TEALIUM | DICONIUM

Agentic AI:

A new strategic roadmap for enterprise leaders

A joint point of view by Tealium and Diconium

Enter

Building the Agentic Enterprise

The rise of Agentic AI and the steps enterprise leaders must take to succeed.

Practical guidelines for enterprises transitioning from GenAI experimentation to autonomous execution

Enterprises across EMEA and beyond spent the last two years capitalising on the benefits of Generative AI, from boosting productivity to accelerating content creation and augmenting the research process. Those gains are real. But today leaders are navigating a more difficult path.

How is it possible to evolve AI from a passive tool into a credible driver of meaningful execution?

Agentic AI changes the nature of that question. It does not simply generate an answer and wait for a human to act. It can interpret context, plan across steps, use tools, remember state, interact with systems and complete defined tasks under human oversight.

This paper at a glance:

- What Agentic AI is, how it differs from Generative AI, and why that distinction matters for enterprise leaders.
- Why Agentic AI is becoming strategically important against the backdrop of operational complexity, talent shortages, rising efficiency pressure and increasing governance demands.
- How enterprises progress through the four waves of AI transformation, from inspiration to implementation, integration and bounded autonomy.
- What leaders must prioritise over the next 12 months to prepare their organisations.
- How the combination of Tealium and Diconium creates the right conditions for AI experimentation to move towards governed, scalable, autonomous execution.

Executive Summary

Enterprises stand on the threshold of a new phase of AI adoption. Two years of Generative AI experimentation have delivered meaningful productivity gains, while exposing a stark gap between intelligence and execution. GenAI can generate concepts, content and insights at scale, yet in most enterprise environments a human must still take the output, decide what to do with it, coordinate the next steps and drive work forward through disconnected systems.

Agentic AI is the next stage of enterprise automation. It moves AI beyond response and into goal-directed execution: agents can support people while also completing multi-step tasks, using tools, remembering context across steps and interacting with enterprise systems under human oversight. The shift is not simply from “chat” to “automation”. It is from individual augmentation to process transformation.

The transition, however, is rarely smooth. Many organisations push agents into production before the prerequisites for safe, reliable scale are in place. When real-time, contextualised data remains fragmented, when governance and operating-model shifts are treated as secondary, and when maturity steps are compressed or skipped, autonomy becomes difficult to trust and even harder to operationalise. Agentic systems also require security architecture, agent evaluation, rollback mechanisms, and disciplined use case scoping. Without those foundations, pilots may look promising but still fail to create enterprise-level transformation.

Unlocking Agentic AI depends on two capabilities that must work together. It starts with a trusted real-time data foundation; without unified, governed, actionable and contextualised data, agents cannot reason, decide or act safely. Tealium delivers the critical data and context layer that keeps intelligence complete, current, and compliant. Alongside that foundation, enterprises need the ability to operationalise autonomy. Diconium provides the strategic, organisational and technical capabilities to move from concept to secure, implemented AI products — including governance design, operating-model rollout and enablement across teams.

Together, Tealium and Diconium help organisations build the critical prerequisites for Agentic AI. The partnership does not replace every layer of the enterprise AI stack. It strengthens the foundations that determine

Moving from GenAI to Agentic AI

When AI stops assisting and starts acting.

Generative AI changed how people work with information. It generates, summarises, translates, drafts and synthesises at speed. It expands what individuals and teams can produce.

But the enterprise challenge is shifting from generating outputs to executing outcomes.

The distinction is not that GenAI is weak and Agentic AI is strong. The distinction is functional. Generative AI responds to a prompt. Agentic AI works towards a goal.

Agentic AI brings together four capabilities:

1. Memory and state persistence
It can retain context across steps, rather than treating every interaction as a disconnected prompt.

2. Tool use
It can interact with systems, APIs, databases, workflows and applications.

3. Goal-directed planning
It can decide which steps are required to reach an outcome, rather than following only a fixed script.

4. Execution and orchestration
It can coordinate actions across systems and workflows, often with other agents and human colleagues involved.

