameworks are operational.

However, penetration testing often reveals a different layer, one that only becomes visible when systems are observed under adversarial conditions.

This includes how authentication mechanisms behave under abnormal request patterns, how segmentation holds under lateral movement attempts, and how detection systems respond to non-standard interaction flows.

The gap is rarely the absence of controls. It is the assumption that controls behave consistently under all conditions.

Seeing Systems as an Adversary Would

A key shift in modern penetration testing is the move away from isolated vulnerability discovery toward behavioural analysis of systems.

Instead of focusing only on whether a vulnerability exists, the emphasis is increasingly on how systems respond when intentionally stressed, manipulated, or chained across multiple vectors.

This includes testing how security layers interact with each other not just in isolation, but as a combined defence architecture.

From this perspective, weaknesses are often not individual flaws, but breakdowns in coordination between controls.

The Reality of Modern Enterprise Environments

Enterprise environments today are rarely static. Hybrid infrastructure, multi-cloud adoption, and rapid digital transformation have created highly interconnected systems with varying levels of visibility and control.

From a penetration tester's point of view, this introduces a consistent challenge: security is often distributed, but risk is interconnected.

A control that appears effective in one environment may behave differently when tested across integrated systems, third-party services, or legacy components.

> **Many real-world exposures emerge not from a single weakness, but from interaction between systems that were never designed to be tested together.**

Where Controls Meet Operational Reality

One of the most important outcomes of penetration testing is not simply identifying vulnerabilities, but validating whether security controls function as expected under real-world conditions.

This includes assessing whether detection systems generate meaningful signals, whether response mechanisms activate correctly, and whether security assumptions hold when systems are actively engaged rather than passively monitored.

In many cases, the difference between theoretical security and operational security becomes most visible during this phase.

The Value of Adversarial Thinking

The most effective penetration testing approaches are increasingly defined by adversarial thinking rather than tool usage.

This means simulating not just technical exploitation, but behavioural patterns that reflect how modern attackers actually operate — adaptive, iterative, and context-aware.

It also requires understanding how defenders interpret signals, and how quickly environments can respond when conditions deviate from expected baselines.

The goal is not disruption. It is clarity - identifying where assumptions about security no longer match operational reality.

Final Perspective

This is the inside view of a pen tester's mind, where security is not defined by controls on paper, but by how they behave when confronted with reality.

In modern enterprise environments, resilience is proven at the point of pressure, where assumptions are challenged, systems interact, and weaknesses emerge not in isolation, but in combination.

Ultimately, security is not what is implemented - it is what still holds when it is tested.

AI-Powered Offensive Security

Redefining Red Teaming in the Age of LLMs

By Mohamad Kontar, Offensive Cybersecurity Manager – KSA

Offensive cyber operations are changing in both structure and execution speed, as the effort required to plan and launch complex attack scenarios continues to decline.

What once required time, contextual understanding, and iterative refinement can now be generated and adapted with significantly less friction through the use of large language models (LLMs), enabling a wider range of attack variations to be produced more rapidly.

This reduces the time required to move from reconnaissance to execution, compressing traditional offensive cycles.

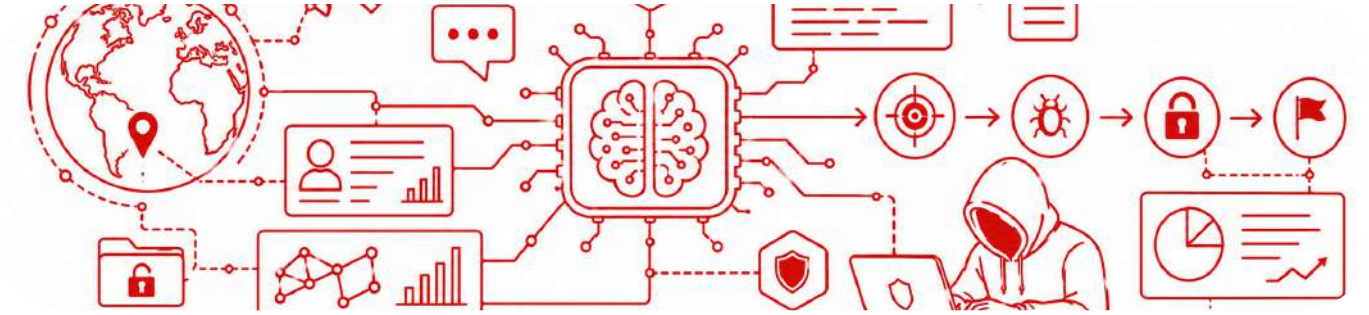

The Offensive Playbook Is Evolving

The core stages of offensive security - reconnaissance, analysis, exploitation, and reporting remain structurally unchanged. What is evolving is how quickly and variably these stages can now be executed.

Agentic Ai are enabling faster synthesis of open-source intelligence, rapid scenario development, and the generation of multiple attack paths from a single starting point. This increases both the speed and variability of offensive activity.

As a result, offensive operations are becoming more adaptive, with techniques that can be modified and restructured dynamically during execution.

At the same time, this variability introduces challenges for traditional detection methods that rely on predictable patterns of behaviour.

Prompt injection, identified by OWASP (Open Worldwide Application Security Project) as one of the most critical risks in generative AI systems, highlights how AI-enabled environments can be influenced when not evaluated under adversarial conditions.

The Enterprise Blind Spot

Across enterprise environments in the region, a consistent gap is emerging.

Organisations are integrating LLMs into customer platforms, internal workflows, and business processes. These systems are typically assessed through standard application security testing approaches, including vulnerability scanning and penetration testing.

However, the behaviour of the AI layer itself is often not evaluated under adversarial conditions.

This creates a structural gap where systems are validated for functionality, but not for behavioural resilience.

Across deployments, AI-enabled systems may operate as expected under normal conditions but behave unpredictably when exposed to manipulated inputs or indirect prompt-based influence.

Regional Acceleration and Exposure

The adoption of artificial intelligence across Saudi Arabia and the UAE continues at scale.

Saudi Arabia's LEAP 2025 event attracted significant AI investment commitments, reinforcing the Kingdom's strategic focus on AI-driven transformation. National initiatives such as HUMAIN, backed by the Public Investment Fund (PIF), are strengthening sovereign AI capability development.

In parallel, regulatory frameworks including the National Cybersecurity Authority (NCA) controls and broader data protection regulations are reinforcing baseline security expectations across critical sectors.

In the UAE, AI adoption continues to expand across government and enterprise environments, supported by national cybersecurity strategies and frameworks such as the Abu Dhabi Global Market (ADGM) Cyber Risk Management Framework.

This level of adoption increases the number of AI-integrated systems across critical infrastructure, financial services, and public sector platforms, each introducing potential exposure if not assessed under adversarial conditions.

While frameworks such as the NCA Essential Cybersecurity Controls (ECC), Saudi Central Bank (SAMA) requirements, and regional data protection laws provide strong governance foundations, AI-specific adversarial validation remains an emerging discipline.

> **AI has not changed the nature of attacker behaviour, but it has significantly compressed the time required to execute it.**

From Traditional Testing to AI Red Teaming

Security validation approaches are evolving to reflect this shift.

AI red teaming is emerging as an extension of traditional security testing, focused not only on system functionality but on behavioural resilience under manipulation. This includes evaluating how AI systems respond to:

- Manipulated or adversarial inputs.
- Attempts to extract sensitive or system-level information.
- Indirect instruction injection into model behaviour.
- Unintended actions triggered through connected workflows.

The objective is to assess how systems behave when actively influenced, not just how they perform under expected conditions. This represents a shift toward continuous adversarial evaluation rather than static validation.

Aligning Defensive Capability with Offensive Acceleration

As offensive actors increasingly leverage Agentic Ai to accelerate reconnaissance and attack development, defensive operations are evolving in parallel.

Security operations environments are adopting AI-assisted capabilities to support faster analysis, correlation of alerts, and identification of anomalous behaviour across systems. At the same time, governance models are evolving to define clearer boundaries around what AI systems can access, process, and execute within enterprise environments.

The focus is shifting from controlling outputs in isolation to understanding and governing system behaviour holistically.

Building AI Security Capability

Across the region, investment in AI capability development continues to grow rapidly.

National programmes in Saudi Arabia have already trained large-scale talent pools in data and AI disciplines, while similar initiatives in the UAE are expanding workforce readiness in emerging technologies.

The next phase of this evolution is the development of specialised AI security capability focused on understanding adversarial behaviour in AI systems and building structured red teaming methodologies.

Early assessments of AI-integrated environments consistently demonstrate that traditional testing approaches alone are insufficient to identify behavioural risks. This reinforces the importance of incorporating adversarial evaluation early in deployment lifecycles.

> **AI deployments are often assessed through traditional security reviews, while the behaviour of the model itself remains insufficiently tested for manipulation or unintended exposure.**

Operational Implications for Leaders

As AI-enabled offensive capabilities mature, enterprise focus is shifting from awareness to structured readiness.

Security programmes are increasingly expected to incorporate adversarial testing of AI systems as part of broader assurance frameworks.

Governance is also moving closer to the system layer, ensuring that AI behaviour is controlled across workflows, integrations, and model interactions rather than only at the application boundary.

From a defensive perspective, matching the speed and adaptability of AI-driven offensive activity is becoming essential. This is driving increased adoption of AI-assisted detection, correlation, and behavioural monitoring within security operations environments.

Finally, developing AI-specific security expertise is becoming a long-term requirement, as traditional security skillsets alone are no longer sufficient to assess adaptive systems.

Final Perspective

Offensive cyber operations are evolving in both structure and execution as AI reduces the effort required to plan and launch complex attack scenarios.

As offensive cycles compress, the boundary between experimentation and exploitation becomes increasingly narrow.

In this environment, security validation must extend beyond traditional application testing to include adversarial assessment of AI systems and their behaviour under manipulation.

Organisations that develop structured AI red teaming capability early will be better positioned to anticipate emerging risk patterns, align with evolving regulatory expectations, and operate effectively in an AI-driven threat landscape.

In an environment shaped by large language models, every AI system becomes part of the attack surface and must therefore be evaluated, understood, and governed accordingly.

05

ARTIFICIAL INTELLIGENCE

SCALING DEFENCE, SPEED, AND DECISION-MAKING

The Human–AI Partnership

Shaping the Next Phase of Cyber Defence

By Majid Ahmed Khan, VP Services - Presales

Security teams have always worked alongside technology. What is changing now is the nature of that relationship. Artificial intelligence is no longer a future capability or a niche enhancement, it is becoming part of the daily operating environment across prevention, detection, and response.

Security analysts are beginning their day with AI-prioritised alerts, AI-generated correlations, and AI-supported investigations. At the same time, attackers are using the same technologies to accelerate reconnaissance, personalise social engineering, and refine their techniques at scale.

This is not a story about replacement. It is a story about partnership.

The next phase of cyber defence will be defined by how effectively human expertise and machine intelligence operate together.

From Automation to Collaboration

For years, automation was introduced to reduce manual effort. Security teams sought faster alerts, quicker triage, and streamlined workflows.

Today, the shift is more fundamental. AI is changing how decisions begin.

Analysts are no longer starting with raw telemetry and building understanding from the ground up. Instead, they are often presented with pre-correlated signals, prioritised incidents, and suggested actions. The starting point of investigation has moved.

This evolution creates a new responsibility.

AI can surface patterns, but it cannot fully understand business context, operational nuance, or organisational risk tolerance. Effective environments recognise this balance:

- AI accelerates analysis and reduces noise.
- Humans apply context, judgment, and intent.

The outcome is shaped by how well these two capabilities are combined.

Where Context Is Created

Cyber defence has always been a problem of interpretation.

Signals rarely arrive with perfect clarity. Alerts compete for attention. Indicators may conflict. Priorities shift depending on business operations, regulatory requirements, and real-world events.

AI excels at recognising patterns across vast volumes of data.

Human analysts excel at understanding why those patterns matter.

When these strengths are combined effectively, organisations gain a more complete picture of risk. Machine insight highlights what could be happening. Human expertise determines what it means and what should happen next.

This partnership becomes particularly important in ambiguous situations — the grey areas where experience, intuition, and operational awareness remain critical.

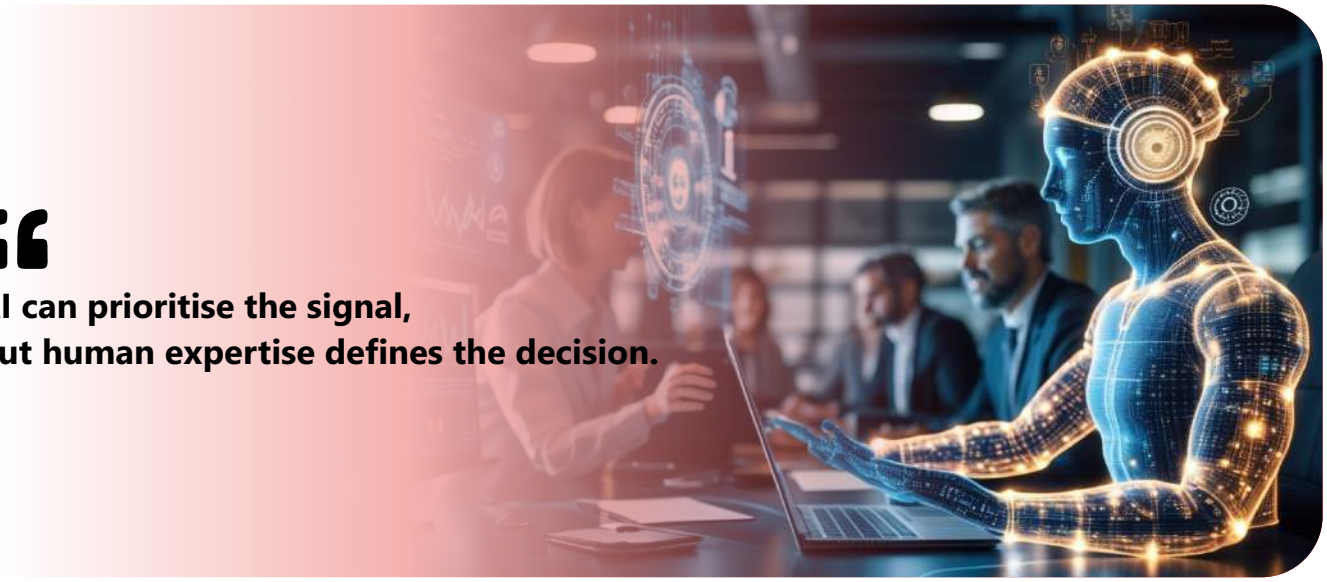

**AI can prioritise the signal,
but human expertise defines the decision.**

Building a Continuous Feedback Loop

The most effective human–AI partnerships are not static, it evolves.

AI models reflect the environments in which they are trained and tuned. As infrastructure changes and threat behaviour adapts, those models require continuous feedback to remain relevant.

