ATHENIAN TECH

NEWSLETTER

March 2025

THREAT INTELLIGENCE BY ATHENA

From the Editor's Desk

In this digital age, one thing which is a constant is the pace at which technology is evolving. The frantic pace at which the realm of technology is witnessing the churn could be frightening to say the least for most. But that's life and whether one may like it not, one has to live in these times and not only live but be in position so that one in empowered enough to stay away from cyber harms, be it at the workplace or at a more personal level.

IN THIS NEWSLETTER

Issue #01

One of the biggest impediments on this road to becoming cyber safe, both at the workplace and at a personal level, especially in an age where technology is and will continue to advance at a pace few can fully grasp, staying informed is and will surely be quite a task. From the rise of artificial intelligence (AI), machine learning (ML) and quantum computing (QC) and their employment for achieving nefarious objectives has profoundly impacted the manner in which cyber criminals operate, thereby greatly transforming the digital threat landscape. In the present-day cyber threats/attacks, thanks to inter-connected world that we live in not only impacts business but also governments and even touches the lives of the common man, which basically means that the cyberattacks about which we have been talking have the potential to impact you, and impact you bad. The only way one can stay ahead of these cyber criminals who employ sophisticated tools to achieve their nefarious objectives is to be aware of how they operate and more importantly whom they target or have targeted. This initiative, in the form of a newsletter, is a small step towards keeping you, the reader, better informed so that you are empowered enough not only to become cyber reliant yourself but are competent and confident enough to spread the word about this important aspect of modern day life. Remember cyber threats and attacks are for real and they have the ability to impact your lives in ways that you can, not even imagine. Most importantly thanks to the digital world that we live in you never know when and where you might be encountered with an attack and its resultant effects.

Against this backdrop, we are reaching out to you with this maiden issue, of this newsletter which is a small but deliberate effort, to empower you, so that you are geared to make sense of a world that is rapidly being redefined by technology. Our goal is not to add to the noise but to bring clarity where confusion often reigns by offering you timely insights on cybersecurity, emerging technologies, and the risks and opportunities that lie in between.

For long, cybersecurity was treated as a technical problem, best left to specialists and engineers. But as the headlines of recent times have shown, us that ransomware attacks have the power to impact even the biggest of corporations and even the best of minds are duped by AI-generated scams, in a jiffy, only means that cybersecurity, in the present day, is not only important for governments and businesses to function properly but also for individuals so that they do not loose their mental sanity over issues over which they have no control on. Today, the line separating corporate boardrooms and individual homes from digital threats has been clearly blurred beyond recognition.

It also needs to be remembered that the kind of churn the technology domain is witnessing will only intensify, going forward, making them more sophisticated and in short more deadly and impactful. The only to counter these sophisticated attacks is by employing the same tools and using them for the greater good of mankind, to fight the nefarious elements. Thus AI is no longer a distant frontier or for that matter even ML: they are now inter-woven into everyday tools that we use, reshaping how we make decisions and how we run businesses. Still in its infancy, Quantum Computing promises breakthroughs that could either unlock untold and unheard of efficiencies or upend the foundations of digital security.

In maiden issue of this newsletter, our mission is simple yet urgent:

- To keep you informed on critical developments in the realm of cybersecurity and emerging technology before they become crises.
- To explain complex issues in clear, meaningful, and actionable ways without oversimplifying their importance.
- And to provide insights that help you stay ahead—whether running a company, leading a team, or simply protecting yourself and your loved ones in an increasingly digital world.

Understanding the digital world is no longer optional. The cost of being unprepared grows higher each day, not just in financial terms but in trust, reputation, and the resilience of institutions and individuals alike.

Each month, we will bring a curated selection of what matters most: a blend of analysis, real-world case studies, and practical guidance - so you can focus on what requires your attention and you steer clear of what does not.

As we begin this journey, we hope to create not just a newsletter but a steady companion, especially for those who wish to understand the digital age for what it is—complex, fast-changing, but navigable with the right insights.

