VOLUME 2, ISSUE 2 | APRIL 2026

Tech Disruptor media.com

techdisruptormedia.com techdisruptormedia

Pg 40

INTERVIEW
Neehar Pathare
MD, CEO and CIO
63SATS Cybertech

WHEN AI HITS 100KW+ PER RACK, DATA CENTER DESIGN MUST BE REWRITTEN

AMIT AGRAWAL

The Leader Scaling India's Digital Heartland

A BRAND OF BNG
bharatnetworkgroup.com

#BharatCoop26

RUGR SIL PVT.LTD.

Formerly known as Saraswat Infotech Pvt. Ltd.

PRESENTS

POWERED BY

INNOVATION IS GOOD

An Initiative of

BNG BHARAT NETWORK GROUP

Concept by

The Banker media.com

Host Partner

NAFCUB

Bharat Cooperative Banking Summit & Bharat Ratna Sahakarita Samman 2026

Empowering Cooperative Banks for An AI-Driven future

5th - 7th June 2026

Ramada By Wyndham,
Hotel & Convention Center, Lucknow

Scan For More Details

For Sponsorship & Exhibition

Abhinav Chaudhary
+91-8700749849
abhinav@thefoundermedia.in

Naman Singhal
+91-9267933240
naman@thefoundermedia.in

www.bharatcoop.com

Tech Disruptor media.com

Your feedback about this magazine is welcomed at

editor@thefoundermedia.com

info@thefoundermedia.com

This magazine is published under/as a part of "HELLO FOUNDER INFOMEDIA PRIVATE LIMITED, an UTTAR PRADESH-based private limited company registered at the Ministry of Corporate Affairs (MCA). The Corporate Identification Number (CIN) of HELLO FOUNDER INFOMEDIA PRIVATE LIMITED is U56210UP2023PTC191833 and registration number is U56210UP2023PTC191833. HELLO FOUNDER INFOMEDIA PVT LTD's registered office address is Flat No 1006 10th Floor, Tulip 3 Gulmohar Garden, Raj Nagar Extension, Ghaziabad, Uttar Pradesh, India, 201017.

VOLUME 2, ISSUE 2 | APRIL 2026

Table of CONTENT

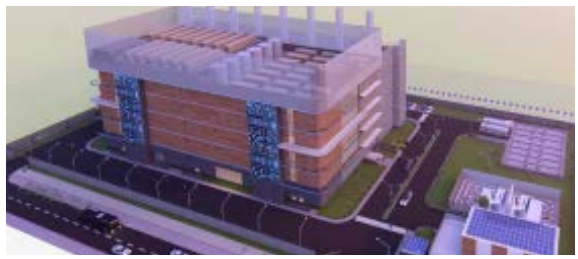

FROM THE FOUNDERS' DESK

Ashish Srivastava (L) and Anupam Gupta (R), Founders, Bharat Network Group (BNG)

RETHINKING THE ENTERPRISE FROM THE CORE

Dear Prime Reader,

Every digital ambition an organisation holds eventually meets the same test: Can its core systems keep up? For years, modernisation sat on the periphery of strategic conversations, treated as a cost to be managed rather than a capability to be built. That thinking is changing.

In this edition of **Tech Disruptor Media**, we examine what it means to modernize core systems for a digital-first future. From cloud strategies and legacy migration to risk management and AI-readiness, the conversation has shifted from whether to modernise to how to do it well.

The contributors featured here have navigated the complexity of legacy migration, cloud strategy, risk management and AI readiness. Their insights offer a clear-eyed view of what modernization demands and what it delivers when done with intention. As the pace of change accelerates, the organisations best placed to lead are not necessarily those with the newest tools but are the ones that have built the foundations capable of carrying whatever comes next. ■

BHARAT™ NETWORK GROUP

UPCOMING EVENTS 2026

#BharatCoop26

RUGR SiL PVT.LTD.
Formerly known as Saraswat Infotech Pvt. Ltd.

PRESENTS

POWERED BY

INNOVATION IS GOOD

Bharat Cooperative Banking Summit & Bharat Ratna Sahakarita Samman 2026

300+ Delegates

5th-7th June 2026
Ramada By Wyndham, Hotel & Convention Center, Lucknow

2nd Chapter

CIO HORIZON

Where Vision Becomes Direction

2026

100+ Delegates

3rd-5th July 2026
Ramada By Wyndham, Hotel & Convention Center, Lucknow

LEGACY IS A CHOICE UNTIL IT BECOMES A CRISIS

Dear Reader,

Enterprise technology has reached an inflection point, one where the gap between organisations that have modernised their core systems and those that have not is no longer a matter of preference. It is a matter of survival.

Every edition of **Tech Disruptor Media** is built around a conversation worth having. This one felt more urgent than most. As we began reaching out to technology leaders for this April edition, what struck us immediately was how much the conversation had matured. A year ago, many of the questions around cloud strategy, legacy modernisation and AI readiness were still being debated in principle.

Today, the leaders we spoke to are past the debate. They are in the middle of the execution, and they had a great deal to say about what that looks like closely. The CIOs and CTOs featured in this edition spoke openly about the complexity of running parallel environments while rebuilding core infrastructure underneath a live business. Technology heads shared their perspective on what hybrid and multi-cloud strategies genuinely deliver once the theory meets the operational reality. This edition is a reflection of where enterprise modernisation stands in April 2026, told by the people closest to it.

Aishwarya Saxena
Sr. Associate Editor
editor@thefoundermedia.com

> **The leaders shaping enterprise technology today are not waiting for consensus. They are building conviction one decision at a time**

COVER STORY

MODERNIZING CORE SYSTEMS FOR A DIGITAL-FIRST FUTURE

Moving from old systems to modern platforms is harder than most roadmaps suggest. **Aishwarya Saxena** speaks with **Amit Patil, MD & Founder, CynalitX Consulting LLP; Suresh Anantpurkar, Founder & CEO, Manch Technologies; and Divyanshu Bhushan, Technology and Business Head – India & SEA, TO THE NEW**, to get an honest picture of the risks, the opportunities and the mindset shift that make modernization work

The hidden costs of standing still: A case for purposeful modernization

Technology modernization is rarely as clean as it looks on a roadmap. **Amit Patil, MD & Founder, CynalitX Consulting LLP** has thought deeply about what actually makes these programmes succeed or fail, from the weight of decades-old systems and the people who depend on them, to the cloud strategies and risk disciplines that determine whether a transformation holds together under pressure.

His views are direct and grounded in what he has seen work and what he has seen quietly derail otherwise well-planned efforts.

Technical debt is only half the problem

The most significant challenge, in Patil's view, is the sheer complexity that accumulates over decades. Legacy systems were not built with modularity in mind. They are deeply interdependent, often undocumented, and critical to daily operations. You cannot simply switch them off. Organizations end up running parallel environments, which drives up both cost

> AI-assisted code migration tools, automated documentation generation, and intelligent testing frameworks are already reducing the manual effort of legacy transformation. Over the next three to five years, this will fundamentally change the economics of modernization

Amit Patil
MD & Founder
CynalitX Consulting LLP

and cognitive load at the same time.

But he is equally clear that the human side is just as hard. Teams that have worked on stable, familiar systems for years often see modernization as a threat to their expertise rather than a chance to grow. That makes change management as important as any architecture decision. Without real investment in upskilling and clear communication, even the best technical strategy will stall.

Data migration is where many organizations underestimate the work involved. Decades of operational data carries inconsistencies, deprecated schemas and compliance obligations that do not map cleanly onto modern platforms. Patil's point is firm: getting data right is not a secondary concern. It is the foundation everything else is built on. Organizations that treat it as a final step almost always find it should have been the first.

Remove the friction and teams move faster

Legacy architecture creates invisible friction at every stage of development. Engineers spend too much time maintaining fragile integrations rather than building new things. When monolithic systems are replaced with modular, API-first platforms, that resistance largely disappears. Teams can develop, test and deploy features independently, cutting out the inter-team dependencies that historically stretched delivery timelines from weeks into months.

Cloud-native platforms remove provisioning bottlenecks that used to require procurement cycles and physical setup. What once took weeks now takes minutes. Patil uses a practical example to make the point: a product experiment that previously needed a business case and a full budget cycle can now be prototyped and validated within a single sprint.

There is also a compounding effect that is easy to overlook. Modern platforms attract modern engineering talent. Developers working with containerization, CI/CD pipelines and observability tools tend to move faster and produce higher quality output. Over time, that talent advantage becomes a durable competitive edge.