Human analysts provide this feedback every day:

- Validating alerts
- Correcting false positives
- Identifying gaps
- Refining detection logic

Over time, this creates a self-reinforcing cycle:

AI accelerates analysis → Humans validate and refine outcomes → Systems improve and adapt

Organisations that treat this as an ongoing cycle develop environments that grow more effective over time, rather than gradually drifting out of alignment.

Scaling the Partnership Across the Ecosystem

Cyber threats rarely stay confined to a single organisation. They move across sectors, supply chains, and national infrastructure.

Across the GCC, collaborative security models are becoming more common, particularly within government and critical infrastructure sectors. Shared intelligence platforms increasingly rely on AI to process and distribute insights at scale.

But interpretation remains local.

AI enables shared visibility across ecosystems. Human expertise ensures those insights are applied in line with regulatory requirements, operational realities, and national priorities.

The strength of ecosystem-level defence depends on maintaining this balance — scale through Agentic Ai, automations, relevance through human judgment.

The Road Ahead

AI will continue to expand what can be processed, analysed, and automated across cyber defence. Its role will deepen across security operations, threat intelligence, vulnerability management, and incident response.

At the same time, the importance of human expertise will not diminish. It will become more focused, more contextual, and more critical to decision-making.

Organisations that lead in this next phase will be those that deliberately design how this partnership operates:

- Defining validation points
- Embedding feedback loops
- Clarifying accountability
- Ensuring human oversight remains central to decision-making

> **The future of cyber defence will be shaped by how well humans and AI learn to think together.**

Final Perspective

Cyber defence is entering a phase where advantage is no longer defined by speed alone.

Organisations should understand how to combine machine intelligence with human judgment to make decisions that are faster, more informed, and grounded in real-world context.

The future of cyber defence will not be built by humans or AI in isolation, but by how effectively they learn to think and act together.

The AI Gateway & AI Firewall

A Strategic Evolution for the CISO

By Mudresh Shah, Manager Solutions Architecture – UAE

Across the region, a growing disconnect is emerging between AI adoption and security readiness. While many organisations have defined AI strategies, far fewer have established dedicated frameworks to secure them.

In the UAE, generative AI is no longer a peripheral capability, it is becoming embedded within critical business and operational environments. According to PwC, 75% of employees in the Middle East have used AI tools at work over the past year, above the global average of 69%, reflecting the region's accelerated pace of AI adoption.[1]

To move from unquantified risk to operational resilience, organisations are increasingly structuring their architectures around two foundational pillars: the AI Gateway and the AI Firewall. At the same time, the human layer of risk remains significant, as employees can unintentionally expose sensitive data when interacting with public AI models.

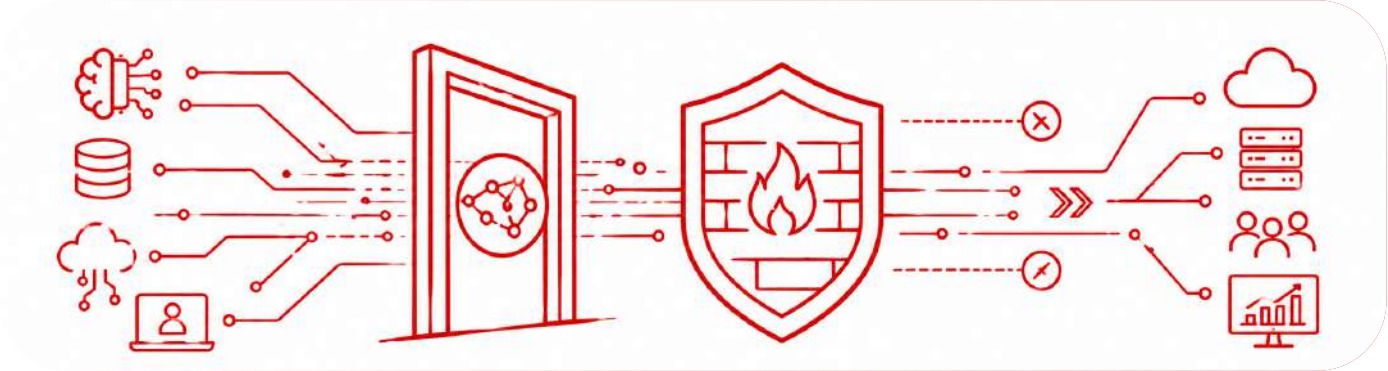

Rather than deploying controls in isolation, organisations are progressing through a structured maturity model that spans visibility, governance, enforcement, identity, and operational scale.

1. Establishing Visibility into AI Exposure

The first stage focuses on visibility. Without it, governance remains limited. Organisations are increasingly identifying Shadow AI by monitoring outbound traffic associated with unauthorised AI usage. Secure Web Gateway and Cloud Access Security Broker telemetry are enabling visibility into where corporate data is being exposed to external AI systems, forming the baseline for controlled adoption.

2. Transitioning to Governed AI Access

As visibility matures, organisations shift toward centralised governance of AI interactions. The AI Gateway is emerging as a unified control layer that standardises authentication, rate limiting, and cost governance across multiple AI providers, reducing fragmentation and improving control.

3. Enforcing Runtime Data Protection

At the next stage, organisations are introducing runtime protection through AI Firewall capabilities. These controls inspect prompts and responses in real time to identify sensitive data exposure and manipulation attempts, including risks such as prompt injection and unintended data leakage.

1. PwC's Middle East Workforce Hopes and Fears Survey 2025

4. Securing the Agentic Identity Layer

As autonomous systems become more prevalent, security boundaries are shifting toward identity. In this phase, AI agents operating through Retrieval-Augmented Generation models are governed under the same Role-Based Access Control principles as human users, with increased focus on monitoring agent behaviour and decision logic.

5. Scaling AI-Enabled Security Operations

At maturity, organisations integrate AI-driven telemetry into security operations environments. In regulated environments where data sovereignty is critical, architectures are evolving toward hybrid and locally governed models, enabling automation while maintaining compliance and control.

This stage increasingly aligns with AI-assisted security operations, improving detection, response speed, and operational consistency.

> **Visibility remains the starting point, followed by the gradual introduction of governed access through AI Gateway and AI Firewall controls.**

Final Perspective

AI adoption is accelerating faster than traditional security architectures can adapt. Organisations should establish visibility first, followed by governed access, enforceable protection, identity control, and operational scale.

In this context, AI security is evolving from a set of controls into a core architectural requirement for operating safely in an AI-driven enterprise environment.

Operationalising this maturity model requires continuous validation across AI gateways, runtime protection layers, and security operations environments. The ability to test, enforce, and adapt these controls under real-world conditions is becoming a defining factor in sustainable AI security.

AI in the Software Development Lifecycle
Scaling Innovation

By Chris Zinn, Principal Cloud & DevSecOps Solution Architect

As AI governance capabilities mature across infrastructure layers, attention is increasingly shifting toward how AI is influencing the software development lifecycle itself. Beyond controlling access to AI models and protecting runtime interactions, organisations are now required to address how AI is reshaping code generation, testing, and deployment processes.

This introduces a new security dimension within DevSecOps environments, where AI is no longer external to development workflows, but embedded within them.

AI is becoming a core part of software development, but the reality inside most DevOps and DevSecOps environments remains far more incremental than widely assumed. While industry narratives often point toward autonomous pipelines and self-securing applications, most organisations are still in an early phase of adoption, applying AI in targeted use cases to improve efficiency rather than fundamentally redesigning delivery models.

At the same time, developers are rapidly adopting AI-enabled tools, driving measurable productivity gains. However, this acceleration is also introducing new risk dynamics, reinforcing the continued importance of established DevSecOps discipline.

Acceleration Without Full Context

One of the most visible shifts is the growing use of Agentic Ai assisted development techniques, where code, scripts, and application components are generated or heavily influenced by AI systems.

While this significantly accelerates development cycles, it also introduces less visible risks, including:

- Increased reliance on patterns that may not be fully understood or validated.
- Embedding vulnerabilities through unchecked open-source dependencies.
- Reinforcement of insecure coding structures through repeated AI-generated outputs.

Over time, these factors contribute to expanding technical debt and a broader software supply chain exposure surface.

Established practices such as Static Application Security Testing (SAST), Software Composition Analysis (SCA), and secure coding standards therefore remain foundational, ensuring that both human- and AI-influenced code is validated before reaching production environments.

AI as an Assistant, Not an Authority

AI-enabled development tools are increasingly evolving into embedded assistants within engineering workflows, supporting activities such as code review, optimisation, and vulnerability detection.

However, these systems remain inherently non-deterministic. Their outputs vary depending on context, training data, and prompting, and therefore cannot be treated as a source of authoritative truth.

In regulated or high-assurance environments, AI-generated or AI-modified code must continue to undergo the same structured review, testing, and approval processes as traditionally developed code, with an increased emphasis on validation and traceability.

Enhancing Testing Within CI/CD Pipelines

AI is also beginning to influence application testing and validation processes within Continuous Integration and Continuous Delivery (CI/CD) environments.

Machine learning-driven capabilities are increasingly being introduced to support:

- Anomaly detection during builds and deployments.
- Vulnerability prioritisation across large codebases.
- Automation of selected elements of security testing workflows.

This evolution aligns with the broader DevSecOps objective of shifting security earlier into the development lifecycle.

However, traditional testing methodologies including static, dynamic, and interactive application security testing — remain essential. AI is best positioned as an augmentation layer that improves signal quality, reduces noise, and accelerates triage rather than replacing existing validation frameworks.

> **AI delivers the most value when it augments existing testing methods rather than attempting to fully automate security validation.**

Expanding the Development Risk Surface

As adoption increases, AI tools are being integrated more deeply into development ecosystems, including direct connections to source code repositories, build systems, and workflow automation platforms.

Without appropriate governance, this creates new risk pathways, including:

- Potential exposure of sensitive code or data to external AI services.
- Increased risk of intellectual property leakage.
- Challenges in meeting data residency and regulatory requirements.

These risks are particularly relevant in regulated GCC environments, where compliance requirements are increasingly shaping technology adoption patterns.

As a result, organisations are beginning to define clearer boundaries around AI usage in development environments, including restrictions on external AI services and increased monitoring of AI-assisted workflows.

From Policy to Embedded Control

In more mature environments, AI risk within DevSecOps is shifting from policy definition to enforcement within the delivery pipeline itself.

This includes:

- Embedding governance directly into CI/CD workflows.
- Establishing visibility into AI interactions with code and data.
- Defining approved pathways for AI-assisted development.
- Strengthening SecOps integration to detect AI-driven anomalies.
- Increasing organisational awareness of secure AI usage patterns.

The emphasis is increasingly on operational enforcement rather than static documentation of acceptable use.

Closing Perspective

The emerging pattern across DevSecOps environments is clear: AI is accelerating software development at a faster pace than corresponding security maturity is evolving.

This creates a widening gap between delivery velocity and assurance capability, expanding both operational complexity and potential exposure.

Organisations that treat AI as a productivity shortcut without governance integration are likely to face increasing security and operational challenges. Those that embed AI within structured DevSecOps frameworks, combining automation, validation, and disciplined engineering practices are better positioned to realise its benefits safely and sustainably.

AI is not replacing DevSecOps. It is increasing its criticality.

What's Next? Agentic AI Enters the Pipeline

Enterprises are adopting agentic AI faster than they are securing it.

Unlike earlier AI assistants that generated code or recommendations, AI agents can take action across development environments, interacting with repositories, calling APIs, executing workflows, and operating inside CI/CD pipelines. This shifts AI risk from information exposure to access, execution, and control.

Over-privileged agents, insecure integrations, and prompt manipulation can quickly expand risk across pipelines, infrastructure, and sensitive codebases.

Securing agentic AI is no longer just about the model. It is about securing identities, permissions, orchestration, tooling, and visibility across the software delivery lifecycle.

With our **AI Security Reference Architecture and Agentic AI Security Blueprint**, Help AG enables enterprises embed AI security directly into modern DevSecOps environments from the start.

The Adaptive SOC

Where AI and Defensive Learning Converge

By Srivatsa Venkatesh, Director of Cyber Defence

Response automation has become foundational to modern SOC operations, improving consistency, scale, and efficiency across detection and response workflows.

But as adversaries compress attack timelines and operate at machine speed, a limitation of this response-centric model is becoming increasingly clear: security still largely begins after detection.

SOC maturity is therefore shifting beyond faster response toward continuous defensive learning, where every investigation strengthens future detection, resilience, and containment.

Automation still operates after a threat has been identified. The sequence remains unchanged: detect, investigate, respond. Automation accelerates execution, but it does not change the underlying operating model.

This limitation is becoming more pronounced as AI-enabled adversaries compress dwell time and adapt faster than traditional detection and response cycles were designed to handle.

The next phase of SOC maturity is no longer about responding faster. It is about building systems that continuously learn from every incident, improve detection coverage in real time, and reduce future attack opportunity with each investigation completed.

The defining question is therefore changing:

Not "How fast did we respond?" but **"How effectively did the SOC learn and how much harder did that make the environment to attack?"**

The Shift From Response-Centric Security to Adaptive Defence

Digital Forensics and Incident Response (DFIR) has traditionally been a post-alert function. An investigation is triggered only after compromise has occurred.

In modern threat environments, timing has become a structural constraint. By the time DFIR processes begin, attackers have often already established persistence, escalated privileges, and moved laterally across environments.

The gap between compromise and investigation is now a structural risk.

This is driving a shift toward embedded DFIR, where forensic capability is integrated directly into the monitoring fabric rather than activated after the fact.

Instead of reconstructing activity retrospectively, organisations maintain persistent forensic context across:

- process execution histories
- identity activity trails
- network session metadata
- system state indicators

This enables detection even when traditional controls produce weak, fragmented, or no alerts.

Detection becomes evidence-led rather than alert-led, driven by behavioural signals, forensic artefacts, and cross-domain correlation.

However, the most important shift happens after the investigation itself.

Forensic outputs have traditionally been treated as closed incident artefacts- documented, reviewed, and archived. In doing so, critical detection intelligence is often lost between investigation and engineering.

In a defensive learning model, this loop is closed.

Every investigation becomes input into detection engineering:

- Attack patterns directly refine detection logic.
- Root cause analysis strengthens behavioural models.
- Each incident reduces the likelihood of recurrence.

Over time, environments become progressively harder to compromise in the same way twice.

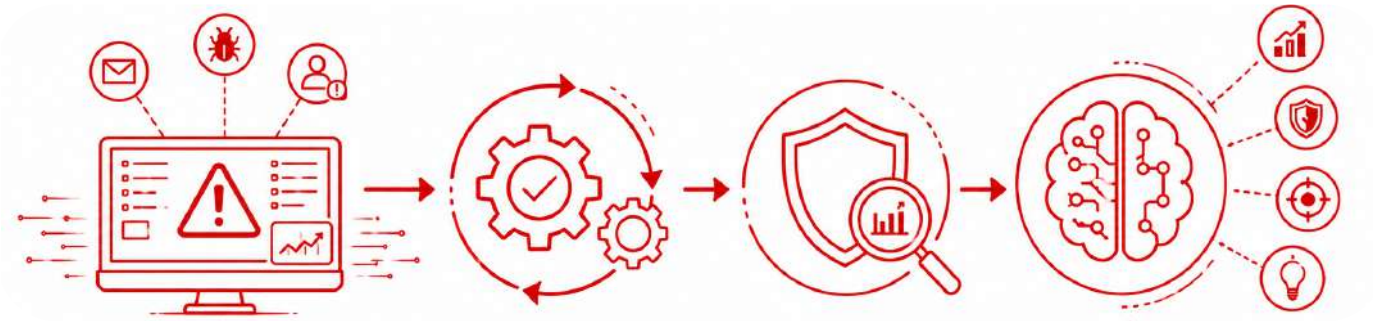

Why Static Detection Engineering Is No Longer Enough

Traditional detection models were built around predefined rules designed to identify known patterns of malicious activity.