This also clarifies the distinction from robotic process automation. Modern RPA is no longer purely rule-based; many automation platforms already include machine learning and intelligent decisioning. The stronger distinction is between process-directed and goal-directed behaviour. RPA generally executes a prescribed path. Agentic AI receives a goal, determines a route, uses tools and adapts its next step as context changes.

Agentic AI is therefore not a replacement for Generative AI. The two are additive. GenAI expands the boundaries of what can be created; Agentic AI navigates those possibilities towards outcomes.

For enterprise leaders, the implication is clear: GenAI augments individuals; Agentic AI has the potential to transform processes. That potential is valuable only when the surrounding data, security, governance and operating model are ready.

Why Agentic AI has become a strategic priority

The opportunity is real. So is the risk of getting the sequence wrong.

Enterprises are faced with a set of interconnected pressures:

- Rising operational complexity, with more systems, channels and data sources.
- Intensifying competitive pressure, compelling organisations to pursue efficiency gains that exceed the incremental improvements unlocked by GenAI alone.
- A growing need for scalable process execution, as manual workflows can no longer keep pace with the volume and complexity of modern operations.
- Persistent talent shortages and re-skilling demands, combined with uncertainty surrounding future workforce readiness.
- Increasing governance, security and compliance demands, particularly in regulated European markets.

Market signals reinforce both the urgency and the risk of moving too quickly without the right foundations. Gartner forecasts that more than 40% of agentic AI projects will be cancelled by the end of 2027 because of escalating costs, unclear business value or inadequate risk controls. That should not be read as an argument against Agentic AI. It is an argument for better sequencing. This high failure rate is exactly what the Tealium and Diconum partnership is designed to prevent, ensuring the foundational data and operational guardrails are built before deployment. [1]

Agentic AI should not be positioned as a shortcut around enterprise transformation. It is the next step for organisations that are ready to connect data, decisioning and execution in a controlled way. Moving too early can create new risks: cascading errors, hallucinated tool calls, prompt injection, unclear accountability, excessive permissions and autonomous actions that are difficult to audit or reverse.

In a European context, this also includes regulatory readiness. The EU AI Act entered into force on 1 August 2024 and introduces a risk-based framework for AI systems. For organisations in sectors such as financial services, manufacturing, healthcare, retail and travel, Agentic AI will need to be designed with governance from the beginning, not added after deployment.[2]

The strategic priority is not autonomy at any cost. It is trusted autonomy introduced in the right sequence, for the right use cases, with the right controls.