Welcome to the first issue. We look forward to walking this road with you. We will ne more than happy to incorporate in this newsletter that you may have. So feel free to share your questions and views on the matter at pankaj.toppo@atheniantech.com .Till we meet again, happy reading.

Pankaj Anup Toppo

Tata Technologies Faces Test of Digital Resilience

On 5 March 2025, the ransomware group *Hunters International* publicly claimed responsibility for a cyberattack on Tata Technologies, one of India's foremost engineering and digital services firms. The attackers alleged that they had exfiltrated approximately 1.4 terabytes of data spread across over 730,000 files. They issued a six-day ultimatum, warning that they would leak the data if Tata Technologies failed to comply with their demand for ransom. Post giving this ultimatum, they released the compromised data on 11 March 2025, on their Onion site.

The nature of the leaked files suggests a breach that runs deep across business layers. The cache includes engineering blueprints, proprietary design documents, internal financial records, vendor contracts, insurance data, legal correspondence, and sensitive HR material, including personally identifiable information about employees' past and present. The breadth of file types, including spreadsheets, presentations, and official PDFs, indicates access not just to isolated systems but also to integrated operational functions.

This public disclosure has prompted renewed attention to a statement issued by Tata Technologies in early January this year. The company reported the detection of unauthorized activity and initiated a precautionary suspension of specific IT systems. At the time, the company maintained that client delivery was unaffected. With the events of March now in view, many cybersecurity analysts believe that the two episodes are linked, signalling that what was described as a precaution in January, this year, may have been the initial phase of this long-running compromise.

The group behind the attack, *Hunters International*, is widely believed to operate on inherited infrastructure from the dismantled *Hive ransomware* group, which had previously targeted *Tata Power* in 2022. Forensic parallels in code and deployment patterns have lent weight to the view that *Hive's* playbook—and possibly its personnel—are now active under a new identity.

The broader context is difficult to ignore. In 2024, India saw a 55 per cent increase in ransomware attacks, 75 per cent of which targeted the industrial sector. With its blend of intellectual property, global client linkages, and sectoral significance, Tata Technologies sits precisely within the crosshairs of these trends.

Interestingly, the market reaction has remained restrained. Tata Technologies' stock held stable and even registered modest gains following the public claim. While some interpret this as a vote of confidence in the firm's containment efforts, experts caution that the more enduring consequences, regulatory scrutiny, client reassurances, or reputational recalibration, often play out over time.

For India's industrial firms, the breach is less of a singular event and more of a structural prompt. In an era where digital design and operational continuity are increasingly intertwined, cybersecurity can no longer remain an auxiliary concern. It must become intrinsic, a built-in pillar of long-term enterprise value.

DRDO Leak Underscores Fragility in Data Governance

On 10 March 2025, the ransomware group *Babuk Locker 2.0* claimed responsibility for a cyber intrusion targeting India's Defence Research and Development Organisation (DRDO), a premier institution under the Ministry of Defence. The group alleged that it had exfiltrated 20 terabytes of sensitive military data, including classified defence documents and credential logs.

The group released a sample of the compromised data, amounting to 753 MB, on the dark web three days later. The leaked dataset included procurement files, engineering diagrams, technical presentations, and documents referencing VVIP evacuation protocols and strategic defence collaborations. One document, marked "SECRET," pertained to an Indian Air Force (IAF) construction project dated December 2020. Another document revealed technical drawings of the RL Mk-III barrel developed by DRDO's ARDE unit in Pune.

Yet, a more nuanced picture emerged as our team examined the material. Much of the data appeared to originate not from DRDO's core IT systems but from a personal device linked to Mr Puneet Agarwal, a former Joint Secretary in the Ministry of Defence (2019–2021). The inclusion personal records, in the compromised data, like Aadhaar details, financial records, and personal travel documents suggests that the compromise stemmed from an individual endpoint, and not from institutional infrastructure.