This model is increasingly strained in environments defined by cloud-native architectures, ephemeral workloads, distributed identities, and highly dynamic infrastructure.

At scale, its limitations are structural:

- Detection logic remains anchored to known behaviours.
- Maintenance overhead increases with environmental complexity.
- Manual engineering cycles cannot match adversary adaptation speed.

Detection-as-Code marked an important step forward. By applying software engineering discipline to detection logic, rules become version-controlled, testable, and deployable through automated pipelines.

This reduces the lag between identifying a threat and operationalising detection coverage, while improving consistency across environments.

However, static detection logic is reaching its limits.

Rule-based systems remain effective for known indicators and repeatable attack patterns. Increasingly, adversaries design techniques specifically to evade them.

This is where AI-driven detection changes the model fundamentally.

Rather than evaluating isolated events against static rules, AI models analyse sequences of behaviour, contextual relationships, and deviations across users, systems, identities, and workloads.

Combined with AI-powered User and Entity Behaviour Analytics (UEBA), detection shifts from indicator-based matching to behaviour- and intent-driven interpretation.

Detection is no longer rule execution. It becomes continuous contextual analysis.

The result is a detection surface that adapts in near real time, rather than waiting for manual updates after incidents occur.

What Defensive Learning Looks Like in Practice

At Help AG, this shift toward defensive learning is reflected in measurable operational outcomes across our Managed Security Services:

>50%
reduction in detection & response time.

>85%
dwell time reduction for automated alerts.

45 min
average zero-day detection-to-deployment.

<6 hours
zero-day to live use case (90% of cases).

Over 60% of alerts are now handled by automation with human oversight reserved for final review. More than 145 security scenarios are fully automated. CTI turnaround time improved by 30% from Q1 to Q4 2025. These are not isolated efficiency gains, they are signals of a system that improves with every incident it processes."

The SOC Operating Model Is Fundamentally Changing

While implementation maturity varies across organisations, the table below illustrates the broader transition shaping next-generation SOC strategies.

Capability	Traditional SOC	Next-Generation SOC
Detection	Alert-driven monitoring with limited correlation.	Context-aware analysis across integrated telemetry sources.
Detection logic	Static rules and manually maintained signatures.	Adaptive models informed by behaviour, risk and threat context.
DFIR	Isolated and reactive post incident investigative process.	Continuous forensic correlation and automated investigation workflows.
Behavioural analytics	Threshold-based anomaly detection.	AI-powered UEBA with adaptive baselines.
Learning cycle	Periodic reviews and manual tuning.	Operational learning loops that improve detection and response over time.
Primary metric	Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR).	Learning velocity and attack surface reduction.

> **When AI-driven UEBA is integrated with defensive learning, DFIR, and detection engineering, it continuously improves both response and long-term resilience.**

Final Perspective

Response automation marked the starting point - bringing speed, scale, and consistency to SOC operations that were once largely manual.

We are now in the next phase, where AI is reshaping SOC operations from isolated detection and response into continuous, intelligence-driven defence.

AI is shifting the SOC from response execution to continuous defensive intelligence.

Detection, response, and visibility are no longer discrete functions — they are becoming a continuous, AI-enabled system where every investigation directly improves the next outcome.

This is no longer about optimisation. It is about compounding advantage.

A SOC that uses AI not just to respond faster, but to learn continuously, adapt detection in real time, and steadily reduce attacker effectiveness with every interaction.

Resilience is no longer built once, it is learned, refined, and reinforced in real time.

In this model, resilience is not static. It is continuously earned through every signal, every incident, every response.

The next-generation SOC will not be defined by how quickly it reacts.

It will be defined by how effectively AI enables it to learn and how that learning shifts the SOC's ability to detect and defend.

The AI-Powered Service Engine

Redefining How Security Operations Work

By Asim Sharfuddin, SVP Service Operations – UAE

As AI becomes embedded across DevSecOps environments, its influence is extending into downstream operational and service delivery models. This is creating a continuous AI-enabled lifecycle that spans how systems are built, secured, and ultimately operated at scale.

Within the DevSecOps layer, AI is already reshaping how software is developed, tested, and validated. The next stage of this evolution is now emerging in cybersecurity service operations, where AI is redefining how threats are detected, analysed, and responded to in real time.

Cybersecurity service delivery is reaching an inflection point. The scale of modern environments, the speed at which threats evolve, and rising expectations for near-instant response are exposing the limitations of traditional reactive models.

For years, service delivery has been structured around detection and response — identify, investigate, and act. While this remains essential, it is no longer sufficient in isolation.

AI is not simply accelerating this model. It is reshaping it by shifting service delivery from reactive execution toward predictive, continuously optimised operations.

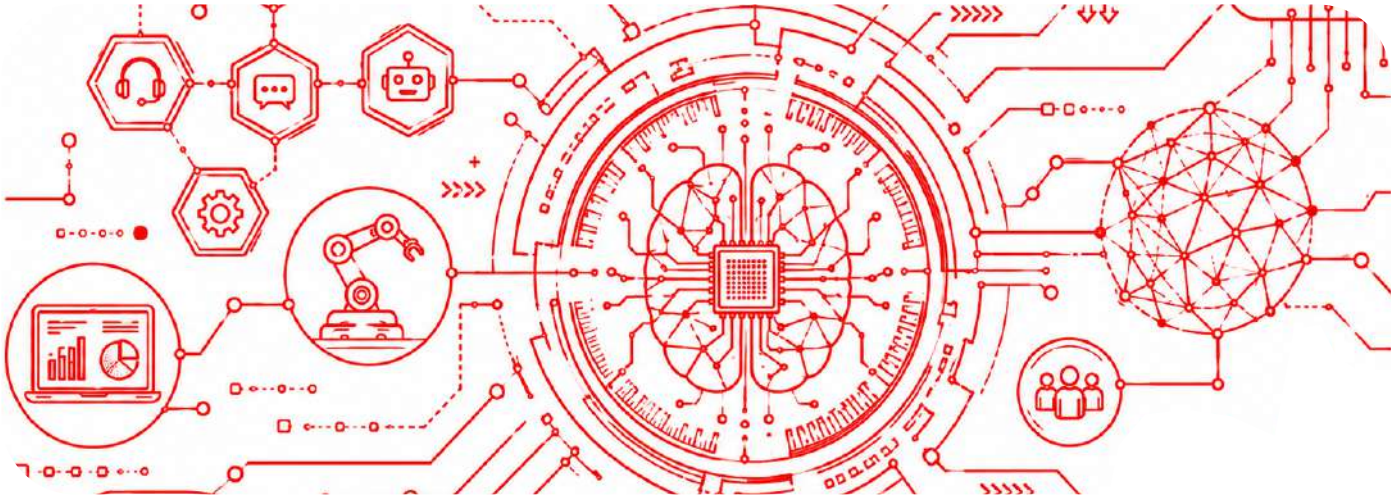

From Automation to an Intelligent Service Engine

AI is often positioned as a tool for efficiency. In practice, its impact is structural. It enables a shift from task-based delivery to an integrated, AI-powered service engine.

This engine operates as a continuous system comprising four interconnected layers:

- Telemetry, providing real-time visibility across environments.
- Decision intelligence, enabling correlation, anomaly detection, and risk prioritisation.
- Execution, supporting automated response and scalable configuration management.
- Governance, ensuring actions remain controlled, auditable, and policy-aligned.

The shift is not in individual capabilities, but in their integration into an adaptive system that continuously improves over time.

> **As services become more autonomous, governance must be embedded by design, ensuring every action remains policy aligned, visible, and auditable.**

Predictive Support in Practice

A key transformation is the move toward predictive support.

Rather than waiting for confirmed incidents, weak signals are now being identified earlier in the lifecycle. These may include behavioural anomalies, configuration drift, or early indicators of compromise.

When correlated at scale, these signals enable service teams to anticipate where failure or compromise is most likely to occur allowing intervention before impact materialises.

This marks a shift from reactive remediation to anticipatory risk management.

Scaling Consistency and Precision

Operational execution remains central, but is increasingly enhanced through automation and AI-assisted decision support.

This includes faster correlation of security telemetry, improved prioritisation of alerts, and reduced manual investigation overhead.

In high-impact scenarios such as Distributed Denial of Service (DDoS) attacks or widespread compromise indicators, automation enables rapid and consistent deployment of protective controls across multiple environments simultaneously.

This ensures speed of response without introducing operational inconsistency or additional risk.

Beyond incident response, capabilities such as zero-touch deployment, parallel upgrades, and automated configuration management are enabling more scalable and reliable service delivery.

Telemetry from operational environments further feeds into continuous service refinement, allowing delivery models, processes, and knowledge systems to evolve over time through validated operational learning.

What proves effective in one environment is validated and applied more broadly, creating a progressively improving operational baseline.

The Evolving Role of Service Teams

As service models evolve, so does the role of security teams.

The focus is shifting away from manual execution toward supervision of intelligent systems, exception management, and higher-value analytical decision-making.

This represents a reallocation of expertise rather than a reduction in human involvement enabling specialists to focus on areas where judgment and experience deliver the greatest value.

Redefining Service Delivery

The value of AI in cybersecurity is often measured in speed. The real transformation lies in how services are delivered.

From reactive operations to predictive support.

From manual execution to intelligent orchestration.

From isolated actions to adaptive, continuously optimised systems.

This is the foundation of the AI-powered service engine, not simply accelerating cybersecurity delivery, but fundamentally redefining it.

> **Service delivery becomes progressively more intelligent when operational telemetry continuously feeds back into and refines delivery models.**

Final Perspective: A Unified AI Security Lifecycle

Across the State of the Market report series, a clear architectural model emerges.

At the control layer, AI governance mechanisms define how AI is accessed and secured. At the development layer, DevSecOps environments are adapting to manage AI-influenced software creation and associated risk. At the operational layer, service delivery is evolving toward predictive, AI-enabled cybersecurity operations.

Together, these layers form a continuous AI security lifecycle spanning control, creation, and operations.

In this model:

- AI governance defines how AI is used and controlled.
- DevSecOps defines how AI influences what is built.
- Service operations define how AI-driven environments are operated and defended.

This represents a shift from isolated security functions to an integrated, lifecycle-based security architecture designed for AI-driven enterprises.

AI Governance in the GCC

From Compliance to Operational Capability

By Talal Wazani, Head of Cyber Trust Advisory

A Region Moving Beyond Policy Framing

AI governance across the GCC is entering a more execution-driven phase. It is now directly linked to how organisations scale AI, operate across regulatory environments, and maintain consistency in how systems are governed in practice.

This shift aligns closely with national priorities. Initiatives such as the UAE's AI Strategy 2031 and the Saudi Data and Artificial Intelligence Authority under Vision 2030 reflect a broader transition toward data-driven economies, where AI is positioned as both an innovation driver and a strategic national capability.

Regulatory structures are evolving in parallel. In the United Arab Emirates, governance is shaped through a layered ecosystem of national strategies, policies, and sector-level initiatives. In Saudi Arabia, a more centralised, risk-based model is emerging, anchored by the Personal Data Protection Law and supported by national AI frameworks and ethical principles.

For organisations operating regionally, the implication is clear: governance must operate consistently across multiple regulatory environments while remaining executable at scale across day-to-day AI operations.

From Oversight to Embedded Execution

Historically, AI governance has functioned as an oversight layer. Policies were defined centrally, assessments conducted at specific stages, and controls applied at predefined checkpoints. While this establishes structure, it does not reflect the pace or adaptability of modern AI systems.

That model is no longer sufficient.

Governance is now moving closer to where systems are designed, deployed, and operated. It is becoming embedded within execution rather than applied around it.

This shift changes how control is exercised and where accountability sits. Monitoring is moving from periodic review to continuous visibility of system behaviour, performance, and exposure to risk. This is increasingly critical as issues such as model drift, bias, and unintended outcomes emerge over time.

Effectiveness is no longer defined by the existence of policies, but by the consistency with which governance operates in live environments.

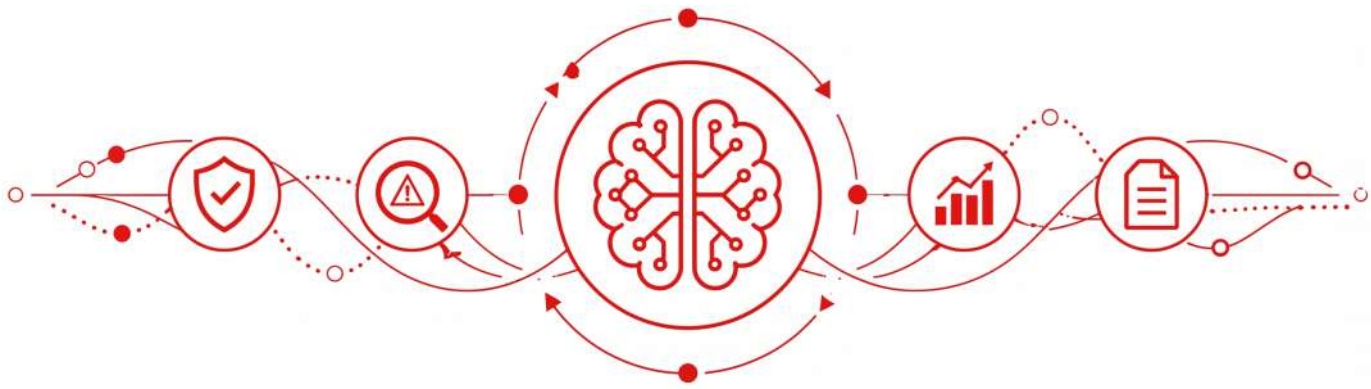

The Structural Challenge: Fragmentation Across Jurisdictions

As AI adoption scales, governance is no longer about alignment to a single framework. It is about maintaining consistency across multiple systems, teams, and regulatory environments simultaneously.

In practice, many organisations still operate fragmented AI initiatives. Governance approaches differ across business units, markets, and use cases. Within the GCC, this is further amplified by varying regulatory models, data requirements, and national priorities.

The result is uneven governance application where controls are strong in some areas, duplicated in others, and inconsistent elsewhere. This creates blind spots, regulatory drift, and reduced visibility across the enterprise.

The core challenge is not regulatory alignment itself. It is the ability to sustain a coherent operating model while adapting to local requirements without fragmentation.