Multi-cloud is a strategy, not a fallback

Patil sees hybrid and multi-cloud

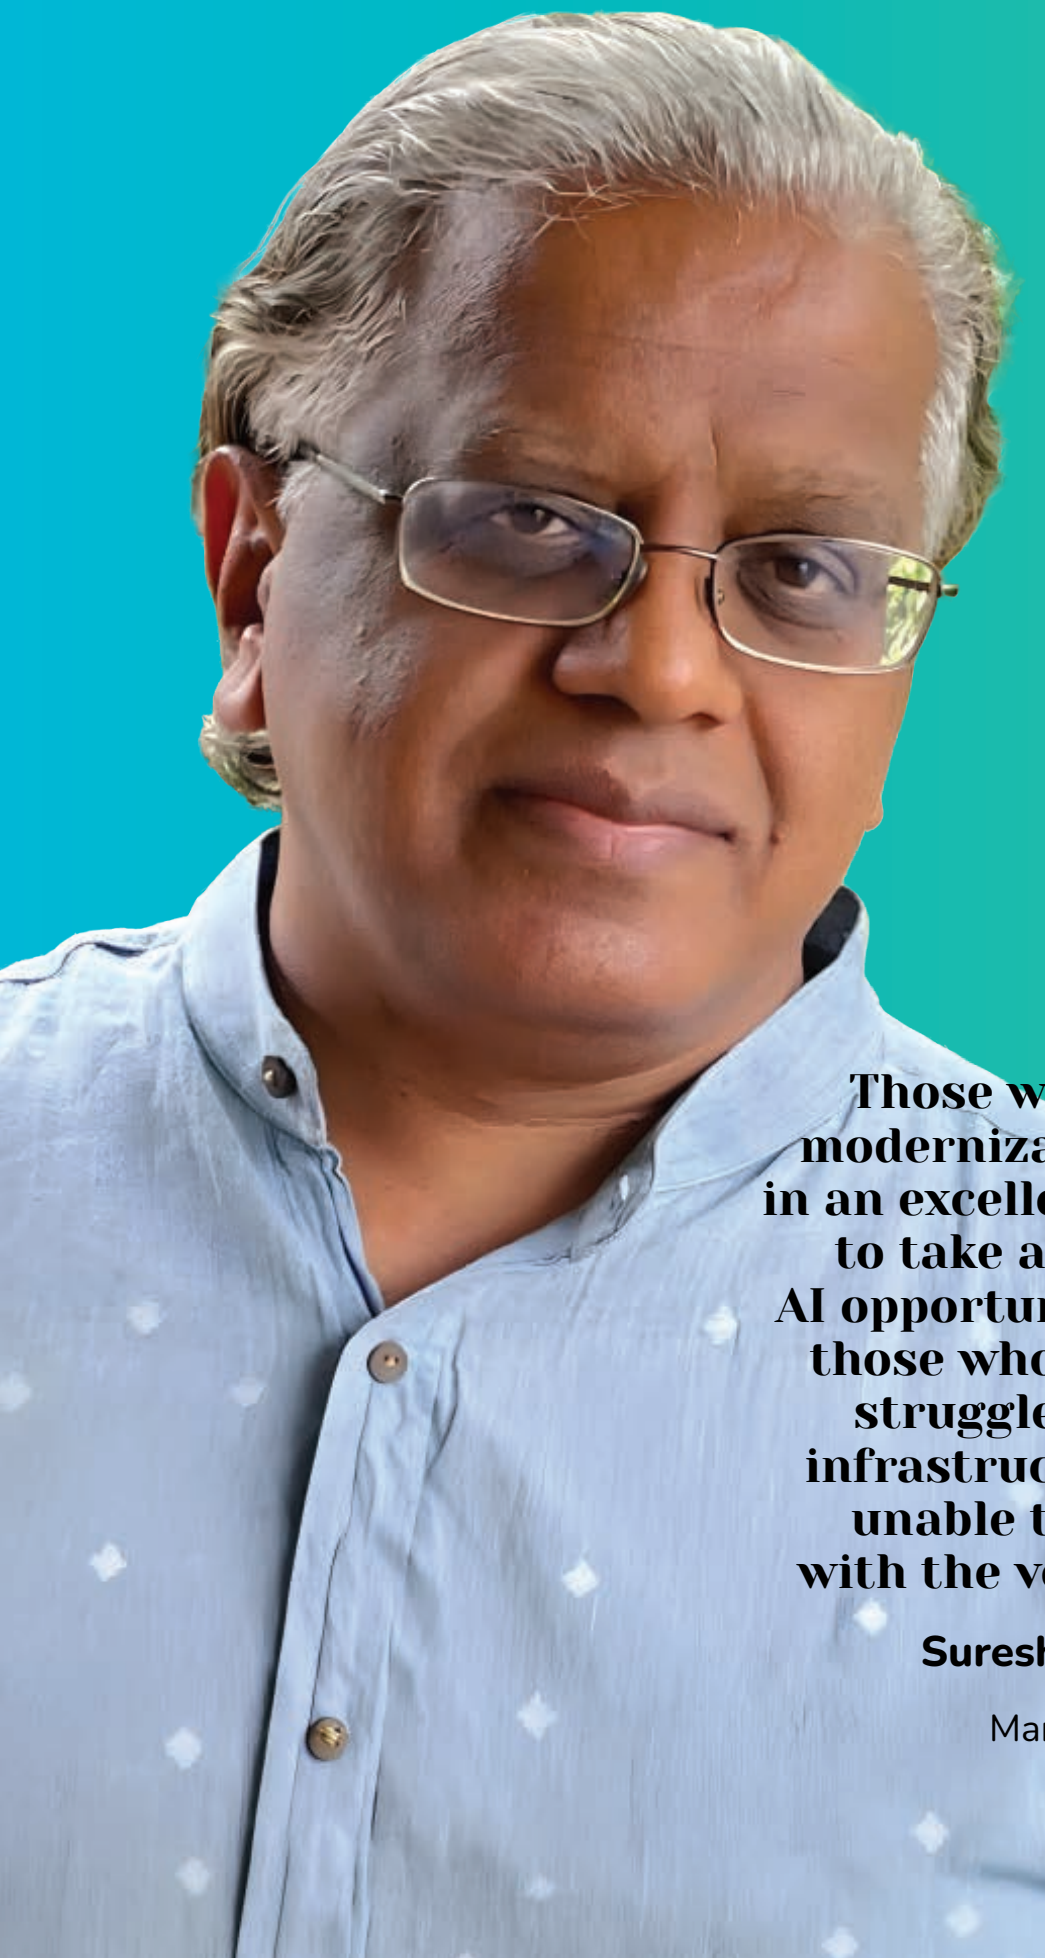

Those who invest in modernization will be in an excellent position to take advantage of AI opportunities, while those who do not will struggle with aging infrastructure that is unable to keep pace with the velocity of AI

Suresh Anantpurkar
Founder & CEO
Manch Technologies

strategies as having moved well past contingency planning. They are now a deliberate architectural choice. No single cloud provider is uniformly strong across every workload. By distributing across providers, organizations can match specific requirements to whoever is best placed to deliver, whether that is machine learning infrastructure, data sovereignty, or low-latency edge computing.

From a scalability standpoint, multi-cloud removes dependence on any single provider's capacity limits or pricing model. Organizations can scale resources precisely where demand arises without overprovisioning. This is especially useful for businesses that deal with seasonal spikes or rapid geographic expansion.

The flexibility argument also holds from a risk angle. Vendor lock-in tends to be a hidden cost that only surfaces at renewal time or when a provider changes its terms. A hybrid architecture keeps negotiating leverage intact and ensures business continuity if any single environment runs into problems.

Risk is a continuous signal, not a periodic audit

Patil is clear that risk management in a large-scale transformation is an operational discipline, not a governance checkbox. The organizations that come through these programmes successfully are the ones that treat risk as something to monitor continuously, not review occasionally. He favours incremental delivery over big-bang cutovers. Progressive migration, running legacy and modern systems in parallel and validating at each checkpoint, limits the damage when something goes wrong. Rollback planning needs to be in place from day one, not added later as an afterthought.

Cybersecurity deserves particular attention during transitions. When old and new environments are running alongside each other, the attack surface expands. Security architecture needs to keep pace with infrastructure changes, not lag behind. Regulatory and data compliance obligations do not pause during a migration either. They need to be mapped, validated and evidenced throughout the entire process.

The next phase of modernization will not just be about upgrading systems; it will be about building organizations that can continuously adapt and evolve

Divyanshu Bhushan
Technology and Business Head for India & SEA
TO THE NEW

AI, sustainability and the expanding edge

Looking at the next three to five years, Patil expects AI to move from being a feature layered onto modernized systems to being embedded in the modernization process itself. AI-assisted code migration, automated documentation and intelligent testing frameworks are already cutting the manual effort involved in legacy transformation. Over time, this will change the economics of modernization and make it viable for organizations that previously could not justify the investment.

Platform engineering will also mature. Instead of each product team assembling its own tools, organizations will invest in internal developer platforms that offer standardized, self-service infrastructure. This creates consistency, speeds up onboarding and frees engineering teams to focus on business problems rather than infrastructure configuration.

Sustainability is another factor that Patil expects to move from reputational concern to structural constraint. Regulatory pressure around energy consumption and carbon reporting will start shaping cloud architecture decisions directly. Organizations that do not account for infrastructure efficiency will face both compliance risk and cost exposure as energy pricing and carbon obligations tighten.

> **Newer approaches, particularly API-first and microservices-based systems paired with DevOps practices, let teams work independently and ship continuously. Organizations that have made this shift have seen time-to-market improve by as much as 50%.**

From legacy to cloud-native: Why modernization can no longer wait

The debate around legacy modernization has been going on for years, but **Suresh Anantpurkar, Founder and CEO of Manch Technologies**, thinks the window for hesitation is closing fast. As AI reshapes what enterprise technology needs to do, the gap between organizations running modern systems and those still holding on to old infrastructure is only going to widen. His perspective across infrastructure, cloud strategy, risk and the road ahead reflects someone who has thought carefully about where this is all heading.

Old systems, new problems

The core difficulty with moving from on-premise systems to modern digital platforms, according to Anantpurkar, is that the two were built in entirely different ways. Legacy systems were developed with older technologies that do not validate data well. Cloud-based systems, by contrast, are designed with interoperability in mind from the ground up.

Making that jump is not straightforward. Data migration, security, access controls and network policy alignment all need careful thought before any move can happen.

But Anantpurkar is firm that the transition cannot be avoided. Staying on old systems is not a safe choice either. It limits scalability and, increasingly, it locks organizations out of the AI capabilities that are becoming central to how businesses operate.

Agility is what modern systems are built for

Anantpurkar points to agility as the defining

advantage of modern core systems, something legacy infrastructure simply cannot offer. These systems are designed for the cloud, driven by APIs, and capable of handling large and complex data across multiple devices and channels.

Better data validation and workflow flexibility mean development cycles get shorter and teams can respond to changing market needs without being held back by the system underneath them.

He also highlights AI-readiness as one of the most important features of modern systems. As AI evolves rapidly, having infrastructure that can support it across the enterprise is not a nice-to-have. It is what separates organizations that can operationalize AI from those that cannot.

Scale on demand, without the maintenance burden

Hybrid and multi-cloud models have given Manch Technologies the kind of scalability and flexibility that enterprise-level growth requires. Anantpurkar describes being able to grow or shrink infrastructure as needed, access better services and keep sensitive data secure, all through a cloud-native approach and the hyperscale providers they work with.

What also stands out for his teams is the reduction in infrastructure management and software maintenance. That overhead used to absorb a lot of time and energy. With it largely gone, teams are free to focus on work that actually moves the business forward.

Risk management is what holds a transition together

Any significant technology changeover brings its own set of risks, and moving to cloud-based systems is no different. Data privacy, cloud security, access controls and cyber threats all come into sharper focus during a transition like this.

Anantpurkar sees this as putting IT management and the CISO role at the centre of the process. Their job is to identify risks, document them and put the right policies and structures in place to address them.

His view is simple: without proper risk management, a transition of this scale cannot be done well. It is not a parallel track. It is part of the work itself.

The next few years will separate the ready from the rest

Looking ahead, Anantpurkar expects the next three to five years to be defined by AI-native platforms, growing cloud adoption, new pricing models and a much stronger focus on data, processes and governance.

For organizations that invest in modernization now, the payoff will be the ability to take full advantage of AI as it matures. For those that do not, aging infrastructure will make it increasingly hard to keep up with the pace at which AI is moving.

Modernizing at speed: Building organizations that continuously evolve

Most conversations about digital transformation tend to get stuck on the obstacles. **Divyanshu Bhushan, Technology and Business Head for India & SEA at TO THE NEW**, takes a different view.

For him, the more useful question is not whether modernization is hard, but how you approach it. Across areas like legacy infrastructure, cloud strategy, innovation velocity and risk, his perspective is consistent: the organizations that treat transformation as an ongoing discipline rather than a one-time event are the ones that come out ahead.

Legacy modernization: complexity worth embracing

Bhushan is clear that modernization should not be seen as a burden. It is one of the most valuable things an organization can do, even if it brings real challenges along the way. Technical debt, integration issues, skill gaps and resistance to change are all part of the picture.

The problem with legacy systems is how deeply rooted they are. Replacing them carries risk, and in many organizations, a large share of the IT budget still goes toward just keeping them running. That leaves little room for anything new.

His view is that modernization works

best when it is treated as an ongoing process rather than a one-off project. The goal is to figure out what still adds value, what needs to be rebuilt and what should simply be let go. Organizations that get this right are the ones that make modernization part of how they operate, not something they do once and move on from.

Modern architecture cuts time-to-market in half

Bhushan believes modernization changes the way innovation actually happens inside a company. In older monolithic setups, even small updates take a lot of effort, which means slower releases and less room to experiment.

Newer approaches, particularly API-first and microservices-based systems paired with DevOps practices, let teams work independently and ship continuously. Organizations that have made this shift have seen time-to-market improve by as much as 50%.

But the bigger change, he says, is in mindset. Innovation stops being something that happens in planned cycles and becomes something that happens all the time. Teams stop waiting on infrastructure and start responding to what the business actually needs, when it needs it.

Cloud as a design principle, not a destination

For Bhushan, the conversation around cloud has moved on. The question is no longer which cloud to use, but which workload belongs where and why. That shift in thinking is what makes hybrid and multi-cloud strategies so useful.

Different workloads have different needs. Sensitive data can stay in controlled environments while high-demand, customer-facing applications scale on public cloud. The result is a setup that is more resilient, less dependent on any single vendor and better aligned with what the business actually requires.

It also means organizations can adapt as their needs change, without having to undo decisions made around a single platform or architecture.

Risk must be built in, not bolted on

Risk is unavoidable in large transformation programmes, but Bhushan thinks most organizations approach it the wrong way. Treating it as a final review or a compliance box to tick at the end rarely works. The better approach is to factor it in from the start.