The group's behaviour also raised doubts about its credibility. Initially demanding $25,000, Babuk Locker 2.0 quickly lowered the ransom demand to $5,000. Despite claiming that DRDO had already paid $300,000 to suppress the leak, they later attempted to resell the compromised data online, a behaviour more consistent with opportunistic extortion than state-aligned cyber espionage.

Still, the nature of the leaked documents cannot be dismissed. Files referencing India's defence procurement strategy, budget allocations, and bilateral partnerships with countries like the United States of America, Brazil and Finland signal strategic sensitivity. The exposure of VVIP evacuation protocols alone underscores the potential national security implications.

Our analysts view the breach as not being a catastrophic compromise of DRDO's secured networks but as reflecting more profound vulnerabilities in data handling and endpoint security. The fact that defence documents were accessible from a personal system against all established cybersecurity norms raises essential questions about compliance, oversight, and internal risk governance.

In an era where cyber warfare blends with misinformation and psychological operations, India's security ecosystem must recalibrate its posture. As this breach shows, it takes only one unsecured device to make national secrets public, and the consequences, though delayed, are rarely minor.

Alleged Breach at ICMR

On 17 March 2025, the ransomware group *Babuk Locke*r 2.0 claimed that it had broken into the defences of the Indian Council of Medical Research (ICMR), and as a result had access to the Aadhaar-linked personal data of 1.46 billion Indian citizens. If this claim by the ransomware group was indeed genuine, the scale of this exposure would make it the biggest breach of data in the country. However, that was not to be and within hours, technical indicators began to point elsewhere.

The group released a dataset via a public file-sharing platform, which contained eight files: six CSVs and two Excel sheets. Our analysts found extensive duplication, padded entries, and multiple corrupted files on inspection. Cryptographic hashes linked the core dataset, shared by the ransomware group, to a previously disclosed breach at ICMR from September 2024, that is already available on public leak forums. The match was conclusive.

Babuk Locker 2.0's claim follows a familiar pattern. In previous incidents, including the leak at DRDO leak. earlier in March 2025, the group relied on high-volume assertions paired with limited or recycled evidence. In the ICMR case, no new material had surfaced. Instead, Babuk Locker 2.0 suggested release of another "50GB" of data that they had access to post their intrusion, without any defined timeline or technical indicators to support the claim.

The group also admitted via encrypted chats that they deployed no ransomware. Their narrative oscillated between political protest, data resale, and public pressure, none of which aligned with conventional extortion behaviour observed in operationally mature ransomware groups.

While ICMR neither publicly confirmed nor denied the alleged breach, independent observers have not reported any signs of system disruption or incident escalation. In this case, the absence of institutional panic may say more than any official statement could.

Still, the claim's effect has been tangible. It has been circulated widely, been amplified online, and tapped into longstanding concerns over the security of Aadhaar data. That such a questionable disclosure could generate national attention underscores the role perception now plays in shaping public understanding of cyber risks.

For India's public institutions, the lesson may lie less in incident response than in confidence management. Today, breaches don't always require technical compromise; sometimes, the mere appearance of one is enough to erode trust.

In a digital climate where facts, fears, and fabrications increasingly blur, it is not just infrastructure that demands resilience—it is institutional credibility.

A Measured Step Towards Autonomous Security Operations

As cyber threats grow in scale and sophistication, their velocity often exceeds the speed at which human teams can interpret and respond, to them. This ever-widening gap is why we continue to monitor advancements in automation within the cybersecurity stack: not for novelty but for their relevance in the real world, in making it cyber resilient.

A recent collaboration between *CrowdStrike* and *NVIDIA* announced at GTC 2025, San Jose, California, illustrates this shift. The two companies will be working to embed agentic AI systems capable of autonomous action into *CrowdStrike's Falcon platform*, with the stated goal of enabling more responsive, decision-capable detection frameworks.