The maturity gap in AI is increasingly defined by this issue, not model performance, but governance consistency at scale.

Toward a Unified Operating Model

Addressing this challenge requires a shift from isolated governance efforts to a unified operating structure.

Frameworks such as Artificial Intelligence Management Systems (AIMS), as defined in ISO/IEC 42001, provide a foundation for this transition. Their value lies not in compliance alone, but in enabling governance to operate consistently across jurisdictions.

In practice, this allows organisations to establish a common governance baseline while adapting to UAE and Saudi regulatory expectations without duplicating effort or creating silos.

Governance is no longer applied to individual AI initiatives. It is structured as a system that operates across them.

Operationalising Governance: A Layered Approach

Leading organisations are increasingly structuring AI governance across three interconnected layers:

The Control Layer defines accountability, ownership, and decision rights. It embeds governance within organisational structures while remaining adaptable to both centralised and distributed regulatory models.

The Lifecycle Layer integrates governance across the full AI lifecycle from data sourcing and model development through to deployment and continuous monitoring. This ensures consistent application of controls regardless of where systems operate.

The Adaptation Layer enables governance to evolve alongside regulatory developments and national strategies. This is particularly important in the GCC, where policy and AI frameworks continue to advance rapidly.

Together, these layers enable governance to function as a continuous operational discipline, unified in structure, but flexible in execution.

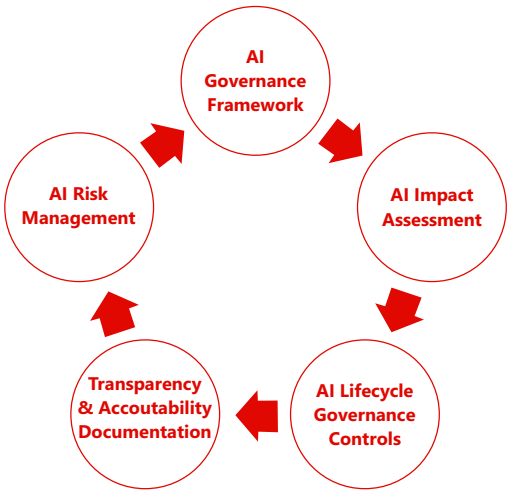

Continuous and Context-Aware Governance

As AI systems evolve post-deployment, governance models are shifting toward continuous, context-aware operation. Static assessments are being supplemented by real-time monitoring of performance, risk indicators, and compliance signals.

This reflects the dynamic nature of AI, where changes in data, usage, and environment directly influence system behaviour. Governance must therefore operate at a comparable pace to remain effective.

Across the GCC, this shift is expected to accelerate as regulatory ambition and digital transformation continue to converge.

The current landscape is redefining control and accountability, shifting monitoring from periodic review to continuous visibility of risk and system behaviour.

Final Perspective

AI governance in the GCC is evolving into a core operating discipline that enables organisations to deploy and scale AI consistently across markets, sectors, and regulatory environments.

The distinction is structural. Governance is no longer defined by control points or documentation but by how effectively it is embedded into execution.

In this model, success is measured by consistency: across jurisdictions, across systems, and across the lifecycle of AI itself.

Organisations that establish governance as a unified operating capability will be best positioned to scale AI responsibly maintaining control, trust, and regulatory confidence as complexity increases across the region.

ARCHITECTURE

ENGINEERING RESILIENT DIGITAL SYSTEMS

Identity as the Control Plane

Securing the Autonomous Enterprise

By Senthil Santhanam, Solution Architect

Enterprise environments are increasingly defined by how access is created, distributed, and used across systems. What was once considered a supporting function now underpins how organisations operate at scale.

As cloud adoption, platform integration, and automation expand, access is no longer tied solely to individuals. It is embedded across services, applications, and operational processes, each interacting through assigned permissions.

The most efficient path through an enterprise environment is no longer through infrastructure, but through the identities that connect it. Movement across systems increasingly depends not on breaching them, but on navigating the access already in place.

Identity has become the foundation through which digital operations are executed and interconnected.

Identity at Scale

This evolution is not only about the growth of identity, but how it now behaves within enterprise environments.

Access is increasingly distributed across systems, embedded within applications, and created as part of how digital services are built and delivered. In many organisations, non-human identities already outnumber human users by a significant margin, operating continuously and often without direct visibility.

A growing share of enterprise activity is carried out through identities embedded within workflows, integrations, and automated processes.

This creates an environment where access is persistent, distributed, and continuously active. When access operates across systems at all times, the distinction between expected use and unintended use becomes less clear.

Identity therefore moves beyond a control point. It becomes the mechanism through which systems interact, processes execute, and outcomes are produced.

The Agentic Shift

AI agents are extending identity beyond automation into autonomous execution.

Unlike traditional systems that perform predefined tasks, agentic systems can decide, initiate, and coordinate actions across multiple environments with limited human intervention.

They do not simply authenticate into systems, they operate within them using credentials, permissions, and interconnected access pathways.

This fundamentally changes identity risk.

Agentic identities are dynamic by design. Their access patterns shift based on context, objectives, and system interactions. A single agent may interact across compute, storage, analytics, and financial systems simultaneously at machine speed.

Static role-based governance models and inherited permissions are no longer sufficient in this environment.

Identity governance must now extend beyond access control into control of autonomous action itself.

From Access Management to Identity Operations

In response, organisations are rethinking how identity is managed as part of core operations. This shift is increasingly being operationalised through structured identity governance models that bring visibility, control, and response together across all identity types.

Leading organisations are moving beyond static access models toward continuously evaluated identity environments.

Visibility is becoming continuous rather than periodic, extending to identities operating in the background of systems and workflows. Ownership is becoming explicit, ensuring system-level identities are tied to defined purpose and accountability.

Access itself is becoming adaptive, shaped by behaviour, environment, and usage rather than granted once and left unchanged. Long-lived access is being replaced with time-bound, purpose-specific permissions that limit how far access can extend if misused.

This reflects a broader shift from managing permissions to operating identity as a continuous control layer.

A Layered Approach to Identity Resilience

Unified Visibility
Real-time view of all identities.

Contextual Access
Access adapts to risk and behaviour.

Credential Control
Access is time-bound and regularly rotated.

AI Identity Governance
Agents operate with defined scope and ownership.

Agent Behaviour Control
Actions align with expected patterns.

Identity Detection and Response
Misuse is identified and contained quickly.

Ecosystem Control
Third-party access is equally governed.

The GCC Context

Across the GCC, rapid digital transformation is accelerating cloud adoption, AI integration, and highly interconnected digital ecosystems.

Sectors such as financial services, energy, logistics, and government are increasingly operating through automated, data-intensive environments where identity activity continues to expand in both scale and complexity.

This shift is elevating identity resilience from an operational requirement to a strategic priority.

As organisations introduce AI-driven and agentic capabilities, identity governance is becoming central to maintaining control, accountability, and trust across distributed environments.

The growing reliance on third-party ecosystems, cross-border platforms, and digital public services further increases the need for consistent governance across human, machine, and autonomous identities.

At the same time, regulatory expectations across the region are evolving toward stronger oversight of access governance, monitoring, and AI accountability.

In this context, identity is becoming the control layer through which enterprise operations are governed, secured, and scaled.

> **At scale, identity is not a control point to manage, but the operational layer through which systems act, decisions are executed, and risk continuously evolves.**

Final Perspective

As digital environments continue to expand, identity is becoming the control layer through which enterprise operations function.

Access is now continuous, distributed, and deeply embedded in how organisations operate.

Consistency, accountability, and control across this landscape are no longer a one-time initiative, they are an ongoing operational capability.

In an agentic environment, this challenge intensifies.

The focus shifts beyond what identities exist and what access they hold to what those identities are authorised to decide and initiate, autonomously and at speed.

Organisations that recognise this shift move beyond fragmented approaches toward a unified model for governing identity across the enterprise.

In the modern enterprise, identity is the operating foundation of stability, trust, and control — not a supporting function.

The Multi-Cloud Control Paradox

The Growing Gap Between Scale and Cohesion

By Majid Ahmed Khan, VP Services - Presales

For years, enterprise cloud strategy has been guided by a consistent assumption: distributing workloads across multiple cloud providers improves resilience, flexibility, and innovation, while leveraging native cloud security capabilities to protect each environment.

In principle, this model aligns with enterprise objectives by combining diversification with operational scale and assuming that stronger platform-level controls would translate into stronger enterprise-wide security outcomes.

However, as multi-cloud adoption matures, a structural limitation is becoming increasingly evident. Strong security controls within individual cloud environments do not automatically translate into a unified security posture across the enterprise. Instead, organisations are experiencing a widening gap between localised visibility and system-wide risk understanding.

This defines the multi-cloud control paradox: as the number of environments increases, the ability to maintain a coherent security perspective decreases disproportionately, exposing a structural limitation in how distributed cloud architectures are governed.

Geopolitical and Regulatory Forces Reshaping Cloud Architecture

This evolution is no longer driven purely by technical design decisions. Increasingly, multi-cloud adoption is being shaped by geopolitical fragmentation, regulatory divergence, and accelerating expectations around digital sovereignty.

Across many regions, cloud strategy is now influenced by sovereign requirements that extend beyond data residency into jurisdictional control, operational transparency, and regulatory accountability. In parallel, global cloud ecosystems are becoming structurally segmented due to export controls on advanced technologies, evolving restrictions on cross-border data flows, and the emergence of regionally aligned digital infrastructure models.

The acceleration of AI adoption further intensifies this structural shift. AI systems operate across cloud boundaries by design, interacting dynamically with data, APIs, and workloads. This introduces a growing tension between distributed execution models and increasing expectations around control, auditability, and sovereignty.

Within this context, multi-cloud is no longer simply an optimisation strategy. It has become a structural response to geopolitical and regulatory conditions that are actively shaping how digital infrastructure is designed and operated.

Why the GCC is Experiencing This Challenge More Intensely

In the GCC, this structural challenge is emerging more rapidly due to the pace and simultaneity of digital transformation across both public and private sectors. Governments and enterprises are expanding sovereign cloud adoption, modernising infrastructure, scaling digital services, and deploying AI capabilities in parallel rather than sequentially. This creates distributed architectures from inception, rather than as an evolutionary outcome.

Key market signals reinforce this trajectory:

- 88% of organisations now operate across multiple cloud environments Flexera[1]
- Sovereign cloud spending in MENA is projected to grow by 89% in 2026 Gartner[2]
- AI adoption across the Middle East continues to exceed global averages PwC[3]

1. Fortinet Cloud Security Report 2026
2. Gartner Sovereign Cloud IaaS Spending (Press Release) 2026
3. PwC Middle East Workforce Hopes and Fears Survey 2025

Taken together, these indicators reflect not incremental change, but a structural redesign of digital infrastructure across the region.

This has direct implications for security operations. Sovereign cloud requirements are expanding beyond data residency into expectations for consistent governance, unified policy enforcement, continuous operational visibility, and cross-environment compliance assurance. Security can no longer be evaluated within individual environments; it must function as a coordinated capability across the entire digital estate.

Multi Cloud has Changed the Nature of Security

Cloud-native security platforms today are highly mature, offering deep telemetry, identity intelligence, workload visibility, and advanced detection capabilities embedded within their ecosystems.

The limitation is not capability, it is architectural scope.

These platforms are optimised to deliver visibility within environments rather than across them.

Modern enterprises now operate across multiple cloud providers, SaaS platforms, hybrid infrastructure, partner ecosystems, and AI-driven operational layers. Within these architectures, workloads, identities, and data move dynamically across boundaries in ways that are continuous rather than discrete.

Threat actors exploit the same structural reality. Movement across environments is no longer exceptional, it is inherent to both legitimate and malicious activity patterns.

As a result, what appears within individual environments as isolated security events may in fact represent fragments of a broader, coordinated attack path. An identity compromise in one cloud, anomalous workload behaviour in another, and data movement across a third may each be interpreted independently, despite forming part of a single intrusion chain.

The underlying issue is therefore not detection capability within environments, but the absence of consistent cross-environment correlation that enables system-level interpretation.

Why Current Approaches Often Increase Complexity

Despite recognition of this challenge, most organisational responses still operate within established paradigms that do not fully address the architectural nature of the problem.

One approach is the introduction of additional visibility or posture management platforms intended to centralise monitoring across environments. While this improves coverage, it often increases integration overhead and operational complexity, adding another abstraction layer without resolving cross-environment correlation.

A second approach involves consolidation around a single cloud provider to reduce fragmentation. While this simplifies governance in the short term, it introduces dependency concentration risk and reduces access to differentiated capabilities across ecosystems, limiting architectural flexibility.

A third approach relies on custom integration frameworks designed to connect disparate environments. Although tailored, these systems often struggle to scale as cloud services evolve, APIs change, and operational complexity increases over time.

Across all three approaches, a consistent pattern emerges: multi-cloud security is still being treated as a tooling problem, when it is fundamentally an architectural coordination challenge.

> **What appears as isolated alerts across multiple clouds is often a single coordinated attack unfolding across the enterprise.**

The Missing Layer in Modern Cloud Security

The next phase of cloud security evolution will not be defined by replacing cloud-native capabilities. It will be defined by coordinating them.

What is increasingly required is an operational coordination layer capable of unifying telemetry across cloud environments, correlating activity into cross-domain attack paths, enforcing governance consistently, orchestrating response actions across systems, and providing a consolidated view of enterprise risk.

This layer does not replace existing platforms. It integrates them into a unified operational model that allows distributed signals to be interpreted within a single system-level context.

Importantly, this model strengthens rather than replaces cloud-native security platforms by embedding them within a broader operational framework that enables coordinated interpretation and response.

The objective is no longer to secure individual cloud environments in isolation, but to establish security across them as a single distributed operating system.

> **The industry spent the last decade scaling cloud adoption. The next decade will be defined by securing cloud interdependence.**

Final Perspective

Cloud security is shifting from environment-level protection to system-level coherence. As enterprises operate across multi-cloud, AI-enabled, and sovereign-regulated infrastructures, the primary constraint is no longer visibility within environments, but consistency across them.

In this context, fragmentation becomes more than an operational inefficiency, it becomes a structural constraint that limits how effectively organisations can interpret, correlate, and respond to risk across distributed environments.

Multi-cloud adoption was originally driven by the need for flexibility, resilience, and strategic optionality. Realising those outcomes at scale now depends on a more fundamental capability: the ability to operate and secure multiple environments as a coordinated whole. That shift defines the next phase of cloud security.

From Cloud Adoption to Sovereign Control

The Rise of the Landing Zone Model

By Sriram Gopalakrishnan, Director - Strategic Product Management

Across the Middle East, digital transformation is entering a more structural phase. The focus is no longer on adopting cloud, data, AI, but on operating them at scale within clear regulatory and strategic boundaries.

As environments expand, complexity is increasing faster than control. Capabilities scale independently, but governance, visibility, and enforcement often do not scale at the same pace.