That means weaving governance, security and monitoring into every phase of the work. Phased rollouts, continuous testing and solid observability practices all help keep systems stable while change is happening around them.

He also points to the human side. When leadership is visible, communication is clear and teams feel trusted, execution risk drops considerably. People who understand why something is changing and how they fit into it are much more likely to make it work.

The next frontier: AI, real-time data, and composable systems

Over the next few years, Bhushan expects architecture, data and intelligence to come together in ways that reshape what core systems can do. AI will stop being something added on top and will instead be woven into systems themselves, helping them adapt, optimise and make decisions on the fly. But that only becomes possible when the underlying data and platforms are modern enough to support it.

Platform engineering will also matter more than it does today. How easy it is for developers to build, test and ship will have a direct bearing on how fast a business can move. Alongside this, composable architectures built from modular, interchangeable parts will make it much easier to respond when the market shifts.

And real-time data will become the standard. Businesses will run on live information rather than reports that are already out of date by the time they are read. ■

INTERVIEW

ALANKIT BALANCING DATA ACCESSIBILITY WITH STRICT OWNERSHIP RIGHTS

As India moves towards a user-centric model of data ownership, **Siteshwar Srivastava, Chief Technology Officer, Alankit Ltd.,** speaks to **Aishwarya Saxena** about what that shift means for enterprises managing high-volume citizen and financial transactions at scale

Alankit operates across financial services, e-governance, and healthcare. How do you architect a unified data governance framework across such diverse verticals?

At Alankit, data governance is a core strategic priority that enables seamless operations across our diverse verticals. Given the scale and diversity of our services, we focus on building a governance framework that ensures consistency, compliance, and interoperability, while still allowing flexibility for sector-specific requirements.

Our approach is built around the following key pillars:

Governance Model

- We employ a federated governance architecture, with centralised policymaking and execution devolved to individual business units.
- This model secures consistency of principles whilst allowing operational flexibility across diverse sectors.
- Governance is regarded as a strategic capability, directly connected to business outcomes and digital transformation.

> **Our approach centres on enabling secure and responsible data usage, while ensuring that ownership and user rights are safeguarded at all times**

Standardisation and Enterprise Controls

- We establish organisation-wide standards for data classification (sensitive, critical, public), security protocols, lineage tracking, and compliance requirements.
- Common governance policies are defined to serve as a baseline across all verticals.
- Regulatory alignment is maintained across sectors including BFSI, healthcare, and public services.

Domain-Level Customisation

- Each vertical applies governance through domain-specific frameworks, tailored to regulatory and operational requirements.
- We enable context-aware controls, ensuring relevance without compromising enterprise standards.

Technology and Interoperability

- We invest in shared governance platforms, including data catalogues, quality tools, and access control systems.
- APIs and integration layers are used to ensure seamless data exchange across systems.
- Interoperability and data consistency are promoted across applications and platforms.

Accountability and Organisational Alignment

- Clear roles are defined: data owners (accountable), data stewards (operational), and custodians (technical).
- A central governance council provides oversight and enforces policy.
- Governance is embedded into business workflows, KPIs, and performance metrics.

This discipline is especially critical given our role in managing high-volume citizen and financial transactions across platforms.

With increasing scrutiny around data privacy, how does Alankit balance data accessibility with strict ownership rights?

At Alankit, we view data accessibility and privacy not as competing priorities but as complementary principles that must be designed to work in harmony. Our approach centres on enabling secure and responsible data usage, while ensuring that ownership and user rights are safeguarded at all times. This begins with a privacy-by-design and privacy-by-default philosophy, where data protection is treated as a core architectural principle from the outset, never as an afterthought.

Access is governed through granular controls, including role-based (RBAC) and attribute-based (ABAC) mechanisms, with the principle of least privilege enforced to ensure users access only what is strictly necessary.

To protect sensitive data, we apply end-to-end encryption both at rest and in transit, alongside masking, tokenisation, and anonymisation techniques that safeguard personally identifiable information while still enabling secure data sharing. Our security systems and controls are continuously strengthened in line with evolving regulatory and cybersecurity standards across financial services and e-governance.

Consent management systems are integrated into our platforms, giving users meaningful control over how their data is used, while audit trails and documentation are maintained to reinforce accountability and ensure compliance with evolving data protection regulations.

Underpinning all of this is a commitment to real-time monitoring for unauthorised access or anomalies, regular audits, risk assessments, and compliance reviews, all of which form the foundation of the transparency and trust we work to build with our customers and stakeholders.

As Alankit integrates AI into its services, how do you ensure that data feeding these systems is accurate, unbiased, and well-governed?

As we progressively integrate AI into selected services, we recognise that the

> **Responsible AI principles guide our approach, ensuring transparency, accountability, and regulatory compliance**

quality and governance of data directly shape outcomes. Our priority is to ensure that the data driving these systems is reliable, transparent, and aligned with principles of responsible use. This is achieved through a structured approach built on the following pillars:

Data Quality and Integrity

- End-to-end data quality pipelines are established, covering validation, cleansing, normalisation, and deduplication.
- Data is continuously monitored for accuracy, completeness, and consistency before being fed into AI systems.

Traceability and Lineage

- Full data lineage is maintained, enabling traceability from source to output systems.
- Auditability of data transformations is ensured, enhancing transparency and accountability.

Bias Mitigation and Fairness

- Diverse and representative datasets are used to minimise systemic bias.
- Periodic reviews and testing are conducted to identify and reduce unintended outcomes.
- Datasets and processes are continuously refined to improve fairness and reliability.

Human Oversight and Controls

- Human-in-the-loop frameworks are implemented, particularly for high-impact or sensitive decision-making scenarios.
- Review and escalation mechanisms are established for outputs generated by automated systems.

AI Governance Approach

- Governance is embedded across the AI lifecycle — from data preparation through deployment and monitoring.
- Documentation, version control, and approval workflows are maintained.

Responsible AI principles guide our approach, ensuring transparency, accountability, and regulatory compliance.

How do you integrate legacy systems into modern governance architectures?

Given the scale and history of our operations, legacy systems remain an integral part of our ecosystem. Our approach is to modernise these systems in a structured, non-disruptive manner, while bringing them under a unified governance framework. Rather than pursuing wholesale system replacement, we follow a phased, risk-based transformation approach that prioritises incremental modernisation to avoid operational disruption.

To connect legacy systems with modern platforms, we employ APIs, middleware, and microservices, creating a bridging layer that enables seamless data flow without requiring changes to core legacy infrastructure. Data standardisation is equally central to this process. Legacy data is carefully mapped and transformed into a unified enterprise data model, ensuring consistency, accuracy, and compatibility across systems. Modern governance policies like covering security, access, and data quality, are extended to legacy environments, with these systems required to adhere to current compliance and audit standards.

Where direct integration is limited, we

rely on data synchronisation, replication, and staging techniques to maintain continuity. Looking ahead, systems are continuously evaluated for upgrade, optimisation, or phased decommissioning, always with a careful balance between cost efficiency and the long-term scalability and governance maturity that our operations demand.

What are the biggest challenges Alankit has faced in implementing governance frameworks at scale?

Implementing data governance at scale is as much an organisational challenge as it is a technological one. It demands alignment across teams, clarity of ownership, and a sustained focus on cultural transformation. The key challenges we have addressed include:

Cultural Transformation

- Driving a shift from viewing data as a by-product to recognising it as a strategic asset.
- Encouraging ownership, accountability, and data-driven thinking across teams.

Organisational Alignment

- Aligning multiple stakeholders, departments, and leadership priorities.
- Managing differing levels of maturity and readiness across verticals.

Ownership and Accountability

- Defining and enforcing clear roles and responsibilities.
- Eliminating ambiguity in data ownership and stewardship.

Operational and Technical Complexity

- Integrating diverse systems, including legacy infrastructure.
- Managing large-scale, distributed data environments.

Balancing Governance with Agility

- Ensuring governance frameworks do not impede innovation or business operations.
- Maintaining the right balance between control and flexibility.

Execution and Adoption Strategy

- Addressing challenges through strong leadership commitment, phased implementation, and measurable KPIs.
- Investing in continuous training, awareness programmes, and governance tools.
- Driving ongoing monitoring and improvement cycles.

How do you see the future of data ownership evolving in India over the next decade?

India's digital ecosystem is evolving rapidly, and we anticipate a clear shift towards more user-centric models of data ownership over the next decade. Individuals will gain greater control over their personal data through consent-based frameworks, driven by rising awareness and growing demand for transparency and data rights. This will be supported by the expansion of secure, interoperable digital public infrastructure and broader adoption of data-sharing frameworks across sectors.

At the same time, the emergence of decentralised identity systems and verifiable credentials will gradually reduce reliance on centralised data storage models, fundamentally changing how identity and data are managed.

For organisations, this evolution will mean transitioning from a position of data ownership to one of custodianship and stewardship, with heightened accountability for ethical data usage and protection.

Strong governance will increasingly become a key differentiator for brand trust and reputation, and regulatory compliance will evolve from a baseline requirement into a genuine strategic advantage. Ultimately, the organisations that will lead in this new landscape are those that prioritise user empowerment, transparency, and secure data exchange, because long-term success will depend not just on technological capability, but on responsible innovation, governance maturity, and the ability to sustain trust over time. ■

editor@thefoundermedia.com

INTERVIEW

HOW BLUE CLOUD SOFTECH SOLUTIONS IS CLOSING THE GAP BETWEEN AI INNOVATION AND ENTERPRISE-GRADE SECURITY

In a candid conversation with **Aishwarya Saxena, Vinod Babu Bollikonda, Group CEO, Blue Cloud Softech Solutions Ltd.** talks about how Blue Cloud Softech is building the frameworks, architectures, and operating models that will define what responsible and resilient enterprise technology looks like in an always-connected world

With your strong focus on AI, cloud, and cybersecurity integration, how do you ensure that security is embedded at the ideation stage rather than retrofitted later?

Security has to start as a design principle, not a deployment activity. The mistake many organisations still make is treating security as a final review before go-live. By then, the architecture, workflows, and data flows are already fixed, which makes protection slower, more expensive, and often less effective.

Our focus is on building systems where identity, data protection, monitoring, response, and governance work together as a connected control fabric

Our approach is to bring security, architecture, and business intent into the same conversation from the beginning. That means threat modelling at the concept stage, defining trust boundaries early, building identity, access, and data protection into the architecture, and ensuring that every new use case is assessed for privacy, resilience, and operational risk before it is built. In modern environments, especially those involving AI and cloud, security has to sit inside innovation.

The most resilient systems are those designed with protection built into their core architecture, transforming security from a reactive control into a strategic enabler of scale, speed, and organisational confidence.

When modernizing legacy systems into cloud-native or AI-driven architectures, what are the most critical security gaps organizations overlook?

The biggest gaps usually appear in the transition layer, not the end state. When organisations modernise legacy systems, they often focus on migration speed, application refactoring, or AI enablement, but underestimate what happens to identity, data, interfaces, and monitoring during the shift.

The most common overlooked areas are weak identity and entitlement design, inconsistent data classification, exposed APIs, poor secrets management, and a lack of visibility across hybrid environments. Legacy systems often carried their own implicit controls, but once workloads move into cloud-native or AI-enabled architectures, those controls must be redesigned, not assumed.

Another blind spot is operational fragility. Many teams build for functionality but not for containment, recovery, or forensic readiness. In a modern environment, security must account for how fast a threat can move, how quickly it can be detected, and whether the organisation can isolate impact without disrupting the business.