While artificial intelligence has long played a supporting role in cybersecurity, primarily in threat pattern recognition, agentic AI introduces a more consequential evolution. These systems are designed not only to flag anomalies but to evaluate context, triage threats, and, in some cases, initiate containment, all with minimal human input.

For example, CrowdStrike's *Charlotte AI* has been re-engineered using *NVIDIA's* NIM microservices. Early results show doubling in detection processing speed and a 50 per cent reduction in computing overhead. More significantly, the system's ability to distil raw telemetry into decision-grade signals marks a notable shift from alert generation to alert interpretation.

The collaboration also includes testing *NVIDIA's Llama Nemotron* reasoning models, which are designed to reduce false positives and enhance signal fidelity across dynamic threat environments. If proven effective, such models could help security teams move beyond a reactive posture and towards selective, high-confidence interventions.

We include this development not as an endorsement but as a reflection of where the field is moving—and why it matters. As defenders, we are often tasked with separating incremental hype from structural change. Agentic AI may not yet be the standard, but it is gradually becoming the direction of travel—particularly for organizations looking to scale without overextending their security operations teams.

The implications for cybersecurity providers is clear: the security architecture itself is evolving. What was once built around perimeter rules and log collection is now being refactored for environments where decisions must be made faster, closer to the point of risk, and with far less human delay.

For clients seeking greater resilience in high-velocity threat landscapes, the shift towards autonomy—however gradual—cannot be ignored. Our role is to help you stay ahead of it, separating what's emerging from what's enduring and empowering you to gear your operations accordingly.

What Incremental Progress Signals for Cybersecurity Preparedness?

Quantum Computing (QC) has long been regarded as a field of immense promise. At NVIDIA's Quantum Day, held during GTC 2025, in San Jose, California saw measured but tangible progress, on this front. For security professionals, these developments are not yet transformative. But they are directional—and increasingly, warrant attention.

Among the announcements, D-Wave introduced a quantum blockchain architecture focused on optimizing security and energy use in distributed systems. Leveraging quantum annealing, the system is designed to address some of the scaling inefficiencies that plague conventional blockchain protocols. While this is still early-stage, the same class of quantum optimization has implications for secure authentication protocols and decentralized data environments—both of which underpin modern identity and trust frameworks.

Elsewhere, *Quantinuum*, in collaboration with NVIDIA, showcased its work on integrating QC with AI. This pairing—QC-enhanced AI—is being explored for use in logistics modelling and risk simulation. In time, it may also influence threat forecasting in cybersecurity, particularly where complex attack paths and probabilistic risk modelling are concerned.

Perhaps most relevant, however, is what QC represents for the future of cryptography. The consensus remains: once quantum systems reach scale, today's most common encryption standards—particularly those based on RSA and ECC—will no longer be viable. While that timeline is not immediate, the transition to post-quantum cryptographic methods is no longer theoretical. For regulated industries or entities with long-lived data, early preparation is both strategic and necessary.

Why are we watching this space? Because our role as a cybersecurity company is not just to react to current threats, but also to help clients prepare for structural shifts. QC—whether applied to cryptographic resilience or AI-driven risk interpretation—will gradually shape what secure systems would look like, going forward.

As with all emerging technologies, the value lies not in anticipating headlines, but in discerning inflection points. NVIDIA's Quantum Day didn't signal disruption—but it did reinforce that the perimeter of cybersecurity is moving, slowly and silently, towards a different kind of architecture.

AI Governance Finds Momentum

In March 2025, new legislations and institutional shifts in Japan and Kazakhstan added notable weight to the global conversation around AI governance. While the two nations differ in approach, both responded to the same underlying shift—the need to regulate how algorithmic decisions shape public life, commercial operations and digital trust.

Japan, long seen as a steward of ethical technology policy, approved its first national AI legislation on 28 February. The Bill promotes research and deployment of AI technologies but anchors its intent on principals of transparency, accountability, and human dignity. Rather than imposing enforcement-heavy restrictions, it opts for collaborative guidance—encouraging organizations to align with voluntary frameworks shaped through public-private dialogue.