This challenge is emerging at a time when cloud adoption across the region is accelerating rapidly, with GCC cloud spending expected to exceed $20 billion in the coming years, and more than 89% of enterprises now operating in multi-cloud environments[1]. As scale increases, so does the difficulty of maintaining consistent control.

This is where sovereignty is being redefined. It is no longer a constraint applied to data, but a design principle that shapes how digital environments are built and operated.

The sovereign landing zone is emerging as the practical mechanism for delivering this model.

> **Sovereignty provides the mechanism through which scale can be achieved without loss of control.**

Why Sovereign Landing Zones Matter Now

Sovereign landing zones emerged to address a gap created by early cloud adoption.

Initial cloud strategies prioritised speed and scalability, while governance and control were added later as separate layers. As environments expanded across multiple platforms and jurisdictions, this model began to fragment.

At the same time, regulatory expectations in the region became more defined, particularly around data residency, operational control, and critical infrastructure protection. Frameworks such as the NCA controls in Saudi Arabia, SAMA cybersecurity requirements, and UAE data protection regulations are increasingly shaping how digital environments must be designed and operated.

What became clear was the absence of a consistent control model that could scale with the environment.

The sovereign landing zone addresses this by embedding governance, security, and compliance into the foundation of cloud environments not as overlays, but as built-in design principles.

It has become a priority because it sits at the intersection of three forces:

- Rapid cloud and AI adoption
- Tightening regulatory requirements
- The need for scalable, consistent control

Sovereignty is no longer a policy discussion. It is now an architecture requirement.

1. Flexera. State of the Cloud Report 2024

Sovereignty as an Operating Model

In practice, sovereignty extends well beyond data residency. It reflects the ability to maintain control across the full digital environment — from data and identity to operations and response.

It operates across multiple dimensions:

- Data sovereignty – where information resides and is processed.
- Operational sovereignty – how access, change, and response are controlled.
- Cyber sovereignty – where monitoring and threat response are executed.
- Legal sovereignty – how environments align with national regulatory frameworks.

Individually, these are controls. Together, they define an operating model.

Sovereignty is therefore not a point-in-time requirement. It is a continuous state of control across the lifecycle of the environment.

From Landing Zone to Control Plane

Traditionally, sovereign landing zones are positioned as secure foundations for cloud adoption. That framing is now too limited.

At scale, control cannot depend on static configurations distributed across identity, infrastructure, data, and security layers. These domains must operate as a coordinated system.

In this model, the landing zone becomes a control plane, not just a deployment framework. Governance is embedded by design and enforced continuously through policy, automation, and monitoring.

This shifts control from something implemented at deployment to something that is maintained throughout operations.

Control is Only as Strong as its Coordination

The effectiveness of a sovereign landing zone is determined by how well control is aligned across layers.

- Identity defines who can act.
- Data policies define what can be accessed.
- Encryption defines how information is protected.
- Security operations define how activity is monitored and responded to.
- Network controls extend enforcement across connectivity.

When these operate independently, control fragments over time.

When they operate together, control becomes systemic.

This alignment closes the gap between policy and execution where most governance breakdowns typically occur in large-scale environments.

> **As digital systems become more distributed, sovereignty shifts from a static requirement to a continuous, coordinated capability embedded across the environment.**

Sovereignty Extends Beyond a Single Platform

Sovereign environments are not built by a single technology layer. They are the result of coordination across an ecosystem.

Governments and regulators define policy direction.

Cloud providers deliver scalable infrastructure.

Telecom operators and security providers operationalise enforcement.

The effectiveness of sovereignty depends on how well these layers converge in practice. Where alignment is strong, control is consistent. Where it is fragmented, gaps emerge in visibility, accountability, and enforcement.

Sovereignty is therefore not delivered by one entity, it is achieved through ecosystem design.

The Role of Telecom-Integrated Security Models

The convergence of connectivity and security is becoming a defining factor in sovereign architectures.

Telecom operators sit at a critical point in this model, anchoring national connectivity, supporting data infrastructure, and operating within regulatory frameworks. When combined with managed security capabilities, they extend control closer to the source of activity.

This enables visibility from network ingress through to application workloads, reducing reliance on isolated enforcement points within cloud or security platforms.

As environments integrate cloud, AI, and digital services, this model strengthens consistency of control across layers and reduces operational fragmentation.

Why This Shift Matters Now

The move toward sovereign landing zones is being driven by a set of converging forces:

- Regulatory expectations are becoming more explicit in defining how environments must operate.
- AI and data-intensive workloads are increasing the need for trusted execution environments.
- Threat activity continues to evolve in both scale and precision.
- National strategies are prioritising resilience, autonomy, and operational continuity.

Individually, these pressures are manageable. Together, they require a structured approach to control at scale.

Without it, scale introduces fragmentation faster than it delivers value.

Sovereignty provides the mechanism to prevent this drift.

Final Perspective

The sovereign landing zone is not a one-time cloud architecture construct. It is an operating model that must continuously evolve alongside the environments it governs.

As digital ecosystems expand across cloud, AI, and distributed infrastructure, the challenge is no longer simply enabling scale — it is sustaining control within that scale without introducing fragmentation.

In this context, sovereignty becomes a design constant rather than a compliance outcome. It defines how environments are structured, how controls are enforced, and how consistency is maintained as complexity increases.

Organisations that embed sovereignty into their foundational architecture will be better positioned to maintain resilience, regulatory alignment, and operational control not as separate objectives, but as integrated outcomes of the same system.

SASE in the Middle East

Bridging Architecture and Execution

By Kosta Skulic, Senior Solutions Architect - UAE

Across the Middle East, digital transformation is advancing at structural scale, driven by national digital agendas, cloud-first mandates, and large-scale smart infrastructure programmes.

This reflects more than increased digital activity. It represents a shift in operating conditions, where applications, users, and data are continuously distributed across hybrid and cloud environments, and where security is embedded directly into how modern systems function.

Secure Access Service Edge (SASE) has emerged as a dominant architectural direction in response to this shift, converging networking and security into a unified, cloud-delivered model designed for distributed enterprise environments.

Yet across the Middle East, SASE adoption is not a linear transition toward fully converged architectures. Instead, organisations are operating in persistent hybrid environments where legacy infrastructure, regulatory obligations, and modern cloud frameworks coexist.

This is not an adoption gap driven by investment or awareness. It reflects structural constraints at the intersection of architecture, regulation, and operational reality.

Most organisations are operating in a deliberate middle state; partially converged, partially modernised, but not fully unified.

This middle state is not simply a temporary transition phase. In many cases, it has become the long-term operating model through which organisations balance agility, governance, and infrastructure modernisation simultaneously.

Architecture vs Operational Reality

The challenge is that many global SASE frameworks are designed around assumptions of architectural uniformity and operational standardisation.

These assumptions do not consistently hold in the Middle East.

Operating environments are shaped by national frameworks such as the National Electronic Security Authority (NESA), the Dubai Electronic Security Center (DESC), and the Abu Dhabi Healthcare Information and Cyber Security Standard (ADHICS). These frameworks directly define:

- Data residency requirements.
- Inspection boundaries.
- Identity governance models.
- Accountability structures.

Security is therefore shaped by regulatory and operational requirements rather than implemented independently of them. As organisations modernise, cloud, on-premises, and legacy environments must often coexist across different stages of transformation.

Convergence is therefore incremental rather than uniform, making operational consistency a greater challenge than architectural design.

> **SASE was designed for simplicity, but in high-growth, regulation-driven markets, it collides with complexity that cannot be standardised or ignored.**

Fragmentation as an Operational Condition

Beyond architecture, the regional environment is defined by velocity.

Policy enforcement becomes distributed across multiple control points, leading to inconsistent access decisions and increased operational overhead. Visibility remains uneven, particularly across encrypted traffic flows where inspection is fragmented across tools and platforms. Identity enforcement also remains inconsistent, with least-privilege models not uniformly applied across cloud, on-premises, and hybrid environments.

As a result, environments designed for convergence can still operate as partially disconnected security domains in practice.

Velocity as a Defining Constraint

The regional operating environment is defined by velocity rather than steady-state conditions.

Organisations are simultaneously managing rapid digital expansion, hybrid workforce distribution, cross-border service delivery, and large-scale national programmes.

Legacy Virtual Private Network (VPN) architectures continue to coexist with emerging SASE deployments, creating overlapping policy domains and duplicated control structures.

In this context, static policy models are increasingly misaligned with operational speed.

This is increasing demand for adaptive architectures capable of making context-aware security decisions in real time.

Convergence as Sequencing, Not Scale

As organisations attempt to reduce this operational complexity, the challenge becomes less about technology acquisition and more about the order in which convergence is operationalised.

Adoption typically begins through one of three paths:

- Identity-first models that establish Zero Trust principles and identity-centric governance as the primary control layer
- Network-first convergence that integrates SD-WAN with security inspection to simplify routing and improve policy consistency
- Data-first approaches that leverage Data Security Posture Management (DSPM) to align access and inspection with data sensitivity

The key constraint is not capability availability, but simultaneous convergence across all domains. Attempting full convergence simultaneously often increases operational complexity without improving integration outcomes. Effective convergence therefore depends on sequencing transformation in line with operational priorities and infrastructure realities.

From Convergence to Adaptive Control

The next phase of SASE evolution is shifting from architectural convergence to operational intelligence.

Market direction indicates increasing adoption of AI-enabled security operations to address scale, speed, and skills constraints.

This marks a shift in operating models:

From static policy systems, where rules are defined and periodically updated, to adaptive systems where decisions continuously adjust based on behavioural and contextual signals.

This transition moves SASE from a control framework to a dynamic security fabric.

At the same time, integration with DPSM is becoming more relevant, ensuring that access decisions reflect both identity and data sensitivity.

Execution as the Differentiator

A consistent gap is emerging between platform capability and operational execution.

While global vendors provide advanced SASE capabilities, these platforms are not inherently designed to account for regional regulatory structures, sovereignty requirements, or hybrid infrastructure dependencies.

The challenge is therefore not capability but consistent operationalisation across distributed environments.

Bridging this gap requires an operational layer that translates architectural design into consistent enforcement across distributed environments, aligning:

- Policy enforcement
- Identity governance
- Data visibility
- Operational telemetry
- Compliance frameworks

Architecture defines strategic direction; execution determines operational effectiveness.

> **The real gap in SASE adoption is not technology capability, but the absence of an execution layer that translates architecture into operational reality.**

Final Perspective

The Middle East SASE landscape is defined not by lack of convergence, but by a structurally embedded hybrid reality shaped by regulation, infrastructure maturity, and sustained digital velocity.

The defining challenge is no longer whether organisations will converge networking and security but how consistently that convergence can be operationalised across fragmented, distributed environments.

SASE is therefore shifting beyond architectural convergence toward adaptive operational models that respond to context in real time.

Leadership in this phase will not be defined by alignment to reference architectures.

It will be defined by the ability to operationalise them in environments they were never designed to assume.

Securing the Hyper-Connected Metropolis

OT & IoT Security in the Smart City Era

By Ahmad Ali, Manager Solutions Architecture - KSA

The smart city vision — where traffic systems, power grids, water networks, and public services operate as an interconnected ecosystem — is no longer a future concept. It is becoming today's urban reality.

As global investment in smart city initiatives continues to grow, the convergence of Internet of Things (IoT), Industrial Internet of Things (IIoT), and Operational Technology (OT) is significantly expanding the cyber-physical attack surface. The global smart cities market size was estimated at USD 877.6 billion in 2024 and is projected to reach USD 3,757.9 billion by 2030, growing at a CAGR of 29.4% from 2025 to 2030.[1]

For CIOs, CISOs, and government stakeholders, this shift introduces a structural change in risk: cyber incidents are no longer confined to IT disruption. They can directly impact physical systems, public services, and citizen safety.

In this context, resilience is increasingly defined by the ability to maintain operational continuity across interconnected digital and physical environments.

The Expanding Risk Surface in Converged OT and IoT Environments

Unlike traditional enterprise systems, smart city environments combine modern digital infrastructure with legacy OT systems that were never designed for external connectivity.

This creates two critical risk dynamics:

1. Large-scale attack surface expansion

Millions of IoT devices operate across urban environments, often with inconsistent patching cycles, limited authentication controls, and fragmented ownership models.

2. Exposure of cyber-physical systems

In OT environments, cyber compromise can translate into physical impact. Attack paths increasingly include:

- Manipulation of sensor telemetry (false environmental or operational readings).
- Interruption of actuator commands (affecting physical infrastructure behaviour).
- Exploitation of legacy industrial protocols lacking encryption or authentication.

In these environments, availability is often prioritised over confidentiality. However, this creates a critical tension: systems designed for uptime are now exposed to threats that target integrity and control.

As a result, traditional IT-centric incident response models focused on data recovery are insufficient when disruption affects real-world services.

From Perimeter Security to Embedded Cyber Resilience

Securing smart city infrastructure requires moving beyond perimeter-based defence toward embedded resilience across the operational stack.

1. Grand View Research Smart Cities Market 2025-2030

This shift is increasingly being driven by three foundational security domains:

1. Device and Protocol-Level Visibility

Security in IoT and OT environments depends on visibility that extends beyond traditional endpoints. This includes:

- Continuous monitoring of device behaviour across industrial protocols (e.g. Modbus, DNP3, OPC-UA).
- Detection of anomalous command sequences or unexpected device states.
- Identification of unmanaged or legacy devices operating outside security baselines.

This level of visibility is essential because many OT environments lack native logging or security telemetry.

2. End-to-End Cryptographic Protection

In smart city architectures, data flows continuously between edge devices, control systems, and central command platforms.

Without end-to-end encryption, these flows become vulnerable to:

- Man-in-the-middle manipulation of telemetry data.
- Injection of false sensor readings into operational systems.
- Unauthorised interception of control commands.

Cryptographic protection across device authentication, data transmission, and storage ensures that operational decisions are based on trusted and verifiable inputs even in partially compromised networks.

3. Data Integrity and System Trust

Beyond encryption, modern OT and IoT environments require assurance that data has not been altered during transmission or storage.

This is achieved through:

- Immutable logging mechanisms.
- Hardware-backed attestation of device outputs.
- Cryptographic verification of sensor and system data.

This ensures that operational decisions, regulatory reporting, and incident investigations are based on verified data integrity a critical requirement in safety-sensitive environments such as energy, transport, and utilities.

> **Securing smart cities requires a shift from perimeter defence to embedded resilience across connected systems.**

Shared Responsibility Across the Smart City Ecosystem

Smart city environments operate through interconnected ecosystems involving government entities, infrastructure operators, and private sector technology providers. This creates a distributed security responsibility model.

1. Government stakeholders

Security requirements are increasingly being embedded into procurement and regulatory frameworks, particularly through Request for Proposal (RFP) requirements that now extend to:

- Device-level integrity assurance
- Encryption standards for operational systems
- Auditability of connected infrastructure

2. Enterprise and infrastructure operators

Organisations operating within smart city ecosystems are part of a shared operational risk surface. A compromise in a connected IoT environment can propagate across critical services.

3. CIOs and CISOs

Security leadership must increasingly unify visibility across IT, OT, and physical infrastructure systems to ensure consistent governance and coordinated response.