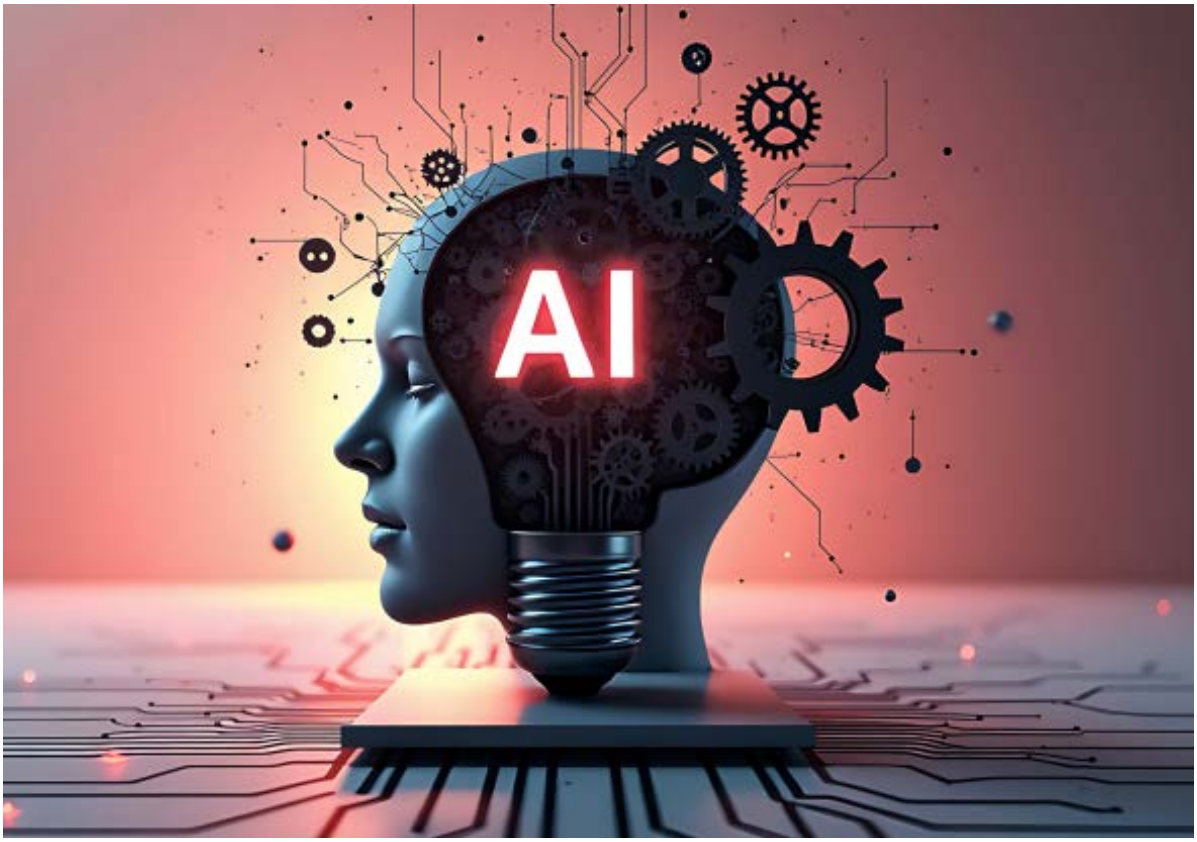

Modernisation ultimately becomes a redesign of enterprise trust, where cloud transformation, AI adoption, and cybersecurity maturity must advance together as a single integrated strategy rather than separate initiatives.

What frameworks do you follow to ensure responsible and secure AI deployment, especially in regulated sectors?

In regulated sectors, responsible AI cannot be left to policy statements alone. It has to be built on a framework that combines governance, explainability, validation, and human accountability. Our starting point is that every AI use case should answer four questions clearly—what it is allowed to do, what data it can use, who owns it, and how it is monitored once it is live.

From there, we look at model risk, data quality, access control, bias exposure, auditability, and the potential business impact of failure. In practice, this means setting clear approval thresholds, maintaining traceability across the data and model lifecycle, using human-in-the-loop controls for high-risk decisions, and

“Modernisation ultimately becomes a redesign of enterprise trust, where cloud transformation, AI adoption, and cybersecurity maturity must advance together as a single integrated strategy rather than separate initiatives

continuously testing behaviour in production rather than assuming the model will remain stable.

For regulated environments, the framework must also map to the relevant compliance and assurance expectations of the sector. The point is not to treat compliance as an afterthought, but to make it part of the operating model for AI.

Responsible AI is ultimately measured not by a model’s capability but by the level of trust an organisation can safely place in its outcomes—marking the difference between experimentation and true enterprise deployment.

How does Blue Cloud Softech approach data sovereignty and localization,

especially in government and enterprise deployments?

Data sovereignty is no longer a narrow legal issue; it is a strategic design requirement. In government and enterprise environments, the question is not only where data resides, but who can access it, under what jurisdiction it is governed, and how control is maintained across its lifecycle.

Our approach is to design for localization, control, and policy alignment from the outset. That means understanding residency requirements, defining access boundaries, applying encryption and tokenization where needed, enforcing role-based controls, and ensuring that sensitive data is handled in a way that aligns with the client's regulatory and operational obligations. In large deployments, sovereignty also depends on architecture choices: how data is segmented, how workloads are hosted, and how visibility is maintained without creating unnecessary exposure.

For public-sector and regulated enterprise use cases, trust depends on the ability to prove that data is not only secure but also governed in a way that respects local requirements and organisational accountability.

In data-sensitive environments, sovereignty extends beyond storage location to encompass control, accountability, and enforceable governance—particularly as organisations operate across increasingly distributed cloud and hybrid ecosystems.

With cyber threats becoming more sophisticated, how do you design systems that anticipate threats rather than react to them?

The shift has to be from reactive defence to continuous exposure management. Threats today are not waiting for perimeter defences to fail; they are exploiting identity weaknesses, misconfigurations, exposed services, third-party dependencies, and human behaviour. So, designing for anticipation means assuming that some form of exposure will exist and building

systems that can detect, isolate, and respond before that exposure becomes material.

This begins with visibility. You cannot anticipate what you cannot see. From there, you need continuous asset awareness, attack surface monitoring, behavioural analytics, threat intelligence, and automated response paths. But equally important is operational discipline: security controls must be tested, monitored, and updated continuously, not just reviewed periodically.

The organisations that are most resilient are the ones that combine prevention with readiness. They know where their critical assets are, what abnormal behaviour looks like, and how to contain an incident without waiting for manual intervention.

Modern cybersecurity is defined less by perfect prevention and more by early visibility, rapid containment, and sustained resilience—making systems significantly harder to exploit and far quicker to recover.

How is Blue Cloud Softech preparing for a future where every system is interconnected and continuously exposed?

The future enterprise will not be a closed environment. It will be interconnected, distributed, API-driven, AI-assisted, and always exposed to some degree of risk. Preparing for that future means designing for resilience rather than assuming isolation. Security has to become adaptive, continuous, and architecture led.

Our focus is on building systems where identity, data protection, monitoring, response, and governance work together as a connected control fabric. In a continuously exposed environment, the goal is to reduce attack paths, improve detection speed, strengthen containment, and ensure that critical services remain dependable under pressure.

That requires cross-functional thinking. Cybersecurity can no longer operate as a standalone function; it has to be embedded

Modern cybersecurity is defined less by perfect prevention and more by early visibility, rapid containment, and sustained resilience—making systems significantly harder to exploit and far quicker to recover

across cloud, AI, data, and application layers. It also requires organisations to move from episodic assurance to continuous assurance, where systems are assessed, monitored, and improved in real time.

In an always-connected digital environment where resilience becomes the new perimeter, the organisations that succeed will be those able to innovate rapidly without compromising trust, control, or operational continuity. ■

editor@thefoundermedia.com

INTERVIEW

WHAT MOST FOUNDERS GET WRONG WHEN THEY DECIDE TO BUILD AN AI-FIRST PRODUCT

Nandagopal P, Chief Technology Officer, Gacsym Ventures, shares with **Aishwarya Saxena** why AI lowers the cost of building but never the cost of bad judgment — and what that means for founders

As CTO in a venture studio environment, how do you define your role differently from a traditional startup CTO?

The honest answer is that a traditional startup CTO has one bet to make. I'm managing several simultaneously, which changes almost everything about how you think.

In a single startup, you can afford to go deep on every decision. You live with the consequences long enough that slow thinking is fine. In a venture studio, that luxury disappears. You have to build judgment systems, not just make judgments. What patterns are reusable? Where does speed matter more than perfection? What can you validate before writing a single line of code?

The role stops being primarily technical and becomes something closer to decision architecture. How do we make better calls under uncertainty, faster, with less waste? After working with 70+ companies, I've noticed that most technical mistakes in early-stage startups aren't really technical mistakes. They're business mistakes that got expressed in code. The CTO's job, especially in a studio, is catching those upstream.


“
Early architecture buys you learning speed and the ability to change direction without starting over

How does your CTO-as-a-service approach influence product architecture decisions in early-stage startups?

The first thing I try to do is separate architecture from ego. A lot of early technical decisions are made to impress other engineers, not to serve the business. That's a slow-burning problem. My actual approach is simple: design for the next credible stage, not for some imagined future where you have millions of users. That usually means keeping things modular, straightforward, and easy to change when the business shifts, which it will.

What working across multiple startups gives you is pattern recognition that's hard to fake. A FinTech product and an AI workflow tool should not be making the same architectural decisions just because the same tools are popular right now. Context matters enormously. The business model, the founder's pace, the compliance exposure, the likely user behavior, all of it shapes what good architecture actually looks like for that specific company at that specific stage.

The goal I keep coming back to is this: good early architecture buys you learning speed. Great early architecture buys you learning speed and the ability to change direction without starting over.

How do you evaluate whether a startup should integrate AI at its core versus using it as an enhancement layer?

I usually start with one question: if you took the AI out completely, would the product still matter?

If yes, AI is probably an enhancement. It can make things faster, smarter, and more personalized. That's genuinely valuable. But the core is something else. If the answer is no, then AI might belong at the center. But even then, I push hard on whether the startup owns anything defensible beyond the API call. This is where I see a lot of founders get confused. Novelty and necessity are not the same thing. Just because AI can be inserted into a product doesn't mean it should define the company's identity. Users don't pay for AI. They pay for outcomes. Speed, trust, clarity, results. If those are what's being delivered, the architecture and the story should reflect that, not lead with the technology.

The four things I actually look at are: how real is the user pain, does the AI advantage get stronger over time with data or usage, do the economics work, and is there anything defensible here that a competitor can't replicate next quarter. If AI materially changes the outcome and compounds with use, it belongs at the core. Otherwise, use it as a force multiplier and build your moat somewhere else.

What are the biggest challenges insurers face when implementing digital-first customer journeys at a scale?

Insurance sounds like a straightforward digitization problem until you get close enough to see what's actually underneath.

The first real challenge is fragmentation. A customer journey in insurance doesn't live in one place. It cuts across quoting, underwriting, KYC, claims, servicing, partner networks, and compliance layers that were never designed to connect cleanly. Digitizing the front-end without fixing the plumbing just makes the broken parts faster.

The second is legacy infrastructure. Most insurers are trying to deliver modern, fluid experiences on systems that were built for stability above everything else. Not agility, not speed, not iteration. That tension is genuinely hard to resolve without significant investment and tolerance for disruption.

But the one that gets underestimated is the emotional layer. Insurance shows up in people's lives during their worst moments. Accidents, illness, loss, uncertainty. A digital journey in that context can't just be efficient. It has to feel human, clear, and trustworthy. That's a much harder design problem than most technological teams are set up to solve. The insurers who will actually win here are not treating this as a front-end redesign project. They're rethinking the whole operating model, data, workflows,

risk systems, service design, and where humans need to be in the loop.