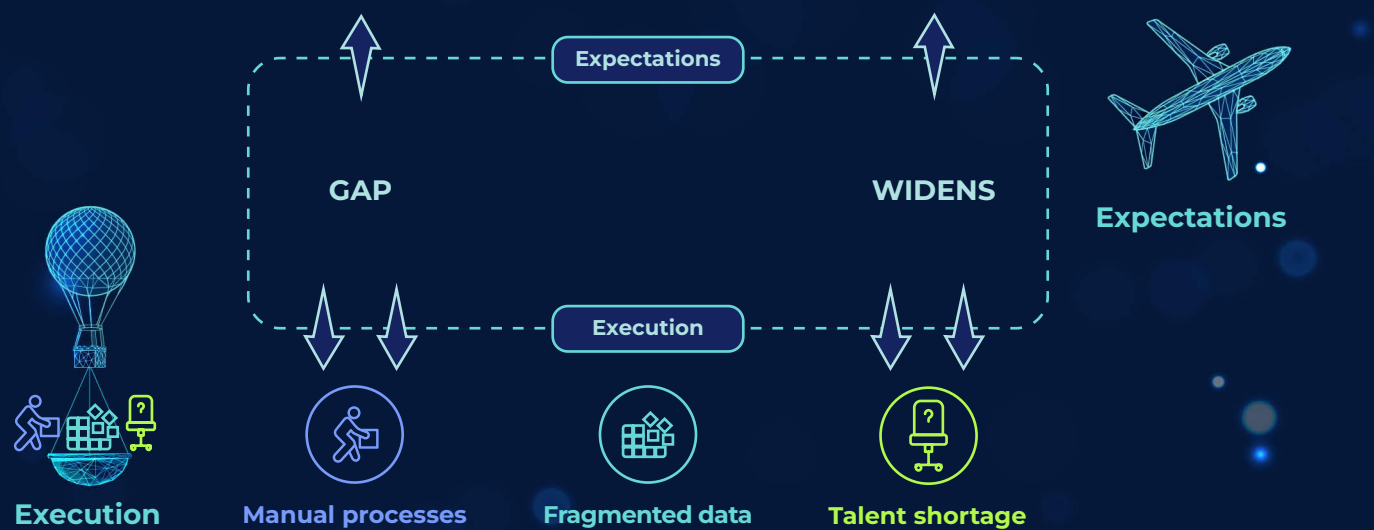

The Four Waves of Enterprise AI Transformation

A sequential maturity journey.

Enterprises will not arrive at Agentic AI through one giant leap. Instead, they are advancing through a logical maturity journey — one that Diconium has identified across different industries. Each wave progressively creates the organisational, technical and cultural foundations for the next.

Most enterprises today are operating at Wave 2 or early Wave 3. The path to bounded autonomy runs through them — not around them.

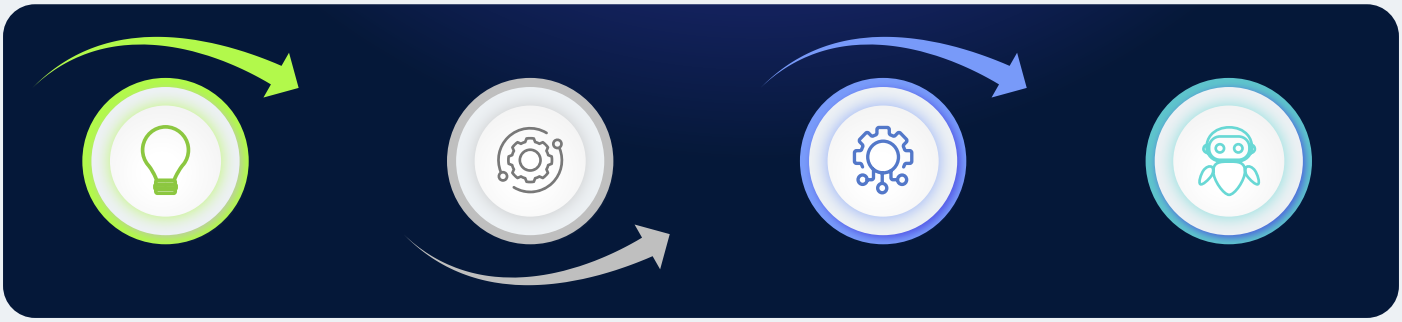

Wave 1 — Inspiration

Experimentation and early ideation

Teams explore GenAI tools at an individual level, test early use cases and build confidence in the technology. This creates momentum, but impact is usually fragmented and difficult to measure.

Wave 2 — Implementation

Productivity gains for individuals and teams

AI starts to support regular tasks, including content creation, analysis and internal workflows. Efficiency gains become visible, but impact is still confined to selected teams or functions.

Wave 3 — Integration

AI embedded into standardised workflows

AI capabilities are integrated across systems, data flows and shared processes. Governance, monitoring and ownership become operational requirements as AI moves from tool to enterprise infrastructure.

Wave 4 — Autonomy

Bounded autonomous execution

Agents begin to act within defined workflows, coordinate across systems and involve human colleagues where oversight is needed. Autonomy is introduced through bounded use cases, not unrestricted delegation.

The transition from Wave 3 to Wave 4 is the hardest step. Leaders need agent evaluation, sandboxed testing, security architecture, rollback mechanisms, auditability and accountability before bounded autonomy can scale.

Enterprises cannot skip waves. The strongest leaders invest in the next wave before they arrive there. Maturity is sequential, but not passive: each stage depends on the right operating model, governance, data foundation and skills.

The foundations for trusted autonomous execution

The data foundation and transformation capability needed for Agentic AI adoption.