> **As urban infrastructure scales, unmanaged complexity becomes an expanding attack surface.**

Building Resilient Urban Infrastructure

As cities become more digitally integrated, resilience must be engineered into the architecture of connected systems.

Key priorities include:

- Full asset discovery across IT, OT, and IoT environments.
- Strict segmentation between enterprise and operational networks.
- Secure, identity-controlled remote access to critical systems.
- Incident response models designed for operational disruption scenarios.
- Security requirements embedded into procurement and supplier ecosystems.

However, technical controls alone are insufficient without alignment across stakeholders.

A shared accountability model is essential ensuring that governments, operators, and technology providers have clearly defined ownership, escalation paths, and resilience testing frameworks.

In highly connected environments, fragmented accountability becomes a systemic vulnerability in itself.

Final Perspective

As smart city ecosystems continue to evolve, OT and IoT security is becoming a core pillar of national and urban resilience.

The convergence of digital and physical systems is fundamentally changing the nature of cyber risk, extending its impact from information systems into critical infrastructure and public services.

In this context, resilience is defined not only by the ability to prevent disruption, but by the ability to maintain control, integrity, and continuity across interconnected environments.

Organisations should embed visibility at the protocol level, enforce cryptographic protection across data flows, and ensure integrity of operational systems while maintaining shared accountability across the ecosystem.

In hyper-connected urban environments, cybersecurity is becoming integral to the safe operation of critical infrastructure systems.

From Scalable to Sustainable

The Next Phase of Cybersecurity

By Sudhir Menon, SVP Business Marketing

Cybersecurity in the GCC is entering a new phase of maturity. After years of building scalable, intelligence-driven security through Managed Security Services, organisations are now shifting their focus toward sustainability ensuring security can be maintained effectively across increasingly complex digital environments.

Unsustainable security is often characterised by reactive tool adoption, fragmented architectures, and compliance-driven activity that struggles to keep pace with evolving risk. By contrast, sustainable cybersecurity is built on consistency, integration, and the ability to scale without increasing complexity.

Across the region, organisations are operating in environments defined by cloud-first strategies, AI adoption, and expanding digital ecosystems. This evolution is also reflected in broader industry direction, where cybersecurity is increasingly moving toward platform-led, outcome-driven models with a strong emphasis on resilience and automation.

Cybersecurity is now a core pillar of resilience, enabling organisations to grow securely in a connected world.

Engineering Resilience into Operations

The nature of cyber risk has fundamentally changed. It is no longer defined by isolated incidents, but by continuous exposure across cloud, identity, applications, and third-party ecosystems.

In this environment, resilience is not about adding more controls, but about strengthening how security operates day to day. Leading organisations are embedding detection, response, and recovery into core workflows, supported by real-time telemetry, automation, and Agentic Ai decision making.

The objective is not only to prevent incidents, but to ensure security operations can continuously adapt and maintain business continuity in the face of persistent disruption.

Scaling Securely in a Distributed World

The expansion of multi-cloud, edge environments, and digital platforms has significantly increased architectural complexity. Sustaining security at scale now depends on how effectively organisations design and standardise their security architecture.

Platform-led approaches enable consistent control across environments while reducing fragmentation. This allows organisations to scale securely without adding unnecessary complexity or operational overhead, while also enabling faster integration of new services and partners.

This requires careful consideration of interoperability, data portability, and long-term architectural flexibility as environments continue to evolve.

> **Unsustainable security reacts to risk. Sustainable cybersecurity is designed to operate through it.**

Extending Trust Across the Ecosystem

The region's digital economy is built on interconnected ecosystems, where cloud providers, technology partners, and third-party services are integral to business operations.

As a result, cybersecurity must extend beyond enterprise boundaries to establish ecosystem-wide trust. This requires stronger third-party risk governance, clearer accountability models, and alignment with regional regulatory frameworks such as the National Cybersecurity Authority (NCA), the Saudi Central Bank (SAMA), and the Dubai Electronic Security Center (DESC).

Key risks include compromised third-party access, exposed application programming interfaces (APIs), and supply chain vulnerabilities that can propagate across interconnected environments if not continuously managed.

Organisations that treat ecosystem security as a core capability, rather than a compliance requirement, are better positioned to manage systemic risk.

Investing for Longevity, Not Just Urgency

Many cybersecurity investments are still driven by immediate priorities such as compliance deadlines or incident response. While necessary, these often result in short-lived solutions.

Adaptive cybersecurity requires a shift toward long-term capability building. Investments in automation, AI-assisted operations, and modular security platforms enable organisations to improve both effectiveness and efficiency over time.

This ensures security programmes evolve with the business, rather than requiring continuous reinvention.

The Regional Imperative

In the GCC, sustainability is closely linked to sovereignty. Data residency, regulatory alignment, and national digital strategies are increasingly shaping security architecture design.

Organisations that embed these considerations early are better positioned to scale with confidence, avoiding costly redesigns and compliance gaps.

Regional frameworks, particularly those led by the NCA, are playing an increasingly important role in shaping cybersecurity as both a regulatory requirement and an operational discipline.

A Natural Evolution of Managed Security

The transition to Managed Security Services established a foundation for scalable, outcome-driven cybersecurity. This next phase builds on that foundation by extending it into a continuous, adaptive, and integrated discipline.

The role of the CISO continues to evolve — from managing security operations to orchestrating resilience across the organisation and its ecosystem.

This shift also extends to the broader security function. Long-term effectiveness depends not only on technology, but also on building skilled teams, supported by automation, and structured to reduce operational fatigue while maintaining high performance.

In this context, partners such as Help AG play a critical role in enabling organisations to operationalise sustainable cybersecurity and scale resilience effectively.

Looking Ahead

As digital ecosystems continue to expand, the ability to sustain security will become a defining factor of business success.

Organisations that invest in resilience, scalability, and ecosystem trust will not only manage risk more effectively, but also enable faster and more confident digital growth.

In a connected world, sustainable cybersecurity is no longer optional, it is foundational.

07

POST-QUANTUM

SECURING THE NEXT ERA

Preparing for the Post-Quantum Era

Rethinking Trust Before It Breaks

By Nikola Kukoljac, Vice President Solution Architecture

Encryption has always followed a familiar cycle. It is trusted until a weakness is discovered, then gradually replaced. For decades, this pattern allowed organisations to respond to cryptographic risk after it became visible. Post-quantum computing disrupts that model completely.

For the first time, encryption can remain secure today while being predictably vulnerable tomorrow. That shift fundamentally changes how CISOs must think about risk, investment and timing.

The Asymmetry of Quantum Risk

Traditional cyber risk is immediate and visible. A vulnerability appears, attackers exploit it, and organisations respond. Quantum risk behaves differently.

The most important threat is not the day a large-scale quantum computer becomes operational. The risk already exists because encrypted data can be captured now and decrypted later, the widely recognised "**store now, decrypt later**" scenario.

This creates a profound asymmetry:

- Data theft can happen today without detection.
- The impact may only materialise years later.
- By the time decryption becomes possible, remediation is impossible.

This changes the central security question.

It is no longer *"Is our data secure today?"*

It becomes "*How long must this data remain secure?*"

Data with long-term value, intellectual property, national infrastructure data, health records, financial transactions, and strategic communications is already within the risk window.

> **The UAE Cyber Security Council is turning quantum risk into a structured, time-bound programme—shifting it from theory to execution.**

The Readiness Gap Emerging Across Organisations

While awareness of post-quantum cryptography is increasing, operational readiness has not kept pace. Across sectors, a consistent and material gap is emerging.

Most organisations still cannot confidently answer three fundamental questions:

1. Where is cryptography used across their environment?
2. Which systems rely on public-key cryptography?
3. How replaceable are those cryptographic components in practice?

The challenge is structural. Cryptography is embedded deep within operating systems, applications, identity frameworks, cloud services, network infrastructure, and third-party dependencies often without clear ownership or visibility.

This creates a critical disconnect: visibility initiatives are beginning to emerge, but the ability to execute cryptographic change at scale is not.

The issue is not algorithm selection. It is operational agility, the ability to evolve cryptography without destabilising business systems.

From Algorithm Strength to Cryptographic Agility

Historically, cryptographic strategy focused on strength and longevity. Algorithms were deployed and expected to remain in place for decades.

Post-quantum readiness introduces a new requirement: cryptographic agility, the ability to replace and upgrade cryptographic mechanisms safely, repeatedly and at scale.

This represents a structural shift in security thinking.

Cryptography is moving from a static control to a living infrastructure capability.

Organisations that are early in their journey often discover that cryptography is tightly coupled to systems and vendors. Replacing it can require application changes, hardware updates, and coordination across complex supply chains.

This is why the post-quantum transition is not a single migration project. It is a long-term operational programme that touches architecture, procurement, governance and workforce capability.

What This Means in Practice

Several practical implications are now shaping early post-quantum initiatives:

- Long-term data becomes exposed earlier

  Sensitive data carries risk from the moment it is created if it must remain confidential for years.

- Visibility does not equal readiness

  Discovering where cryptography exists is only the starting point. Readiness depends on the ability to replace it without disruption.

- Dependencies shape the timeline

  Migration speed will depend heavily on vendor roadmaps, platform support and supply-chain alignment.

- Prioritisation becomes harder

  Because the threat is not yet visible, post-quantum initiatives must compete with immediate and active risks for funding and attention.

The organisations making the most progress are those translating cryptographic discovery into structured transition planning, controlled testing and coordinated execution.

Why This Matters Now

The most difficult aspect of post-quantum security is not technology — it is timing.

If organisations wait for the threat to become urgent, it will already be too late for long-lived data. Yet acting too early without structure risks wasted investment and fragmented adoption.

This places CISOs in a familiar but complex position:

making long-term security decisions before the risk becomes visible to the wider organisation.

The post-quantum transition therefore becomes a leadership challenge as much as a technical one. It requires building awareness, securing executive alignment and embedding cryptographic agility into long-term security strategy.

Regulatory Momentum in the Middle East

One of the most significant developments in the global post-quantum landscape is the emergence of structured national guidance. Few jurisdictions have translated the theoretical risk of quantum decryption into formal obligations as clearly as the United Arab Emirates.

Regulatory Framework and Classification

National Security Advisor Decision No. (71) of 2024 concerning the Executive Regulation of the Federal Decree BY Law No. (8) of 2023 on Cryptography. 9.

The National Advisory focuses on categorising UAE entities into Category 1 and Category 2, while clearly outlining the differing cryptographic requirements and methods that should be applied to each category. At the same time, the legislation defines and classifies critical assets into Category A and Category B, effectively providing organisations with practical guidance on where to prioritise their post-quantum migration efforts. In addition, the document identifies technologies and cryptographic algorithms that are strictly prohibited.

The Executive Regulation describes National Cryptography Center (NCC) as an administrative authority established by decision of the National Security Advisor to implement the Cryptography Decree and its Executive Regulation. In the UAE, the NCC is exercised by the Cyber Security Council (CSC).

National Encryption Policy (September 2025)

UAE Cyber Security Council published "National encryption policy" (September 2025) in order to enhance data security, aligned with the UAE's national priority as a global leader in cybersecurity; and enhance the security posture of organisations and individuals within the UAE dealing with critical and personal data. The policy aims to secure data at rest and in motion by providing the mandatory encryption controls that shall be applied by the entities. The policy further defines requirements for secure key management and paves the way for quantum-safe computing.

Post-Quantum Cryptography Requirements (Section 6)

Section 6 of the policy named "Post Quantum Cryptography" is focused on providing insights how entities should prepare to protect their non-open information against future threat posed by quantum computers.

Entities governed by the Cyber Security Council should establish and maintain a comprehensive cryptographic asset inventory for all systems and applications using public key cryptography. It also mandates lab-based testing of new post-quantum cryptographic standards before any production deployment, followed by the development of a structured transition plan. In addition, entities must create post-quantum acquisition policies, including defining new service requirements, assessing vendor readiness, integrating post-quantum capabilities into strategic roadmaps, and identifying foundational technologies needed for adoption. The section further emphasises organisational readiness by requiring entities to notify IT departments and vendors of the upcoming transition, while also educating and training the workforce to support effective implementation.

Compliance, as per the policy, will be monitored through the UAE Information Assurance (IA) Standard controls.

Phased Migration Approach

Based on the information from the market, post-quantum cryptographic migration is likely to be executed in multiple structured phases:

- Initial phase (approximately six months)

Entities governed by the Cyber Security Council would be expected to establish a comprehensive cryptographic asset inventory, identify gaps between current cryptographic practices and regulatory requirements, assess the operational and technical impact of migration, and evaluate associated risks.

- Planning and approval phase

This phase would involve submitting a formal migration plan to the Cyber Security Council for review and approval.

- Execution phase (typically 24 to 48 months)

Once approved, entities would enter a long-term execution phase during which they would transition to post-quantum cryptographic algorithms in accordance with the approved roadmap.

- Validation and certification

Once migration is successfully completed, implementation will be audited and awarded with a compliance certificate.

This approach moves post-quantum security from a future concern to a defined programme of work. It also creates alignment pressure across vendors and partners, extending its impact beyond regulated entities.

What Comes Next

Post-quantum adoption will be gradual and uneven. Regulated sectors will move first, while broader ecosystems follow over time. Hybrid cryptography, combining classical and post-quantum methods will define early deployments.

The defining challenge will not be algorithm selection, it will be execution.

Organisations that succeed will focus on:

- Maintaining visibility of cryptographic dependencies.
- Enabling controlled testing before deployment.
- Aligning internal teams and external vendors.
- Treating cryptography as adaptable infrastructure.

Final Perspective

Post-quantum cryptography is changing the security baseline, and not just the security roadmap.

In the UAE, where the Cyber Security Council is driving a national approach to resilience, the expectation is clear: cryptographic readiness cannot be deferred to a future migration cycle. It must be built into how systems are designed, procured, and operated today.

This is not about replacing encryption at scale in response to a deadline. It is about ensuring that critical systems are not structurally exposed to a known future risk. That requires cryptographic agility as a standard capability, not a specialised project.

Post-quantum risk is already defined. The differentiator now is execution and how quickly organisations align with national direction and embed quantum-safe readiness into their security architecture before exposure becomes irreversible.

INDUSTRY VOICES

INSIGHTS FROM CYBER LEADERS

Cyber Resilience in UAE Banking's Extended Ecosystem

Prashant V Jethwa, Chief Information Security Officer (CISO), Commercial Bank of Dubai

As AI, geopolitics, and supply chain dependencies converge, cyber resilience is shifting from internal control to ecosystem-wide accountability.

Over the past few years, our view of cyber risk has changed fundamentally.

In banking, security was once largely defined by what we owned and controlled. Today, that boundary has expanded far beyond the organisation itself. The services we deliver and the trust customers place in them now depend on a complex ecosystem of cloud platforms, service providers, fintech partners and digital supply chains.

This shift has changed how we think about resilience.

A Convergence of Pressures

We are operating in an environment shaped by several forces at once.