How do you see the venture studio model evolving in a world where product development is increasingly commoditized by AI tools?

Honestly, I think AI strengthens the studio model more than it threatens it.

When building gets cheaper, the real bottleneck moves. It shifts from "can we build this" to "should this exist, will people actually care, and can this team execute repeatedly over time." Execution capacity alone stops being a differentiator. Judgment becomes the scarce resource.

The studios that will struggle are the ones whose primary value was helping founders build things faster. That advantage compresses quickly. The ones that will get stronger are those with real domain knowledge, sharp thesis formation, genuine distribution insight, and the operational systems to turn a new idea into a real company without starting from scratch every time. What I think the next generation of studios looks like is less of a shared dev shop and more of an operating system for company creation. Fast product velocity, yes, but combined with better validation frameworks, tighter feedback loops, and smarter experimentation on go-to-market. AI lowers the cost of building. It does not lower the cost of bad judgment. That gap is where studios either earn their value or don't.

What common architectural mistakes do you see founders make when building their first product?

The most consistent one is overbuilding before earning the right to complexity. I see it constantly. Founders design for scale before they've confirmed anyone cares. They build multi-service systems for products with no real users yet. They introduce abstraction layers nobody needs. They optimize for hypothetical problems two years away while real friction exists right now in front of them.

> **A FinTech product and an AI workflow tool should not be making the same architectural decisions just because the same tools are popular right now**

A close second is letting tool popularity drive architectural decisions instead of business reality. The team picks what's fashionable rather than what they can actually operate well under pressure. That creates fragility at exactly the wrong time.

The one that surprises founders most when they catch up with them is weak data design. Everyone focuses on UI, features, and flows. But the data model is quietly making decisions for you the whole time. Poor data structures become expensive constraints that are genuinely painful to fix later. The thing I try to instill early is this: build the cleanest version of what you actually know to be true today, while keeping the door open to change tomorrow. Most first products don't fail because they were too simple. They fail because they became too rigid to adapt. Simplicity isn't a limitation in the early stage. It's a strategy.■

editor@thefoundermedia.com

INTERVIEW

HELO.AI BUILDING THE FUTURE OF AI-FIRST OMNICHANNEL COMMUNICATION

Vikram Raichura, Founder and MD, Helo.ai, sits down with **Aishwarya Saxena** to reveal the bold architectural choices and privacy-first principles powering Helo.ai's vision for the next generation of omnichannel communication

Helo.ai positions itself as an AI-first omnichannel platform. What specific architectural decisions differentiate this from legacy CPaaS platforms?

Helo.ai is designed with AI as the core decisioning layer that operates before message orchestration, rather than being added on top of existing delivery systems. The architecture is built on event-driven workflows and native model integration, moving away from traditional channel-centric pipelines where AI is typically an afterthought.

How do you ensure scalability of AI inference across millions of concurrent conversations?

Helo.ai follows a tiered inference strategy where lightweight models manage the majority of interactions, while heavier LLM calls are deployed selectively. This approach, combined with asynchronous processing, dynamic model routing, and horizontal scaling, ensures low latency and consistent performance even at massive scale.

What challenges arise in unifying structured and unstructured data across channels, and how do you address them?

One of the primary challenges is handling

fragmented data sources while maintaining fast access during live interactions. Helo.ai is developing a centralized data store for frequently accessed structured data, alongside a lightweight RAG layer that dynamically retrieves and contextualizes unstructured information as needed.

How does Helo.ai handle ambiguity, slang, and regional language variations in user queries?

Helo.ai is developing a framework that leverages multilingual LLMs to better interpret mixed language inputs, slang, and informal communication styles. This includes exploring on-premises deployable open-source models such as LLaMA, Mistral, and Indic-focused models like IndicBERT.

How do you ensure compliance with data privacy regulations while using AI to process customer data?

Helo.ai integrates a PII masking and tokenization layer to safeguard sensitive data before it enters AI processing pipelines. The platform is built with minimal data exposure, audit logging, and support for private or on-premises deployments to meet enterprise compliance needs. Security and compliance measures include:

- End-to-End Encryption
- PII Masking & Tokenization
- Role-Based Access Control (RBAC)
- Audit Trails & Logging
- Compliance Alignment
- Local data residency support for India

AI models are also designed to avoid storing sensitive data unnecessarily and use anonymized datasets wherever possible.

> Helo.ai integrates a PII masking and tokenization layer to safeguard sensitive data before it enters AI processing pipelines

How is AI improving engagement in BFSI use cases like alerts, onboarding, and customer support?

AI is enabling a shift from static, one-way communication to more interactive and adaptive customer experiences — where alerts become actionable, onboarding turns conversational, and support becomes more responsive and intelligent. This helps improve overall engagement, reduce drop-offs, and streamline customer interactions across all BFSI journeys.■

editor@thefoundermedia.com

INTERVIEW

FLEET MANAGEMENT'S NEXT ERA IS BEING BUILT ON AI, NOT GPS

Bhanutej Mallangi, Chief Product Officer, ROQIT, tells **Aishwarya Saxena** why carbon tracking, asset telemetry, and predictive intelligence are no longer future priorities but are the present competitive edge in fleet management

Building a platform that serves fleet managers, OEMs, vehicle makers, and financiers simultaneously is complex. How do you prioritize product features when each stakeholder group has fundamentally different needs?

ROQIT approaches this as a modular platform design problem rather than a feature prioritization exercise. Different stakeholders require different outcomes but building separately for each lead to fragmentation. Instead, the platform is structured around core modules, such as asset tracking and trip management, that can be configured across use cases. This allows us to serve multiple stakeholders without over-engineering for any single one.

At the core is a simple abstraction: an asset represents any operational unit, and a trip represents its movement or utilization cycle. Whether this is a vehicle on a route or equipment operating within a defined cycle, the same building blocks apply.

Today, customers are increasingly relying on ROQIT to understand how much of their asset movement is actually productive versus incidental, and this is already starting to change how operators identify