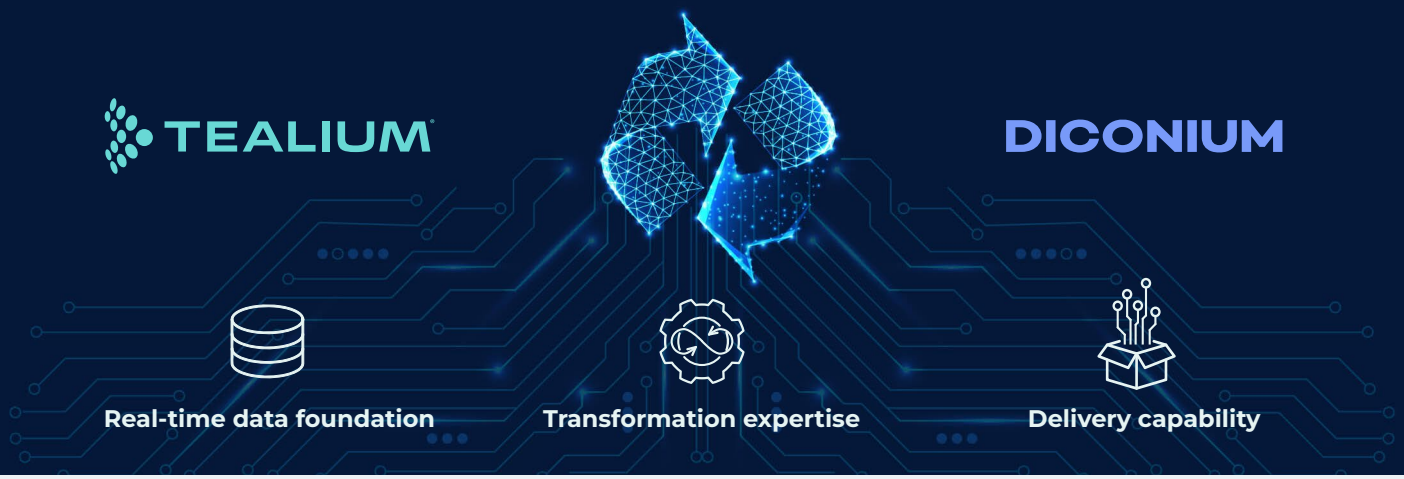

Agentic AI cannot operate without trustworthy, real-time customer and operational data. Agents need to understand what is happening, who or what is involved, what permissions apply, which systems they can use and which action is appropriate in the moment. Tealium provides the governed, real-time data layer that Agentic AI depends on to reason, decide and act responsibly. For autonomous execution to work and succeed long-term, enterprises must have data that is:

- **Unified** - Connected across relevant systems, touchpoints and customer or operational records.
- **Real-time** - Available when decisions need to be made, not after the moment has passed.
- **Governed** - Collected, processed and activated with consent, control, lineage and compliance.
- **Actionable** - Structured in a way that systems, models and agents can use.
- **Contextualised** - Described with meaningful attributes, labels and metadata that help agents interpret signals correctly.

Tealium does not simply supply data. It activates it. Tealium helps organise, enrich, govern and activate data in real time, giving AI systems the context they need to interpret signals, trigger the right actions and reduce unsafe or misaligned outcomes.

Tealium has also expanded its Model Context Protocol capability – a framework allowing AI to securely access enterprise data - to support more dynamic agentic workflows and managed access to live attributes and audiences.[4]

> *"Move the data, move it fast, enrich it, contextualise it, activate it."*
>
> **Nick Albertini, Tealium**

Technology alone does not change a business model; people and processes do. Diconium helps organisations turn AI ambition into operational reality by:

- Designing use cases around business value.
- Integrating AI into end-to-end customer journeys.
- Building the AI target operating model, including roles, governance, compliance, monitoring and ownership.
- Supporting change, enablement and adoption across teams.
- Developing and implementing AI workflows, AI assistants and value-oriented AI products.
- Applying software and data science expertise to build value-oriented AI products.

The combined proposition is not to replace every layer of the enterprise AI stack. It is to strengthen the foundations that determine whether autonomous execution becomes viable: trusted data in motion, clear operating models, governed use cases and delivery capability that can move from strategy to

Agentic AI in action

Insights on how Agentic AI operates in the real world.