Geopolitical uncertainty is increasingly reflected in cyber activity targeting regional institutions. Artificial intelligence is amplifying the scale and sophistication of social engineering. At the same time, supply chains are becoming a preferred path for attackers seeking indirect access to highly protected environments.

Together, these dynamics are reshaping the risk landscape in a way that goes beyond traditional security boundaries.

It is no longer enough to secure what we directly control. The resilience of partners, providers and platforms has become inseparable from our own.

Moving Beyond Compliance

Regulation continues to play a critical role in shaping our approach, but compliance alone is no longer the destination.

Cyber resilience is evolving into a continuous, threat-informed discipline. Security-by-design is becoming an expectation across the organisation and the wider ecosystem that supports it.

This is especially visible in how the banking sector is approaching artificial intelligence. AI offers clear opportunities to strengthen detection and response, but it also introduces new governance and risk considerations. As adoption accelerates, we are learning to treat AI as both a capability to harness and a technology that demands rigorous oversight.

Cybersecurity as a Business Enabler

One of the most noticeable changes has been how cybersecurity is discussed internally.

Across the banking sector, cybersecurity is now firmly positioned at board level, framed in terms of risk, operational continuity, and customer trust. The conversation has moved beyond technical controls toward the broader impact of disruption and the confidence required to operate and scale in a highly regulated environment.

Investment priorities across the sector increasingly reflect this shift. Alongside the adoption of new capabilities, there is a strong focus on maturing and integrating existing controls, expanding SOC coverage, strengthening partnerships, and ensuring alignment across the security stack.

The goal is a posture that is cohesive, resilient and continuously calibrated to the threats we face.

From Strategy to Continuous Execution

The pace of change in the threat landscape has made one thing clear: resilience must be treated as an operational capability.

Supply chain assurance needs to be approached with the same rigour as internal controls. AI governance must be proactive. Response and recovery must be tested and improved continuously.

The human layer remains central throughout. As AI-driven attacks become more targeted and convincing, awareness must evolve accordingly, particularly for higher-risk users. At the same time, cybersecurity teams must expand their expertise across AI, data security and governance.

Most importantly, resilience at this scale cannot be achieved alone. Collaboration across regulators, industry peers and trusted partners is becoming essential to maintaining visibility and responding effectively to shared threats.

Final Perspective

Cyber resilience in banking is no longer built within organisational boundaries.

It is shaped by how effectively institutions secure, govern and collaborate across the ecosystems that underpin modern financial services.

Quick Takes

Q. What matters most when choosing a cybersecurity partner?
A. Trust, evidenced through execution.

Q. What is the first thought that comes to mind when you think of Zero Trust?
A. Zero Trust is a mindset, not a product.

Q. What is the most underestimated risk today?
A. Supply-chain risk.

Speed, Resilience, and Clarity

Rethinking Cybersecurity for the Agentic AI Era

Sarith Bhavan, Head of Cybersecurity & Technology Platform Operations Mubadala

Cybersecurity entered 2026 with a fundamentally different relationship to time.

For years, security teams operated in environments where threats evolved quickly, but response cycles still allowed room for analysis, mitigation, and recovery. That balance is now shifting. AI adoption, increasingly sophisticated attacks, and geopolitical instability are compressing the timelines cybersecurity teams once relied on.

Cybersecurity has always operated in a catch-up cycle — technology evolves first, and security frameworks and controls follow after. What is different today is the pace at which that gap is narrowing, forcing CISOs and technology leaders to rethink how resilience is sustained across the business.

Within that context, leadership alignment becomes a defining factor. At Mubadala, cyber is firmly positioned as a core pillar of the organisation, reinforced by strong executive backing that ensures security strategy remains ahead of change rather than reacting to it. While not driven by sector-specific mandates, there is active contribution to evolving national information assurance standards, alongside a clear internal commitment to global best practice by design.

What was once treated as preparedness planning is increasingly becoming operational reality.

Cyber resilience is becoming non-negotiable. It is no longer just about protection. It is about sustaining the business through disruption.

> **We've been fortunate to have leadership that recognises cyber risk as business risk. Cyber has always been one of our critical pillars within the organisation and it continues to be so.**

Velocity has Rewritten the Rules

One of the most defining shifts in cybersecurity today is speed.

For decades, the operating model was relatively consistent: detect, analyse, patch, contain. Even when adversaries moved quickly, defenders still retained meaningful time to respond before escalation. That assumption is now eroding.

Agentic AI and increasingly autonomous capabilities are compressing attack cycles dramatically, accelerating reconnaissance, vulnerability discovery, phishing, and social engineering at a scale traditional operating models were never designed for. The nature of attack is also becoming more coordinated, adaptive, and continuous rather than isolated.

This fundamentally changes the operating requirement for security teams. The challenge is no longer simply identifying threats, but ensuring that defence can operate at the same velocity as the environment itself, without losing governance or stability.

This becomes even more complex as AI is embedded across enterprise systems. While it drives efficiency and scale, it also expands the attack surface, introducing new identities, permissions, and governance dependencies that must be continuously managed.

Cybersecurity therefore shifts into a continuous state by default not as a design preference, but as an operational requirement in environments where innovation and exposure evolve in parallel.

From Preparing for Disruption to Operating Through it

In 2025, resilience discussions across the region were still largely centred on preparedness planning for disruption, testing continuity, and anticipating geopolitical volatility.

By 2026, that framing has materially shifted. The question is no longer how to prepare for disruption, but how to operate through it.

In large-scale operating environments, this is already a lived reality. It directly shapes how continuity is designed under conditions where disruption is no longer a scenario, but a variable that must be accounted for in real time. "The question became: how does cyber sustain itself and continue supporting the business without disruption?"

That shift has repositioned resilience as part of the operating fabric itself, not a recovery layer activated after disruption, but a capability embedded into business continuity, operational stability, and executive decision-making.

Within Mubadala, this is reinforced by a tightly integrated security operating model, including a Cyber Fusion Centre, closely aligned with the business. That proximity enables faster, context-aware decision-making where timing and clarity are critical.

The emphasis has therefore moved beyond protection alone. It is now centred on endurance, continuity, and sustaining operations without fragmentation under sustained uncertainty.

Complexity is Quietly Becoming a Structural Constraint

Alongside speed, operational complexity has emerged as one of the most persistent constraints on cybersecurity effectiveness.

Security environments have expanded organically over time in response to evolving threats, resulting in layered architectures that often include overlapping tools and inconsistent visibility. While this expansion was necessary, it has also introduced friction into operational response.

A key but often underestimated factor is underutilisation of existing capabilities. "One of the approaches we have taken is to leverage as much native functionality as possible. Many organisations don't realise what they already have. We spend time optimising and using 100% of the capabilities we have in existing platforms and tools before adding anything new."

This approach materially changes the security equation. By fully activating existing functions, the real gap becomes narrower, more precise, and significantly more manageable.

At the same time, there is a disciplined focus on rationalising tooling ensuring every technology in the environment delivers clear value. More tools do not automatically equate to stronger security. In many cases, they introduce integration overhead, slow response, and reduce clarity at the moment it matters most.

The direction of travel is therefore clear: optimise first, expand only where there is a clearly defined and validated gap.

> **"The modern threat landscape we operate in is persistent, reputational, and increasingly sophisticated. As a highly visible organisation, we continuously encounter targeted attack attempts, including phishing, identity compromise, deepfake impersonation, DDoS activity, attempted data exfiltration, and brand-targeting campaigns."**

Resilience is Becoming the Operating Model

The definition of cyber resilience continues to evolve in parallel.

It is no longer confined to incident response or recovery planning. Instead, it is becoming embedded into how organisations sustain continuity, enable transformation, and maintain trust under conditions of disruption.

This is now reflected across identity governance, AI oversight, operational continuity, modernisation planning, executive accountability, and long-term security sustainability.

Identity remains one of the most critical domains in this evolution. Human identities are now joined by machine identities, autonomous systems, and dynamic access patterns that traditional governance frameworks were never designed to handle at scale.

At the same time, CISOs are increasingly required to balance speed with stability — ensuring that modernisation does not introduce fragility into the environment. The objective is no longer simply faster transformation, but sustainable transformation that strengthens resilience over time.

Final Perspective

Cybersecurity in 2026 is defined not by a single shift, but by the convergence of speed, AI-driven autonomy, and operational complexity.

AI will continue to accelerate both innovation and adversarial capability. Threat activity will continue to increase in scale and sophistication. Geopolitical uncertainty will continue to shape resilience planning and operational decision-making across industries.

Within that environment, cybersecurity is no longer positioned as a supporting function that protects the business from disruption.

It is becoming the condition that determines whether the business can operate at all.

Quick Takes

Q. Biggest cyber risk ahead?
A. AI, agentic AI, and maintaining resilience during geopolitical and operational disruption.

Q. What do organisations most commonly overlook?
A. The native security capabilities and controls they already possess.

Q. What's your immediate thought on AI in cybersecurity?
A. AI is simultaneously advancing both offensive cyber tactics and defensive capabilities at an unprecedented pace.

Defend Forward

The Case for Continuous Cyber Resilience

Mitul Pansuriya, Group IT Manager
Samena Capital Investments Limited

My view of cybersecurity has evolved significantly over the past few years. It is no longer about protecting individual systems, it is about protecting how the business operates, scales and absorbs risk in real time.

In 2026, priorities sit across five key areas: AI security, identity protection, cloud security, ransomware defence and regulatory compliance. What stands out is not the list itself, but how interconnected these domains have become. Decisions in one area now directly impact exposure in another.

As businesses continue to shift into the cloud, identity has effectively become the control plane governing access, trust and accountability across increasingly distributed ecosystems. At the same time, AI is now embedded across workflows, which means security has expanded beyond infrastructure into how decisions are made, automated and governed.

Ransomware has also changed in nature. It is no longer a technical disruption to recover from, but a sustained business continuity threat. In many cases, the question is no longer if operations can be restored, but how quickly and under what constraints they can continue safely.

Our focus areas include cloud security, SOC modernisation and Zero Trust which are now increasingly treated as a single transformation journey. The real shift is toward reducing fragmentation across tools, teams and visibility layers.

Identity security continues to stand out as a critical exposure point. User identities, service accounts and privileged access remain the most exploited entry paths. As third-party ecosystems expand, identity has effectively become the perimeter — and in many cases, the weakest link if not tightly governed.

Regulatory compliance remains a constant, but its role is changing. It is no longer just a control requirement, it is increasingly shaping architecture decisions. The challenge is embedding security early enough so compliance becomes an enabler of scale rather than a constraint on it.

> **The biggest opportunity for cybersecurity leaders in the region right now is leveraging AI and automation to strengthen security — not just as a defensive tool, but as a strategic capability.**

Building a Security Culture That Keeps Pace

A key shift in awareness for me has been recognising that security maturity is shaped as much by people and operating discipline as it is by technology particularly in environments that are highly distributed and cloud-driven.

The skills we prioritise reflect where risk is concentrated and these are not just technical capabilities, they determine how quickly an organisation can adapt under pressure.

Some of the most meaningful gains have come from strengthening fundamentals; expanding multi-factor authentication, tightening endpoint controls and improving security awareness. These remain foundational because attackers continue to exploit consistency gaps rather than new technologies.

Security is also becoming more continuous by design. Awareness programmes and phishing simulations are no longer periodic exercises but now an ongoing operational rhythm, reinforcing behaviour rather than just knowledge.

Even in leadership discussions, I have found that real-world examples are significantly more effective than abstract risk models. Incidents affecting peers or comparable organisations tend to compress decision-making cycles and accelerate investment alignment.

> **Implementing Zero Trust tools does not constitute a complete Zero Trust strategy. The tools create the conditions — governance, behaviour and continuous validation are what make it real.**

From Periodic Defence to Continuous Resilience

If there is one clear takeaway for 2026, it is that reactive security models are no longer sufficient in environments defined by constant change.

The pace of both technological adoption and threat evolution has outgrown traditional review cycles. Security now needs to operate as a continuous function built on resilience, automation and real-time risk visibility rather than periodic assessment.

For organisations operating across multiple jurisdictions and regulatory environments, this shift is structural, not optional. The data we protect, the relationships we maintain and the trust placed in us all depend on sustained operational resilience.

Cybersecurity is no longer a periodic exercise. It has become a continuously operating business capability.

Quick Takes

Q. In one word, what defines your cybersecurity strategy in 2026?
A. Continuous improvement

Q. One area of cybersecurity that deserves more attention today?
A. Identity security

Q. One word that describes the biggest challenge ahead?
A. Agentic AI attacks

Resilience by Design

Securing Healthcare in the Age of AI

Sageer NK, General Manager - IT
Burjeel Holdings

At Burjeel Holdings, we operate a broad network of hospitals and healthcare facilities, serving over seven million patient visits annually across the UAE, Oman, and Saudi Arabia. Every one of those interactions depends on digital infrastructure that must be continuously available, secure, and trusted. That operational reality shapes how we think about cybersecurity — not as a cost centre or a compliance function, but as a direct enabler of patient safety and clinical continuity.

AI is accelerating both attacks and defences, and in doing so, it is increasing the need for stronger governance, clearer oversight, and secure-by-design adoption. The organisations that will lead on security in 2026 are not necessarily those with the most advanced toolsets. They are those that have embedded security earlier: into digital transformation, procurement, AI adoption, cloud architecture, and third-party onboarding.

For security leaders in the region, the biggest challenge is managing accelerating complexity across AI, regulation, supply chains, and workforce capability, all at the same time. But the opportunity is significant: cybersecurity now has a stronger mandate than ever to influence business strategy, resilience investment, and digital innovation.

From Perimeter to Identity: The Structural Shift

One assumption organisations need to challenge in 2026 is that traditional perimeter-led security remains sufficient. The market is shifting decisively toward identity-centric, cloud-aware, and continuously validated security models — and that shift has real consequences for how security programmes are structured and resourced.

In a healthcare environment like ours — where clinical staff, third-party suppliers, biomedical devices, and AI-enabled systems all require access to sensitive patient data — identity is not an abstract concept. It is the most consequential control point we manage. Protecting user identities, privileged access, service accounts, and machine identities is no longer a specialist sub-discipline. It is the foundation upon which everything else rests.