For cybersecurity professionals, Japan's direction is instructive. As AI tools increasingly power detection engines, triage systems, and response orchestration, questions of interpretability and control move to the forefront. Japan's model suggests that governance will assess outcomes, how decisions are made, and who remains accountable when machines act.

Kazakhstan, meanwhile, has taken a more prescriptive step. In March, this year, its Parliament introduced a draft law banning autonomous systems from making decisions without human oversight. Alongside this, the Government of Kazakhstan announced the establishment of *Alem.AI*—a national AI research center in Astana—and laid the groundwork for a broader national strategy. Together, these moves position Kazakhstan as a digital hub that welcomes AI innovation—but insists on structural guardrails.

The regulatory divergence between the two countries is less of a contradiction than a complement. Both affirm that AI cannot remain a purely technical domain. Legal and ethical frameworks will follow as these technologies shape governance, risk modelling, and even national identity systems.

The implications for cybersecurity teams are straightforward. AI-enabled systems—whether in threat detection, fraud analytics, or content moderation—will face growing demands for explainability and control. Systems will not only need to work but also show how they work and for whom.

In that sense, the shift we are witnessing is not just regulatory. It is architectural. For defenders operating at the intersection of automation and trust, it signals a future where resilience will depend not only on technical strength but also on institutional clarity.

5G Advanced Reshapes the Architecture of Trust

In March 2025, a quiet shift is underway: less visible than a product launch but no less significant. Across sectors, the rollout of 5G Advanced, in tandem with edge computing, is beginning to redraw the boundaries of how and where decisions are made.

The appeal is not difficult to understand. With latency now reduced to near-instantaneous levels and compute power positioned closer to where data originates, systems that once relied on cloud round-trips now operate in real-time. In manufacturing, machines now adjust their processes mid-cycle. In diagnostics, medical images are interpreted at the point of capture. The edge, once auxiliary, is fast becoming central.

But with that shift comes a quiet inversion of risk. What once flowed inward, data, alerts, and decision trees now happen outward, distributed across thousands of devices and locations. That decentralization, while essential to speed, also complicates visibility. In March 2025, researchers publishing in academic and industry journals raised a recurring concern: the edge is growing, but the tools to secure it are still catching up.

For cybersecurity practitioners, the challenge is not abstract. Edge nodes often operate with reduced overhead—minimal logging, limited fallback, and constrained computing. A compromised device at the periphery may not trigger alarms in the core, but it may still move laterally, gather context, or remain undetected. In short, the edge is not just a new performance layer. It’s a new attack surface.

What makes this moment important is not the technology itself but what it quietly replaces. The traditional model, centralized infrastructure, well-defined perimeters, and consolidated response yield something more agile but more diffuse. As data becomes more local, so too must security.

For firms managing complex digital environments, this is not a call to retreat, but to reorient. Defending the edge is not about replicating old controls at new endpoints. It is about redesigning trust closer to where the action begins.

Because when systems begin to think and respond in milliseconds, detection and integrity must move as fast and close.

Over the past year, the chance of a leak coming from a personal weak link has become more prominent than the entire IT infrastructure getting compromised, as was seen in the DRDO data breach. Cyber threats are no longer confined to the edges of enterprise networks—they now live in inboxes, voice calls, and everyday habits.

In isolation, each of the threats below may seem avoidable. But taken together, they outline a pattern: human behaviour, quietly manipulated at scale. The breach is no longer a sudden event. It's a slow compromise enabled not by systems, but by habits.

And that's where defence begins: not with new tools, but with sustained attention. Because in a world of autonomous systems and real-time transactions, trust remains both the asset and the target.

Fraudulent Authority Calls

It often begins with a phone call, seemingly from a law enforcement official or a regulatory agency. The voice is calm, scripted, and precise. You're told that your name has surfaced in a case involving financial irregularities. There's a sense of urgency, but also clarity. All they need is verification. In that moment, the instinct to cooperate overrides the impulse to confirm. No official investigation operates like this, but in that pause, the damage is often already done.