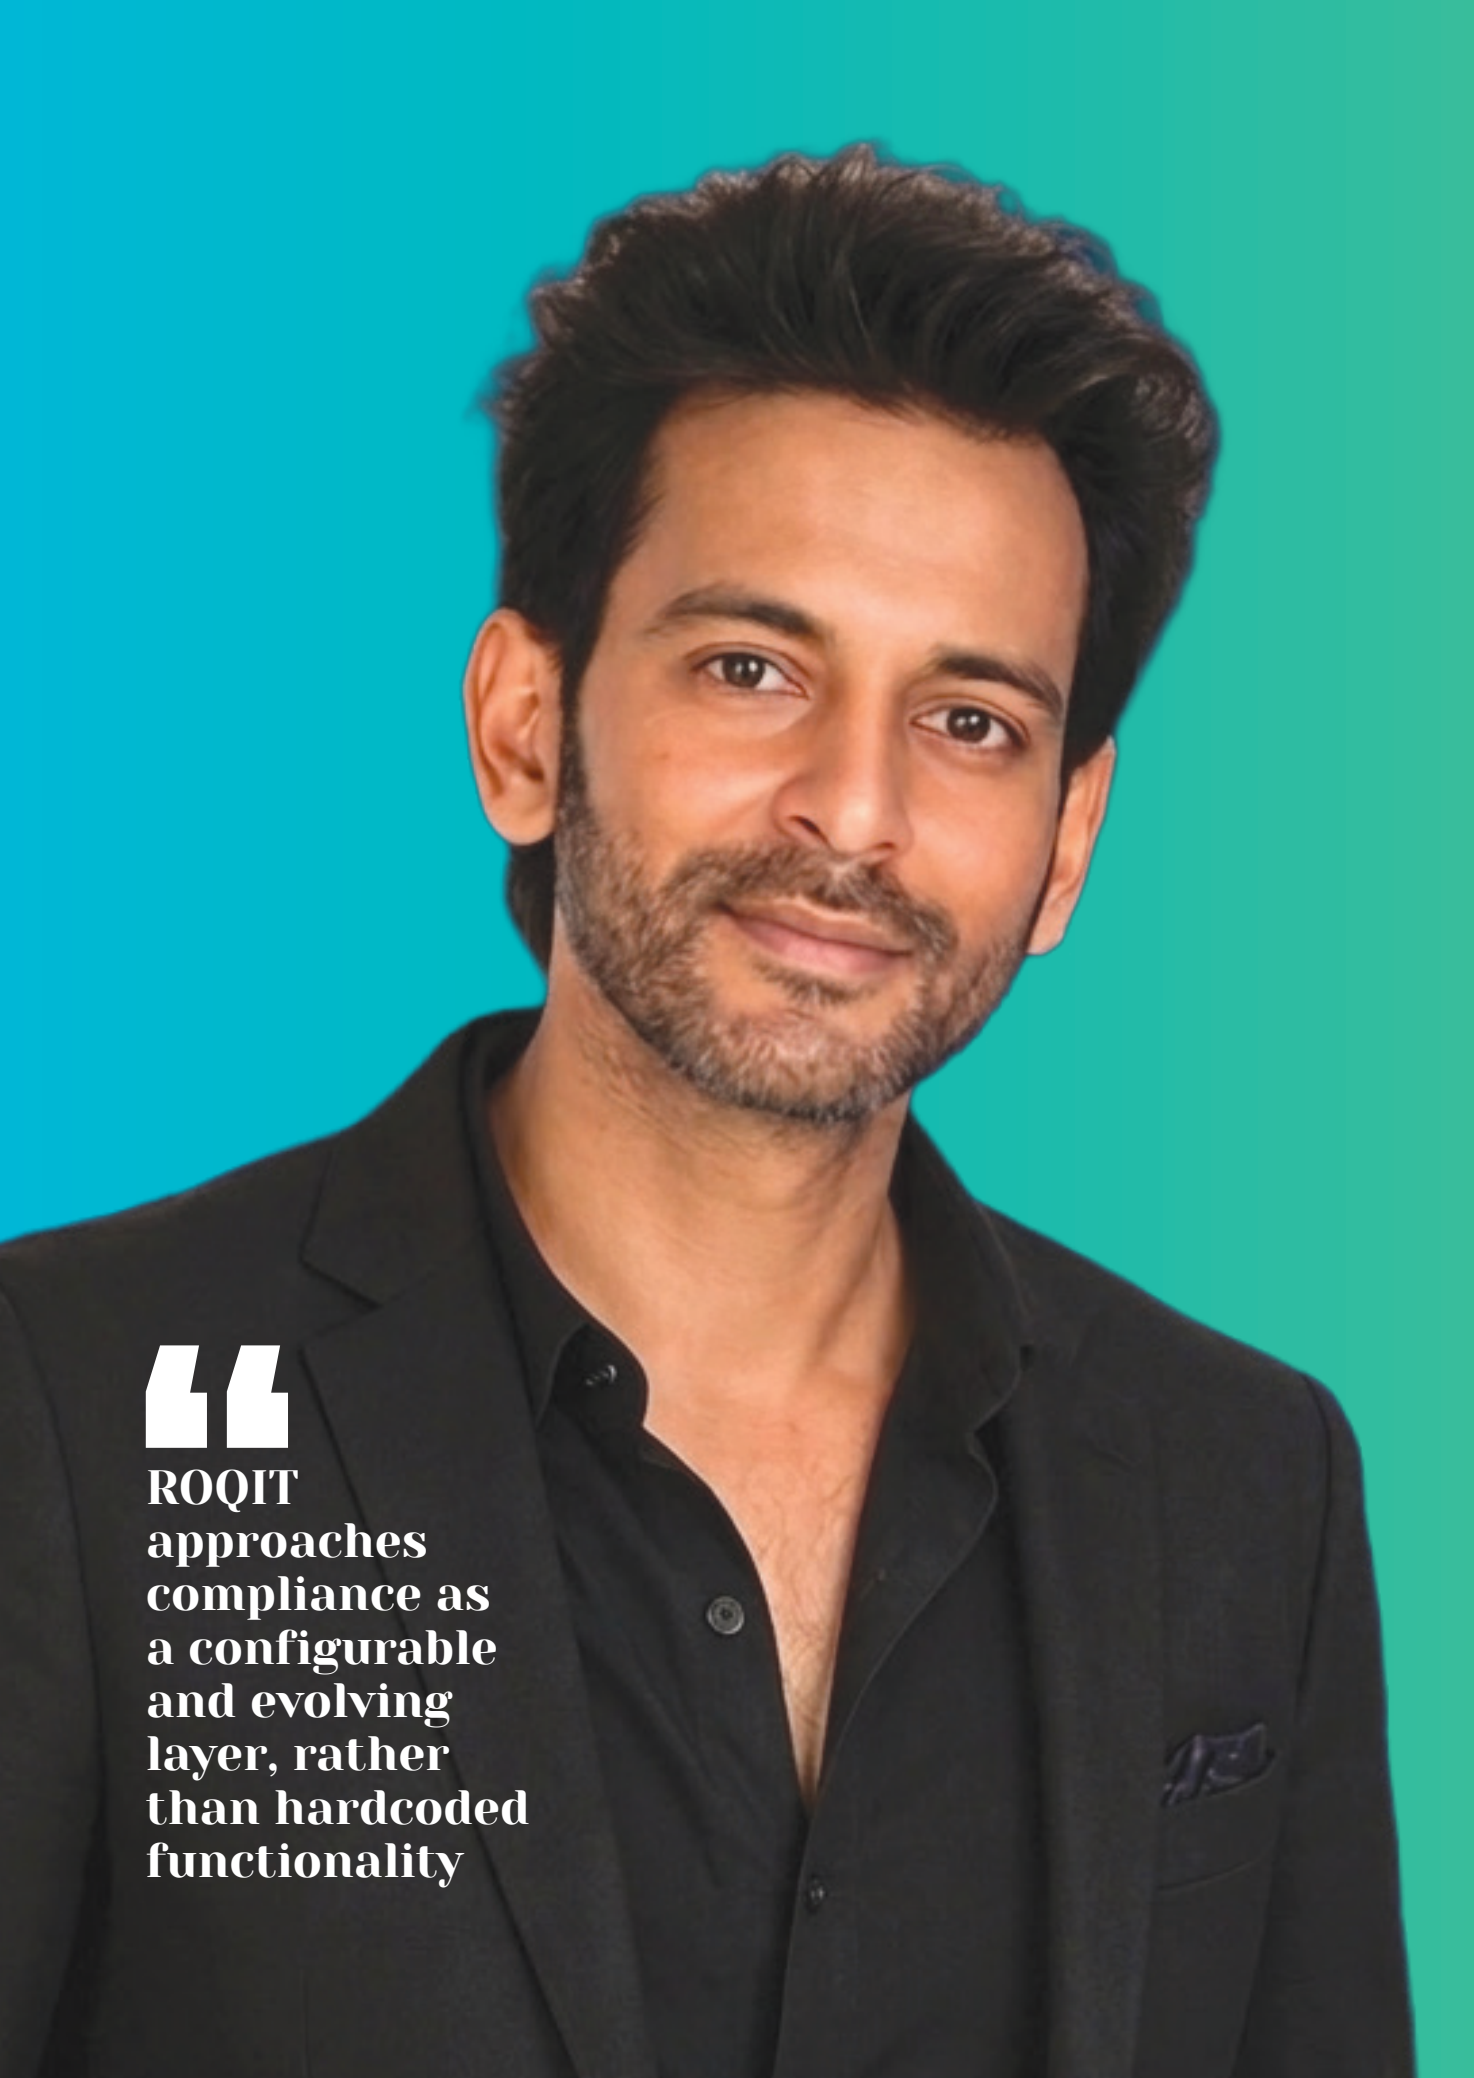

ROQIT approaches compliance as a configurable and evolving layer, rather than hardcoded functionality

and reduce non-productive movement in their operations.

Most systems focus on where an asset is. ROQIT focuses on whether its movement is actually creating value. The industry has largely solved for visibility—the next challenge is enabling better operational decisions. This approach is already being applied across varied operating environments, including large-scale electric fleet operations that are closely tied to broader logistics networks, where the same system can adapt to very different movement patterns without structural changes.

India's EV charging infrastructure is still patchy. How does the platform factor in charging infrastructure gaps when doing route planning or fleet scheduling?

ROQIT is designed to operate within real-world constraints rather than ideal conditions, particularly in environments where EV infrastructure is still evolving. Today, the platform focuses on trip intelligence and asset utilization, helping operators distinguish between productive (revenue-linked) and non-productive movement. In current deployments, this is already starting to change how operators identify and reduce non-productive movement, and how routes and operations are structured.

In one such environment, the platform is being used across a large electric fleet operating at scale, where trip creation and execution at the fleet layer are directly aligned with downstream logistics flows. Operational systems can generate trips, which can be orchestrated through ROQIT and fulfilled through driver-facing applications, ensuring continuity between planning, movement, and fulfilment.

A "trip" in this context is not just distance covered, it represents a unit of work or output, whether that's a delivery route or an operational cycle. On the EV side, ROQIT integrates available telemetry and is evolving to incorporate charging-aware insights, including station availability through ecosystem integrations.

Looking ahead, we are working towards enhancing route intelligence by incorporating factors such as energy feasibility, terrain, payload, and traffic conditions, progressively moving towards more predictive and EV-optimized planning.

ROQIT positions itself at the intersection of electrification and AI-led transformation. But for fleet operators who still run mixed fleets like part diesel, part EV. How does the platform handle hybrid fleet management without penalizing operators still in transition?

Most fleets today operate in a hybrid state, and ROQIT is designed to support this transition without introducing operational complexity. The platform provides a unified operational layer, enabling operators to manage different asset types within a single system, with consistent visibility across utilization, trip patterns, and efficiency. This abstraction allows fundamentally different asset behaviors to be compared through a

common lens—how effectively each asset is being utilized within its operational cycle.

At the same time, we recognize the need for deeper, asset-specific intelligence. Capabilities such as battery health analytics, degradation tracking, and advanced diagnostics are part of the next phase of the platform's evolution.

Our approach is to establish a strong operational baseline first, and then progressively layer in specialized intelligence, ensuring that decision-making improves over time without disrupting existing workflows.

ROQIT is built to comply with India's national and state-level EV policies, which are still evolving rapidly. How do you build regulatory compliance into the product without making every policy update a major re-engineering effort?

ROQIT approaches compliance as a configurable and evolving layer, rather than hardcoded functionality. While this capability is still being built out, the direction is towards a rule-driven framework where regulatory requirements can be translated into configurable guardrails, reducing the need for repeated engineering changes. Today, the platform captures granular operational and usage data at the asset and trip level, which forms the foundation for compliance and reporting.

Building on this, we are working towards enabling policies to be applied dynamically based on factors such as geography, asset type, and usage context, ensuring that compliance aligns with how assets are actually being operated. This becomes increasingly important in environments where regulatory frameworks vary significantly across use cases, requiring flexibility without compromising consistency.

Looking ahead three to five years, where do you see the biggest technology disruptions in fleet and logistics management and how is ROQIT's product strategy positioning itself to lead rather than react?

> **Most fleets today operate in a hybrid state, and ROQIT is designed to support this transition without introducing operational complexity**

Over the next 3–5 years, fleet and logistics management will be shaped by three shifts: AI-driven decision-making, deeper asset intelligence, and the emergence of carbon as a measurable and monetizable layer. A key shift will be how movement is interpreted—not just as distance covered, but as output generated. Platforms that can contextualize asset behavior within its purpose will define the next phase of operational intelligence.

ROQIT is being built around this principle. By structuring the system around assets and their operational cycles and acting as a layer through which trips can be orchestrated across systems, the platform is designed to extend across industries where utilization, efficiency, and output are tightly linked. This allows ROQIT to evolve from a fleet management system into a broader asset intelligence and orchestration platform, without changing its core architecture. The focus is on enabling customers to measure, optimize, and derive financial and operational value from how their assets are used.■

editor@thefoundermedia.com

INTERVIEW

63SATS CYBERTECH SECURING ENTIRE AI LIFECYCLE FROM DATA INTEGRITY TO MODEL VALIDATION

Neehar Pathare, MD, CEO and **CIO, 63SATS Cybertech**, tells **Aishwarya Saxena** why the moment an organisation stops assuming the intruder is already inside is the moment it becomes genuinely vulnerable

What inspired the positioning of 63SATS as a "Cyber Security Force" for enterprises and governments?

Think of the traditional approach to security like a fire department, they are world class at putting out fires, but they only arrive after the smoke appears. Positioning 63SATS as a Cyber Security Force (CSF) represents a shift from being a "firefighter" to being a "specialized defense unit".

In a world where digital threats are constant, you cannot afford to be reactive. We operate with a "defense-oriented" mindset, focusing on continuous monitoring and rapid response. Much like a physical security force patrolling a perimeter, we use real-time threat intelligence to stop intruders before they breach the gates. For our partners in government and enterprise, this means we aren't just a vendor they call during a crisis; we are a strategic partner integrated into their daily operations to ensure their systems remain resilient and their data stays trusted.

Having transitioned from core IT leadership roles into cybersecurity, what

> When we build new technology, we don't "add on" security at the end; we frame it in from the very first line of code

mindset shift was required?

In traditional IT leadership, the goal is "Performance and Uptime". You want the engine to run smoothly, efficiently, and fast. However, in cybersecurity, the starting assumption is much darker: the "engine" is already under attack, and the intruder might already be inside the cabin.

The shift required is moving from a "Reliability" mindset to a "Resilience" mindset. It's no longer enough to build a strong wall; you have to think like the person trying to climb over it. This means embedding security into the very DNA of an organization from how a server is configured to how an employee handles an email. At 63SATS, we "stress-test" this mindset using Red Teaming, where we simulate real-world attacks to find gaps before a malicious actor does.

With AI increasingly embedded in enterprise systems, how do you mitigate risks like data poisoning and adversarial attacks?

AI is a double-edged sword. It's incredibly powerful, but its "brain" is only as good as the data its fed. Data poisoning is like someone slipping "false memories" into an AI's education, causing it to make dangerously wrong decisions.

To mitigate this, we secure the entire AI Lifecycle.

Data Integrity: We verify the "purity" of the data at the very first stage to prevent manipulation.

Model Validation: We treat the AI model like a black box that needs constant inspection to ensure it hasn't been tampered with or developed "biases".

Zero-Trust for AI: We apply Zero-Trust principles, meaning no input is trusted by default. Every interaction with the AI must be verified, and we use our own AI-driven detection to spot patterns of "adversarial attacks" that a human might miss.

How does your use of frameworks like MITRE ATT&CK enhance your ability to simulate real-world attack scenarios?

In the past, security testing was often a "checkbox" exercise. Did you lock the door? Did you close the window? Frameworks like MITRE ATT&CK allow us to move beyond checkboxes to Adversary Emulation.

Think of MITRE ATT&CK as a massive, globally updated encyclopedia of every move a "burglar" has ever used. Instead of

just looking for a "weak lock," we map out the entire "heist". We can simulate how an attacker would move through a network, what they would try to steal, and how they would try to hide their tracks. This gives organizations a clear map of not just their weaknesses, but their gaps in detection and recovery, providing actionable insights to harden their defenses effectively.