Retail

Agents monitor live behaviour

The Challenge

Retailers operate in high-velocity environments with massive transaction volumes and constantly evolving customer needs and wants. Even the most technically mature digital platforms struggle to turn rich first-party data into real-time, personalised experiences that unlock measurable commercial outcomes.

The Agentic Solution

Agentic systems track live behaviour, infer intent and dynamically orchestrate journeys, tailoring content, offers and messaging in the moment. This real-time optimisation can help improve conversion and margin by reconciling customer demand with inventory and commercial priorities.

The Enabler

The missing link is alignment between use case, data foundation and operating model. Agentic AI becomes operational when those elements are embedded into clear ownership, enablement and execution.

Travel and Hospitality

Agents identify unusual signals and infer intent.

The Challenge

Traditional personalisation in travel relies on batch data and manual orchestration, making it difficult to react to real-time shifts in traveller intent, particularly in the context of high-value segments like elite business travellers.

The Agentic Solution

An elite business traveller visits a booking site on a Saturday, searches for four passengers and selects a leisure destination. An agent can recognise the shift in intent, validate it against historical profile data and adjust the experience within seconds.

The Enabler

A real-time, governed data layer gives the agent the context needed to distinguish between assumed identity and live intent. The experience can shift from business-travel defaults towards family packages, relevant offers and destination content.

Industrial Manufacturing

Agents predict service needs.

The Challenge

OEMs are frequently held back when it comes to AI adoption due to fragmented data ownership and disconnected legacy systems. This fragmentation gets in the way of a unified view of assets, customers and service needs.

The Agentic Solution

Once datasets are connected and contextualised, agents can consolidate multi-owner histories, predict service requirements and trigger proactive outreach to customers, dealers or service partners. This can evolve service from reactive to predictive, strengthening retention and unlocking new service-revenue streams.

The Enabler

The key enabler is data alignment across owners, systems and infrastructures. With a harmonised data layer and clear governance, Agentic AI can operate with the accuracy and context required for advanced, high-value industrial use cases.

Financial Services

Agents support regulatory monitoring.

The Challenge

Global financial institutions must navigate complex regulatory environments, sophisticated fraud risk and constantly evolving compliance obligations. Monitoring regulatory change across markets is labour-intensive, while customer and transaction data must remain tightly governed.

The Agentic Solution

AI agents can support regulatory monitoring and triage by surfacing relevant changes for human review. In fraud and advisory contexts, agents can help prioritise signals, surface risk and recommend next actions within defined controls.

The Enabler

Financial services requires strong data lineage, human oversight and auditable governance. Trusted data foundations and clear accountability structures help keep agentic workflows safe, explainable and aligned with regulatory expectations.

What next steps should enterprise leaders take today?

How to start moving from AI experiments to Agentic-ready operations.

AI is not waiting for organisations to feel ready. Companies do not need to start big, but they do need to start with the right sequence, the right use cases and the right controls.