Romance-Based Fraud

Some attacks don't rely on panic, they rely on patience. A message, a friend request, a slowly built connection. Fraudsters invest weeks, sometimes months, in cultivating rapport. The eventual ask: money, personal details, or intimate material doesn't come as a red flag. It arrives as a continuation of trust that has already been earned, in due course. When exploitation is disguised as affection, the breach leaves more than just technical scars.

Messaging-Based Banking Scams

On platforms like WhatsApp or SMS, attackers mimic institutional language with surprising accuracy. You're informed that your bank account is at risk, your PAN card is flagged, or your KYC is incomplete. A link is provided to update the same. The whole game seems and looks very plausible. But these messages don't aim to inform, they aim to rush you. Banks never use such channels for sensitive requests with their customers. But the interface is close enough. And when trust is compressed into seconds, mistakes become measurable.

Payment Requests via QR or UPI

The same speed that enables frictionless payments also enables real-time deception. Fraudsters send altered QR codes, unexpected refund requests, or UPI push prompts disguised as customer support. The transaction feels routine, until it isn't. In systems designed for immediacy, fraud doesn't need time. It just needs silence.

Job Offers Demanding Fees

A high-paying remote role. A legitimate-looking offer letter. A small fee for onboarding. The details are convincing but the channels are informal, and the verification almost always breaks down under scrutiny. The goal isn't employment. It's exposure. In a job market flooded with urgency and optimism, exploitation often dresses itself as opportunity.

Phishing and Support Impersonation

Scammers now pose as helpdesk agents from e-commerce platforms, banks, and subscription services. They call or message first, anticipating problems that don't exist. When you respond, they offer to fix it but require remote access or verification codes. Official support teams never initiate contact this way. But when a problem is pre-empted before it's even felt, suspicion can feel misplaced.

Unsecured Public Wi-Fi

Free Wi-Fi in cafes, hotels, or airports offers easy connectivity, and subtle vulnerability. On open networks, attackers can intercept traffic without needing direct access. Sensitive logins, transaction credentials, and session cookies can be captured in transit. A VPN helps—but awareness helps more. Not all threats are visible on screen.

Password Reuse and Weak Credentials

Even strong passwords become vulnerable when reused across services. If one site is compromised, attackers often test the same credentials elsewhere—a tactic known as credential stuffing. Password managers help break that chain. They ensure that a breach in one place doesn't become a breach everywhere.

Multi-Factor Authentication (MFA)

Multi-Factor Authentication (MFA) has moved from it being recommended to it being essential. Yet many still rely solely on SMS verification—susceptible to SIM swaps and interception. App-based authenticators, though slightly more effortful, provide stronger control. In systems where credentials are shared or federated, the second layer is often the one that holds.

Delayed Updates and Unpatched Systems

Security patches often arrive before public awareness of the vulnerability they address. But update prompts are delayed, ignored, or postponed for convenience. Attackers depend on this lag. Systems left even a few weeks behind can become targets. The habit of updating—automatically or proactively—remains one of the most quietly powerful defences.

Unlocked Devices and Poor Endpoint Hygiene

Unlocked laptops. Phones without biometrics. Browsers that remember too much. Physical access remains one of the easiest paths to compromise. Encryption, locking, and local authentication aren't high-friction measures—they're continuity enablers. Devices, like doors, should close behind you.

Manipulated Media and Deepfake Voice Cloning

With tools now able to recreate faces and voices convincingly, attackers no longer need passwords to impersonate. A voice note from a manager. A video call from a known contact. The impersonation may not be perfect, but it only needs to be believable long enough to trigger action. Verification must now extend beyond familiarity. Second-channel confirmation is no longer excessive. It's expected.

Athenian Tech

Contact Us

www.atheniantech.com

https://www.linkedin.com/company/athenian-tech/