With increasing reliance on third-party vendors and SaaS tools, how do you embed supply chain security into modernization strategies?

Modern businesses are "interconnected." When you use a third-party software or a cloud service, you aren't just buying a tool; you are opening a door into your house for a stranger. Supply chain risk is the danger that this "stranger" (the vendor) might have their own security flaws that lead back to you.

At 63SATS, we treat supply chain security as a non-negotiable part of digital transformation.

Lifecycle Assessment: We don't just vet a vendor when they sign a contract; we monitor them continuously.

Extended Zero-Trust: We extend our security perimeter to include these external tools. Every time an external tool tries to talk to your internal systems, it must be authenticated, authorized, and verified. This allows companies to scale and innovate without worrying about "hidden" risks lurking in their ecosystem

> **We can simulate how an attacker would move through a network, what they would try to steal, and how they would try to hide their tracks**

As MD, CEO, and CIO simultaneously, how do you balance business growth, technology innovation, and security governance?

People often see security as a "brake" on a car, it slows you down. I see it as the reason you can drive at 100 mph with confidence. Balancing these roles requires understanding that they are all interconnected drivers of success.

Growth is built on trust: You cannot grow a business if your customers don't trust you with their data.

Innovation requires "Security- by Design": When we build new technology, we don't "add on" security at the end; we frame it in from the very first line of code.

Governance as a blueprint: Governance provides the rules of the road, allowing us to scale responsibly while staying compliant with global standards.

By integrating these, security becomes a business enabler rather than a constraint. It gives an organization the "Digital Sovereignty" to evolve and compete on a global stage, knowing their foundation is rock-solid.

editor@thefoundermedia.com

VIEWPOINT

CLOSING THE SKILLS GAP IN AN EVOLVING TECHNOLOGY LANDSCAPE

In this thought-provoking piece, **Amit Patil, MD & Founder, CynalitX Consulting LLP**, highlights why reskilling is no longer optional but a strategic necessity.

Artificial intelligence, cloud computing, cybersecurity, and data analytics are no longer niche specialisations. They are the operating terminology of modern business. Yet a significant and widening gap exists between the skills organisations need and the talent available in the market. For business leaders, this is not a future problem. It is a present crisis demanding strategic action today.

The scale of the challenge

The World Economic Forum estimates that over 1 billion people will need reskilling by 2030. Closer to the boardroom, the story is just as straightforward; technology roles go unfilled for months, projects stall, and competitive advantage erodes, not because of a lack of ambition, but because of a lack of capable hands.

The irony is that technology itself is partly responsible for this gap. Automation and AI are simultaneously displacing certain roles and creating entirely new ones often faster than traditional education systems can respond. A degree earned five years ago may already be partially obsolete. The life of a technical skill is shrinking before it matures.

The case for internal upskilling

The most resilient organizations are not just hiring their way out of the skills gap but are building from within. In the consulting and

> India's policy momentum and the sector's own innovators offer both inspiration and a practical roadmap

professional services sector, this imperative is especially critical. Client expectations are evolving rapidly; engagements that once required domain expertise now demand fluency in AI-driven analytics, automation tooling, and digital transformation frameworks.

A leading global professional service organisation have responded decisively. Its Future Talent Platform has committed to training over 300,000 or more professionals annually in emerging technologies, from generative AI to cloud architecture, embedding learning directly into project delivery cycles rather than treating it as an offline activity. The result is a workforce that remains billable, relevant, and ahead of client needs simultaneously. For leaders, this is the benchmark worth measuring against.

Effective upskilling strategies share three traits

They are continuous, not occasional. A one-time training initiative does not build a learning culture. The best programs are embedded into the rhythm of work which has a structured time for learning, mentorship pathways, and access to on-demand digital platforms. In organisations, where practitioners move between engagements, competency credentialing and modular learning paths are particularly effective.

They are role specific and outcome driven. Generic digital literacy workshops rarely move the needle. Leaders must identify the precise skills their business will need 18 to 36 months from now, which might include AI-augmented advisory, data-led strategy, cybersecurity consulting, etc, and design learning journeys that map directly to those outcomes. They are visible from the top. When the Managing Director or COO champions learning openly by discussing their own reskilling journey it signals to the entire firm that growth is valued, not just utilisation rates.

The role of policy— India's growing ambition

Internal efforts alone, however, cannot close a gap this large. For leaders operating in or from India, the policy landscape offers

both a signal and a resource worth taking seriously.

Skill India Mission

India's Skill India Mission represents one of the most ambitious workforce development programmes in the world, targeting the skilling, reskilling, and upskilling of hundreds of millions of citizens across sectors. For firms with large delivery centres and talent pipelines rooted in India, this initiative is directly relevant as it is shaping the entry-level and mid-career talent pool that the industry depends on.

National Skill Development Corporation (NSDC)

The NSDC plays a pivotal coordinating role, bridging government intent and industry execution. Its partnerships with private training providers and technology platforms create a framework that firms can actively plug into by co-designing curricula, sponsoring certification pathways, or accrediting internal programs. Forward thinking leaders are already doing this, turning NSDC partnerships into a talent pipeline advantage.

PM Vishwakarma Scheme

While focused on traditional craftsmanship and artisanal skills, the PM Vishwakarma Scheme indicates something broader and important; the government's recognition that skilling is not one-size-fits-all, and that digital tools must be made accessible across the full spectrum of India's workforce. For leaders in manufacturing, MSME, or rural development sectors, understanding this scheme is increasingly part of delivering credible, grounded recommendations.

Taken together, these initiatives represent a national infrastructure for workforce development that business leaders should engage with and not observe from a distance. That means contributing to policy consultations, co-investing in skilling ecosystems, and advocating for frameworks that reward enterprise-level training investment.

> **The pace of technological change has never been faster, and the workforce is struggling to keep up**

The leadership imperative

Closing the skills gap is not an HR agenda item. It is a boardroom priority. Leaders in consulting and professional services face a particular version of this challenge; their product is their people. A firm whose practitioners cannot speak the language of AI, data, and digital transformation will not merely struggle to grow but will struggle to remain relevant to clients who are already ahead.

The organisations that will define the next decade of professional services are those investing now in the capabilities their clients or product will demand tomorrow. India's policy momentum and the sector's own innovators offer both inspiration and a practical roadmap.

VIEWPOINT

RETHINKING ENTERPRISE CLOUD IN INDIA: FROM ADOPTION TO STRATEGIC CONTROL

As enterprises grapple with rising cloud costs, data sovereignty concerns, and AI-driven workloads, **Manoj Dhanda, Founder and CEO, Utho Cloud**, speaks on what it truly takes to build future-ready infrastructure in India

India's cloud journey has reached an interesting inflection point. Over the past decade, enterprises moved rapidly toward cloud platforms, largely driven by ease of deployment and global best practices. But today, the conversation is changing. It is no longer about whether to adopt cloud—it is about how to use it intelligently.

What we are seeing now is a shift from adoption to optimization. Enterprises are becoming more deliberate in how they design infrastructure, how they manage costs, and how much control they retain over their data and systems.

Scalability is no longer about one cloud

Earlier, scalability meant choosing a large cloud provider and building everything within that ecosystem. That model worked well when workloads were simpler and largely CPU-driven. But today, things are different, especially with AI. Now, workloads are not uniform. Some require

India is at a stage where it should not just consume global cloud services but also build its own infrastructure ecosystem

specific GPU configurations, and those are not always available with a single provider. This is one of the main reasons why multi-cloud is becoming the default approach.

In fact, most organizations today are either already using multiple providers or actively planning to do so. The decision is no longer about brand preference—it is about availability, performance, and cost at a given point in time.

At the same time, hybrid cloud—combining private and public infrastructure—has not scaled as much as expected. One key reason is that private cloud environments often slow down innovation. They struggle to match the pace of cloud-native development, which limits how quickly teams can adopt new technologies. So, scalability today is not about scaling within one system. It is about building distributed architectures that allow flexibility across environments.

The growing importance of sovereign cloud

Another shift that is becoming very clear is the importance of data sovereignty. Enterprises in India are increasingly concerned about where their data is stored and under whose jurisdiction it falls. This is not just a regulatory issue—it is a strategic one.

When businesses rely entirely on global infrastructure, they also inherit certain risks—whether it is compliance complexity, pricing unpredictability, or geopolitical uncertainty. This is where sovereign cloud becomes critical. It ensures that data is stored within India, governed by Indian laws, and aligned with local regulatory frameworks. But more importantly, it gives businesses greater control and confidence.

India is at a stage where it should not just consume global cloud services but also build its own infrastructure ecosystem. As digital adoption accelerates, having a strong, indigenous cloud backbone is becoming foundational for long-term growth and resilience.

Security begins with infrastructure, but doesn't end there

When enterprises evaluate cloud, one of the biggest concerns is security. But in practice, the first concern is usually compliance, followed by trust. Large organizations want to ensure that the infrastructure they use meets global standards. This is where Tier

III and Tier IV data centers play a critical role. They provide high availability, redundancy, and operational reliability—forming the foundation of secure cloud environments.

However, security is not just about the data center. It is also about how systems are managed. Cloud operates on a shared responsibility model. While providers secure the infrastructure, enterprises are responsible for how they configure and use it. That means governance becomes extremely important—things like access control, monitoring, and resource management. Trust is built over time. It comes from consistent performance, transparency, and the ability to clearly understand how systems are operating—not just from certifications alone.

Cloud vs Dedicated: A more practical approach

Another area where thinking is evolving is the choice between cloud servers and dedicated infrastructure. In the early days, many companies moved to cloud without fully understanding the cost implications. Free credits made it easy to experiment, but once those credits expired, the real costs started to surface.

If you break it down, most cloud expenses come from three areas—compute, storage, and bandwidth. And without proper governance, these can scale quickly.

Today, businesses are far more cost-conscious. Instead of defaulting to cloud, they are asking more practical questions:

- Do we need flexibility or predictability?
- Is the workload stable or dynamic?
- Do we need full control or managed infrastructure?

Dedicated environments still make sense for predictable workloads where control and performance are critical. Cloud environments are better suited for dynamic workloads where scalability is required. What we are seeing is not a shift toward one model, but toward balanced architectures, where decisions are made based on actual workload needs.

“What we are seeing is not a shift toward one model, but toward balanced architectures, where decisions are made based on actual workload needs

The road ahead: A more independent cloud ecosystem

Looking ahead, the cloud ecosystem in India will become more distributed, flexible, and independent. Enterprises will move away from being locked into a single provider and toward architectures that allow them to choose the right environment for each workload. At the same time, infrastructure will become more abstracted, with organizations interacting through APIs rather than managing hardware directly.

AI will play a big role in this shift, pushing demand for specialized infrastructure and accelerating the move toward distributed systems. However, while technology will evolve quickly, enterprise adoption will remain cautious. CIOs and leadership teams will continue to balance innovation with risk. India’s cloud journey is entering its most important phase. The focus is no longer just on scaling—it is on scaling intelligently. That means optimizing costs, ensuring data sovereignty, strengthening security, and building flexible infrastructure strategies. The future will not belong to a single cloud model or provider. It will belong to organizations that can adapt, distribute, and make informed infrastructure decisions aligned with long-term goals.■

VIEWPOINT

WHY AI-POWERED VOICE IS THE NEXT FRONTIER IN ENTERPRISE CUSTOMER ENGAGEMENT

Alok Anibha, Founder, Girikon.AI, explains why AI-powered voice technology is no longer just a call-routing tool but a transformative force reshaping how enterprises connect with customers at scale

In recent years, enterprises have invested heavily in digital customer engagement tools: chatbots, messaging platforms, and self-service portals. These tools now play a standard role in enterprise operations across industries.