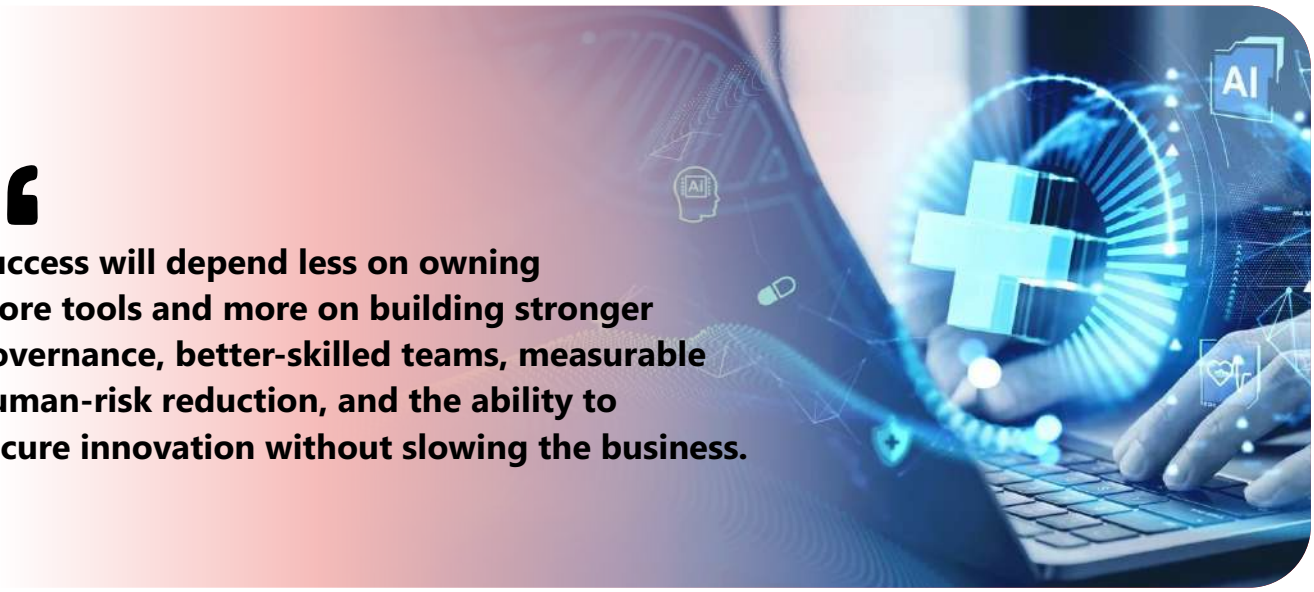

> **Success will depend less on owning more tools and more on building stronger governance, better-skilled teams, measurable human-risk reduction, and the ability to secure innovation without slowing the business.**

AI Governance: The Mandate

Burjeel Holdings is actively investing in AI-enabled care delivery as part of its digital transformation strategy. That makes AI governance a live operational requirement, not a future consideration.

As AI becomes embedded in clinical workflows, diagnostics, and patient management, it is changing how we operate in real time — not just how we plan for the future. Security and governance frameworks therefore need to be built into these systems from the outset, not layered on after deployment.

In practice, this means we are not only asking whether AI is secure, but how it behaves inside live clinical and operational environments — how it is monitored, how decisions are influenced, and how trust is maintained as systems evolve continuously.

The scale of the challenge is not hypothetical. According to the WEF Global Cybersecurity Outlook 2026, 94% of respondents identify AI as the single most significant driver of cybersecurity change this year — and AI-related vulnerabilities were flagged as the fastest-growing cyber risk of 2025 by 87% of security professionals surveyed.

Yet roughly one third of organisations still have no process to assess AI tool security before deployment. The gap between AI adoption and AI governance is widening across every sector, but in healthcare, where AI systems are directly involved in patient care, closing that gap must be the top priority.

The supply chain picture tells a parallel story. In healthcare, third-party access sits at the heart of clinical operations.

Regional data reflects a growing awareness of this: organisations are increasingly prioritising supplier security maturity assessments and involving security teams earlier in procurement decisions. Security is moving upstream, becoming a condition of partnership rather than a periodic review exercise.

Human Capability: The Foundation

47% of organisations in the Middle East and North Africa report lacking the workforce skills required to meet their current cybersecurity objectives. But the skills gap in 2026 is no longer only about headcount — it is about practical capability. Leading organisations are combining vendor certifications, hands-on labs, cross-functional drills, intelligence-led briefings, and continuous learning across cloud, identity, AI governance, and zero trust.

The most in-demand roles reflect where the risk is concentrating: cloud security architects, identity and access specialists, security automation engineers, threat hunters, digital forensics and incident response professionals, and third-party risk specialists. Equally important are practitioners who can govern AI usage safely, a role that barely existed at scale two years ago.

In healthcare specifically, the capability requirement goes further. Security professionals need to understand clinical workflows, biomedical ecosystems, and the operational demands of always-on patient services.

Resilience by Design: The Strategic Imperative

The market direction for 2026 is clear: organisations need to move faster from reactive security improvement to resilience by design. For a group like Burjeel, expanding across the UAE and Saudi Arabia, integrating new facilities, and accelerating digital transformation — resilience is not a retrospective consideration. It is being built into every layer of how we grow.

Measuring cyber resilience in business terms matters here. Operational continuity, recovery readiness, regulatory standing, and trust — these are the metrics that connect security investment to organisational outcomes. For a healthcare provider, they are also the metrics that connect security directly to patient welfare.

Across the market, security awareness is also evolving in the right direction. Moving beyond annual compliance training toward a continuous, behaviour-based model, embedded into daily operations through phishing simulations, role-based learning, executive briefings, and incident-response exercises. The shift is from generic awareness toward measurable human-risk reduction. That shift in culture, as much as any technology investment, is what resilience by design looks like in practice.

> **As AI becomes embedded in clinical workflows, diagnostics, and patient management, the security and oversight frameworks around those systems must be built in from the outset — not bolted on after deployment.**

Outlook

As AI adoption, cloud transformation, and connected healthcare environments continue to evolve, cybersecurity is becoming inseparable from operational resilience and patient trust. Success in 2026 will depend on the ability to secure innovation without slowing it — embedding governance, identity security, workforce capability, and resilience into every stage of digital transformation. In healthcare, cybersecurity is no longer only about protection. It is fundamental to continuity of care, business stability, and patient safety.

Quick Takes

Q. What is becoming a non-negotiable in cybersecurity?
A. Identity-first security — continuous verification of every user, device, and machine identity is now critical in today's AI- and cloud-driven environments.

Q. What comes to mind when you think of AI in cybersecurity?
A. Dual-use — AI is amplifying everything in cybersecurity, from threat detection and response to the scale and precision of attacks.

Q. What matters most when choosing a cybersecurity partner?
A. Trust backed by proven expertise — a partner that combines transparency, deep technical capabilities, and measurable security outcomes.

06

PEOPLE

SUSTAINING TRUST, SKILLS, AND CULTURE

People Behind Resilient Systems

From Hiring to Capability Design

By Tarun Kumar, Director Resident Resourcing

Across the GCC, cybersecurity investment has accelerated significantly as governments and enterprises continue expanding digital infrastructure, cloud adoption, and connected services.

As cybersecurity environments become more complex, the challenge facing organisations is also beginning to shift. The issue is no longer limited to access to security technologies or talent acquisition alone. Increasingly, it is the ability to build teams and operational models that can sustain performance over time.

This challenge is unfolding against a backdrop of rising global demand. Although the cybersecurity workforce grew by approximately **8–9% year-on-year**, the global talent gap still exceeds **four million professionals**[1]. As a result, organisations are placing greater focus not only on hiring talent, but on how expertise is developed, integrated into operations, and retained as part of long-term organisational capability.

That shift is becoming increasingly important across the GCC, where cybersecurity resilience now depends as much on operational maturity and team cohesion as it does on technology investment.

From Talent Acquisition to Organisational Strength

Cybersecurity strategies frequently begin with hiring plans. Roles are defined, gaps are identified, and recruitment follows. While this approach addresses immediate needs, it does not automatically translate into sustained performance.

What differentiates more mature environments is how effectively individuals integrate into the operating model. How quickly they develop judgement. How consistently teams function under pressure. How experience is retained and built upon.

This introduces a new leadership priority. Cyber talent is no longer viewed solely as a resourcing requirement. It becomes part of how the organisation maintains stability, adapts to change, and manages risk in real time.

The shift is subtle but significant: focus moves from filling positions to shaping how those positions perform collectively.

Bridging Learning with Operational Reality

Across the GCC, investment in cybersecurity education and certification has expanded significantly. Access to training is no longer the primary constraint. What matters now is how effectively knowledge is applied.

Cybersecurity environments are defined by pace, ambiguity, and consequence. Decisions are rarely made with complete information, and outcomes often depend on how individuals respond in unfamiliar situations.

As a result, learning is increasingly embedded within operational settings. Teams are exposed to real scenarios, shared problem-solving, and continuous knowledge exchange that reflects actual conditions.

This approach strengthens practical judgement and ensures expertise develops in alignment with the realities teams face, rather than in isolation from them.

Structuring Teams for Sustained Performance

Cybersecurity functions operate under continuous demand. Maintaining clarity and consistency in this environment depends heavily on how teams are organised and supported.

Strong teams are typically characterised by balanced responsibility, clear prioritisation, and leadership that maintains direction without creating unnecessary pressure. Workloads are structured to preserve focus, enabling individuals to

1. ISC2 Cybersecurity Workforce Study

operate with precision over extended periods.

This is not about lowering expectations. It is about designing teams so that high standards can be sustained even as demands fluctuate.

Over time, this stability becomes a defining factor in how effectively organisations respond to evolving threats.

Nationalisation as a Strategic Lever

Across the GCC, national workforce initiatives are increasingly influencing how cybersecurity capability is developed. As organisations focus on resilience and operational maturity, localisation is becoming more closely tied to workforce continuity and skills development across the regional ecosystem.

Within cybersecurity, expertise is built through operational experience, mentorship, and continuous exposure to evolving threats. Nationalisation programmes support this by creating structured pathways for talent development and long-term career progression within the sector.

Across Help AG, this approach includes graduate programmes, university partnerships, career engagement initiatives, and structured early-career pathways into cybersecurity. In the UAE, more than 15% of the workforce are UAE nationals, while in Saudi Arabia Saudi nationals represent more than 65% of the workforce.

As cybersecurity environments continue to evolve, sustaining capability becomes increasingly important. In this context, nationalisation contributes not only to workforce development goals, but to stronger organisational resilience across the region.

> **“As cybersecurity becomes more central to national transformation, organisational strength increasingly depends on the ability to sustain capability, not just acquire it.”**

Aligning Talent with Strategic Intent

As cybersecurity becomes more closely linked to board-level priorities, alignment between strategy and execution becomes increasingly important.

Cyber teams operate at a point where decisions have immediate and far-reaching impact. When closely connected to organisational priorities, their role extends beyond execution to informed contribution.

This alignment is reinforced through visibility. Clear communication of strategic direction, opportunities for engagement with leadership, and recognition of the function’s role all contribute to stronger cohesion.

In a region where cybersecurity is directly linked to national transformation agendas, this connection carries additional weight. It reinforces purpose, sharpens decision-making, and supports long-term retention.

Final Perspective

Cybersecurity across the GCC is entering a more advanced phase. Investment is established, regulatory direction is clear, and expectations continue to rise.

The distinguishing factor now lies in how organisations convert this momentum into sustained performance.

Organisations will not be defined by how quickly they hire, but by how effectively they build environments where expertise develops, teams remain stable, and knowledge accumulates over time.

from short-term fulfilment to long-term structure,

from isolated roles to integrated teams,

from access to talent to the ability to sustain it.

In a region advancing at speed, this is what will ultimately shape resilience and define leadership.

Cybersecurity 2026: AI, Cloud, Sovereignty, and the New Security Operating Reality

Cybersecurity in 2026 is no longer defined by isolated domains of change, but by the convergence of multiple structural shifts occurring simultaneously. Across cloud adoption, sovereign regulation, Agentic Ai acceleration, post-quantum readiness, SOC modernisation, identity-centric security models, and the evolution of attack surfaces such as DDoS and automated threat systems, a single pattern is becoming clear: security is being reshaped as a continuously interconnected operating environment.

What were once treated as separate challenges, infrastructure protection, threat detection, incident response, compliance, and emerging cryptographic readiness are now part of a unified operational reality. These domains no longer evolve independently; they influence and reinforce one another, creating environments where complexity is not episodic, but persistent and structural.

In this context, the definition of cybersecurity maturity is shifting. It is no longer measured by the depth of individual capabilities or the scale of deployed tooling, but by the degree of cohesion achieved across them. The primary determinant of resilience in an environment defined by constant change.

At the same time, the operating landscape itself is accelerating. AI is increasing both the speed and autonomy of systems. Sovereign requirements are redefining how environments are structured and governed. Cloud architectures are expanding across multiple providers by necessity rather than choice. And post-quantum cryptography is introducing a long-term but unavoidable transformation in how trust and security will be established.

Together, these forces are not simply adding complexity, they are redefining the assumptions on which modern cybersecurity has traditionally been built.

Within this shift, cybersecurity is becoming less about perimeter definition and more about systemic coherence. It is no longer a function positioned around the enterprise, but an embedded capability that shapes how systems are designed, connected, and scaled.

For organisations across the region and globally, the challenge is therefore not only to keep pace with change, but to maintain coherence within it. Growth without structure leads to fragmentation. Structure without adaptability limits progress. The balance between the two is where modern cybersecurity now operates.

This year has also reinforced a broader reality shaped by geopolitical volatility and global uncertainty: change is no longer linear or predictable. Events can shift operating conditions rapidly across technology, regulation, supply chains, and threat landscapes. In this environment, preparedness remains essential but awareness becomes equally critical. The ability to sense, interpret, and respond to change in real time is now as important as the ability to defend against it.

As this landscape continues to evolve, cybersecurity will remain in a state of expansion rather than stability. New technologies will emerge, threat models will evolve, and regulatory expectations will continue to intensify.

Within this environment, Help AG will continue to play a central role in enabling organisations across the Middle East and beyond to navigate complexity, strengthen cohesion across their security operations, and build resilience across cloud, identity, SOC, and emerging threat domains. The focus extends beyond preparedness, towards enabling greater awareness, clarity, and control in environments defined by constant change.

That is the direction of cybersecurity in 2026 and the foundation on which its next phase of maturity will be built.

About Help AG

Help AG, the cybersecurity arm of e&, is a leading cybersecurity provider in the Middle East, delivering advanced security services to enterprise and government clients across the region. The company enables organisations to strengthen cyber resilience and securely accelerate digital transformation in an increasingly complex threat landscape.

As part of e&'s AI, Networks & Solutions (AI-N&S) portfolio, Help AG is positioned at the core of the Group's cybersecurity and digital resilience strategy, focused on delivering AI-driven security, sovereign cyber capabilities, and intelligent infrastructure that support secure digital growth.

Established in 2004 and operating in the Middle East ever since, Help AG became part of e& (formerly Etisalat Group) in 2020, forming a leading regional cybersecurity and digital transformation powerhouse.

CYBERSECURITY SERVICES

- Managed Security Services
- Agentic AI Security
- Offensive Security
- Security Posture & Assessments
- Cyber Trust & Advisory
- Cyber Integration & Support Services

E2E SOLUTIONS

- Post-Quantum Cryptography
- Intelligent SOC Automation
- SASE (Secure Access Service Edge)
- Cloud & Application Security
- OT & IoT Cybersecurity
- Identity & Access Security
- Data Protection & Privacy

SOVEREIGN AI PLATFORMS

HYPERSCALERS

CUSTOMER SUCCESS & SERVICE EXCELLENCE

HELP AG
an e& company

Want to learn more?

Visit www.helpag.com

HELP AG
an e& company