While these solutions have undoubtedly enhanced efficiency and streamlined processes, they have not replaced one fundamental reality: when interactions become important or complex, customers still prefer to speak rather than type.

And voice remains the most natural and intuitive way humans communicate. Therefore, for many years, enterprise voice systems have struggled to keep up with rising customer expectations. Anyone who has navigated a long IVR menu knows how quickly frustration can build when technology fails to understand the customer's intent.

Today, artificial intelligence is changing how voice operates in enterprise engagement. AI-powered voice systems

For enterprises investing in AI voice technology, measurement is essential

now turn voice from a call-routing tool into a conversational interface. These systems can understand context, emotion, and intent.

For every technology leader, this shift actually represents a major opportunity to rethink how enterprises interact with customers at scale.

Moving beyond the limitations of traditional IVR

Earlier, Traditional Interactive Voice Response systems were designed primarily for efficiency, allowing organisations to route calls, automate simple tasks, and manage high volumes of inbound requests.

While this approach helped reduce operational costs, it often came at the expense of customer experience.

IVR systems typically rely on rigid menus and predetermined workflows. Customers must select from predefined options, even when their needs do not fit neatly into those categories.

As a result, interactions become slower and more frustrating, and customers often repeat information before reaching the right agent

AI-powered voice systems address this limitation by enabling natural conversation. Using natural language processing and machine learning, these systems can interpret spoken language, detect intent, and respond in ways that feel far more intuitive.

For example, instead of asking the customers to “press one for billing” or “press two for support,” AI voice systems now allow users to simply describe their issue in their own words and language.

After that, the system analyzes the request and determines the most relevant response or routes the conversation to the appropriate human expert.

So, from a technology leadership perspective, this shift is not simply about automation but about creating a more intelligent interface between customers and enterprise systems.

Why enterprises are re-evaluating voice

Several forces are driving the renewed focus on voice technology. First, customer expectations have changed dramatically. People are now expecting interactions with businesses to be as seamless as the digital tools they use every day.

Waiting on hold, navigating complicated menus, or repeating the same information across channels no longer feels acceptable.

Second, organizations are under constant pressure to improve efficiency while maintaining service quality. Contact centers must handle large volumes of interactions while controlling costs and reducing response times.

AI-powered voice systems can help manage routine inquiries, assist agents during complex calls, and shorten resolution cycles.

Third, recent advancements in artificial intelligence have significantly improved the capabilities of voice technology. Speech recognition accuracy has improved, multilingual support has expanded, and sentiment analysis now allows systems to detect emotional cues during conversations.

For enterprises operating across diverse markets and languages, these improvements make voice technology far more practical and scalable than it was even a few years ago.

Real-world use cases across industries

AI-powered voice systems are already creating measurable impact across several industries. For example, in banking and financial services, voice AI is being used to handle account queries, transaction confirmations, and fraud alerts. These interactions require both accuracy and

compliance, making intelligent automation particularly valuable

In Healthcare sector, they are using voice assistants to streamline administrative tasks, like appointment scheduling, patient reminders, and initial triage interactions. This reduces operational pressure while allowing healthcare professionals to focus on patient care.

Also, telecom companies and service providers are deploying voice AI to help customers troubleshoot connectivity issues, understand billing details, and navigate service upgrades without waiting for a live agent.

Across these use cases, the role of AI is not to eliminate human involvement but to ensure that human expertise is applied where it matters most, handling complex or sensitive interactions that require empathy and judgment.

What technology leaders should consider before implementation

While the potential of AI voice technology is significant, successful implementation requires thoughtful planning.

One critical consideration is data governance. Voice systems capture large volumes of customer data, including potentially sensitive information. Organizations must ensure that their systems meet regulatory requirements related to data protection, privacy, and security.

Integration is equally important. AI voice platforms should connect seamlessly with enterprise systems such as customer relationship management tools, analytics platforms, and service management applications.

Without this integration, voice interactions remain isolated and fail to contribute meaningful insights to the broader organization.

Scalability is another key factor. Contact centers often experience unpredictable spikes in call volumes. A well-designed AI voice architecture must be able to scale dynamically while maintaining consistent performance and reliability.

> **AI-powered voice systems can help manage routine inquiries, assist agents during complex calls, and shorten resolution cycles**

Finally, user experience design should remain a central priority. The goal is not to replace human interaction but to create conversations that feel natural, helpful, and efficient.

Measuring the business impact of voice AI

For enterprises investing in AI voice technology, measurement is essential. Operational metrics such as average

resolution time, call deflection rates, and agent productivity can help quantify efficiency improvements.

At the same time, customer-focused metrics such as satisfaction scores and interaction quality provide insight into whether the technology is actually improving the experience.

Voice interactions also generate valuable conversational data. When analyzed effectively, this data can reveal patterns in customer behavior, highlight recurring issues, and provide early signals about emerging service challenges.

For technology leaders, these insights can play an important role in shaping both operational strategy and product development.

Looking ahead: Voice as a strategic interface

As AI capabilities continue to evolve, voice is likely to become an increasingly important interface between enterprises and their customers.

Advancements in generative AI, contextual understanding, and multilingual processing will make voice systems more adaptive and capable of handling complex conversations.

Over time, voice could become a primary gateway through which customers interact with enterprise systems, services, and workflows.

For technology leaders, the real opportunity lies not just in adopting new tools but in rethinking how voice fits into the broader customer engagement ecosystem.

Organizations that approach voice strategically, combining automation with human expertise, will be better positioned to deliver the kind of responsive, personalised experiences modern customers expect.

In the end, the goal is simple: technology should make communication easier, not harder. AI-powered voice is bringing enterprises closer to that ideal.

VIEWPOINT

WHEN AI HITS 100KW+ PER RACK, DATA CENTER DESIGN MUST BE REWRITTEN

Amit Agrawal, President, Techno Digital, points out that cities like Chennai are emerging as strategic AI infrastructure hubs, but location alone means little without purpose-built engineering within it

India's data center industry is entering a new phase, one defined not by capacity alone, but by the ability to support growing AI Factories. As AI workloads move into production, they are revealing a gap between what traditional data centers were designed for and what modern compute actually demands. And for enterprise IT leaders, this shift is already visible.

AI deployments are moving into production across analytics, automation, customer platforms, and real-time decision-making systems. What is emerging is not an incremental shift, but a structural one that challenges the very physics data centers were built on.

The breaking point of traditional design

Conventional data centers were designed for stability under moderate, predictable loads. Rack densities typically ranged between 6 to 10 kW, cooling systems were designed for distributed heat loads, and power systems were engineered for gradual variations.

Today, GPU clusters are pushing rack densities to 100 kW and beyond. More importantly, these workloads introduce sharp load variability, continuous high thermal output, and far tighter tolerance requirements across both power and cooling systems.

At these densities, inefficiencies that


“
Modern infrastructure design is moving toward tighter distribution, where control is maintained closer to the point of consumption

were once negligible begin to compound rapidly. Voltage drops across distribution paths, minor electrical losses, or airflow imbalances can directly impact performance. Thermal instability is no longer an operational inconvenience, it is becoming a constraint on compute output.

Why Chennai is emerging as an AI infrastructure hub

India's data center growth is expanding beyond traditional concentration in a single metro. Chennai has emerged as one of the most strategically positioned destinations for AI-scale data centers.

The reasons are structural.

- A strong industrial-grade power backbone capable of supporting sustained high-density loads.
- Proximity to submarine cable landing stations enables low-latency connectivity across Asia-Pacific markets.
- Policy momentum in Tamil Nadu is accelerating approvals and infrastructure readiness.

For enterprises deploying AI workloads particularly those requiring a balance between latency, power availability, and scalability, Chennai represents a compelling combination. But location alone is not enough. What matters is how infrastructure is engineered within that location.

Designing for AI from first principles

One of the key lessons in building for AI-scale environments is that legacy models cannot simply be stretched to accommodate new demands. They must be chosen and rethought as per the current requirements.

Power architecture: Stability under dynamic load

In AI environments, power is no longer just about availability. It is about behaviour

under stress. GPU-driven workloads do not follow predictable patterns. They create sudden spikes and sustain high demand, which expose weaknesses in traditional power systems. At high densities, even minor inconsistencies in voltage or distribution can affect performance.

This changes how power architecture is approached. The focus shifts toward tighter control, reduced losses, and ensuring consistent electrical conditions across the system. Stability becomes critical not at peak load alone, but across every fluctuation cycle.

Distribution strategy: Stability closer to the rack

One of the key challenges in high-density environments is maintaining voltage stability at the point of consumption. In high-density environments, voltage instability often originates closer to the rack rather than at the source. The longer the distribution path, the higher the chances of fluctuation and inefficiency. Modern infrastructure design is moving toward tighter distribution, where control is maintained closer to the point of consumption. This reduces variability, improves efficiency, and ensures that compute systems operate within stable electrical limits, particularly during peak demand.

Resilience reimagined: Continuity without disruption

Traditional resilience models focus on surviving failures. AI workloads require something more continuity without interruption. AI workloads require a different approach. Even short transitions during power events can interrupt processing and affect outcomes. Recovery is no longer sufficient. The focus is shifting toward continuity. Infrastructure needs to ensure that disruptions do not translate into workload instability. This requires

> **In AI-driven environments, performance depends on how well power, cooling, and distribution work together under continuous load**

coordinated systems where backup power, energy storage, and transition mechanisms operate seamlessly, without visible impact on compute.

Cooling architecture: Designed for thermal intensity

At densities approaching 100 kW per rack, cooling becomes a primary constraint. It is one of the first systems to reach its limit as AI workloads scale. At densities approaching and exceeding 100 kW per rack, heat is no longer a secondary consideration. It becomes a defining constraint that directly influences how efficiently compute can operate over time.

- Cooling systems must handle continuous and concentrated thermal loads, not intermittent peaks
- Approaches such as adiabatic cooling help reduce water dependency while maintaining performance under high ambient conditions
- Airflow design needs to ensure uniform distribution to avoid hotspots across densely packed racks
- Thermal systems must be built to scale with increasing rack densities without compromising stability
- Consistent thermal conditions are critical, as even small variations can impact workload performance

This ensures that thermal conditions remain stable, allowing AI workloads to operate at their intended performance levels.

Redefining the role of data centers

The industry must now move beyond viewing data centers as physical shells that house compute. In AI-driven environments, performance depends on how well power, cooling, and distribution work together under continuous load. This shifts the focus from building capacity to sustaining efficiency, stability, and control at scale. As AI adoption accelerates, infrastructure is no longer a background layer. Because at AI scale, infrastructure does not simply support compute. It is the system that defines performance.■

To know more about Techno Digital, visit: https://technodigital.in/

BNG BHARAT™ NETWORK GROUP

Our Publications

The Founder, The Educator and The Banker—three insightful magazines—delivering expert perspectives on business and finance, education, banking and IT, to empower industry leaders and professionals

To Read our Digital Editions,

Log on to: www.bharatnetworkgroup.com

For Print Editions, Contact:

info@thefoundermedia.com

editor@thefoundermedia.com

Announcing

2nd Chapter

CIO HORIZON

Where Vision Becomes Direction

2026

Meet 100+ tech leaders at the industry's most premium summit

3-5 JULY 2026 | RAMADA BY WYNDHAM, HOTEL & CONVENTION CENTER, LUCKNOW

Lucknow

An Initiative of

BNG BHARAT™ NETWORK GROUP

Concept by

Tech Disruptor media.com

For partnership opportunities, please contact

Naman Singhal
naman@thefoundermedia.in
+91 9267933240

Abhinav Chaudhary
abhinav@thefoundermedia.in
+91 8700749849