TEALIUM DICONIUM

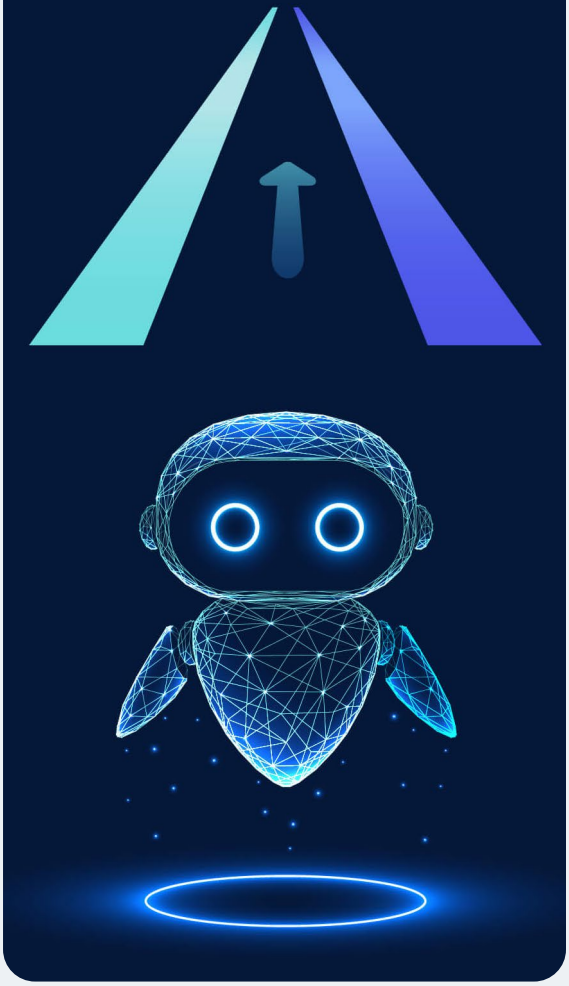

1 Audit data readiness
Wave 1 to 2

Assess whether your data is structured, consented, API-accessible and contextually enriched enough to support autonomous execution.

2 Define the value case
Wave 1 to 2

Identify where autonomy can improve cost, quality, speed, risk reduction or customer experience. Prioritise use cases that are high-frequency, data-ready and bounded.

3 Assess organisational readiness
Wave 2 to 3

Map gaps across data, governance, processes and people. A team operating at Wave 2 should not design as if it is already at Wave 4.

4 Build the control layer
Wave 2 to 3

Prioritise real-time data pipelines, orchestration, observability, agent evaluation, sandboxed testing, security architecture, access controls and rollback mechanisms.

5 Design the operating model
Wave 3 to 4

Define who owns, monitors and governs the agent. Set escalation paths, risk frameworks and change-management structures before deployment.

6 Pilot before scaling
Wave 3 to 4

Select a contained, measurable use case. Test it in a sandbox, deploy with human oversight and use predefined metrics to prove it is useful, governed and repeatable.

Governance, security and regulation

The conditions for trusted autonomy.

Agentic AI introduces new governance and security requirements because agents do not only generate outputs. They can take action. That changes the risk profile.

An incorrect AI-generated answer can mislead a user. An incorrect agentic action can trigger downstream consequences across systems, customers, partners and compliance workflows. Enterprises therefore need governance that operates at the speed of agentic execution.
At a practical level, organisations need to separate two governance questions.

The first is whether the data feeding the system is permitted, traceable and reliable. The second is how the system is allowed to act, who remains accountable and where human oversight sits.
The former is a data governance challenge. The latter is an AI governance challenge.

Data governance
Data governance determines whether agents are acting on information that is current, consented, permitted, traceable and contextually meaningful. It requires consent management, data lineage, security, compliance, filtering, enrichment and auditability.

AI governance
AI governance determines how autonomous workflows are designed, monitored and controlled. It requires ethics frameworks, human oversight, risk controls, decision accountability, ownership structures, escalation paths and operating principles.

Security architecture
Agentic AI also requires specific security controls. Prompt injection, excessive tool permissions, unclear agent identity, unmanaged access to systems and missing audit trails can all create material risk. Enterprises need identity and privilege management for agents, controlled tool registries, secure testing environments, logging, monitoring and rollback mechanisms.[4]

Regulatory readiness
For European enterprises, the EU AI Act reinforces the need for responsible AI design. Its risk-based approach gives organisations a clearer framework for understanding where AI systems create higher obligations and where stronger controls are required. Agentic AI programmes should therefore be designed with compliance in mind from the beginning, particularly in regulated sectors or use cases where automated decisions affect customers, employees or operations.[4]

Governance should not be treated as a blocker to AI progress. Done properly, it is what allows enterprises to move faster with confidence.

A partnership to power the Agentic Enterprise

Trusted data in motion. Transformation that reaches production.

Moving from the “Inspiration” phase of AI to bounded autonomy requires more than a software licence. It requires a shift in both data architecture and organisational capability.

Tealium and Diconium bring together two critical pillars of the Agentic Enterprise.

The real-time data foundation

An agent is only as useful as the data and context it can access. Tealium provides the governed, real-time and contextualised data layer that enables agents to reason from current, consented and meaningful signals.

- Contextual accuracy: Agents can act on current, first-party data rather than delayed, fragmented or incomplete signals.
- Trusted activation: Data can be enriched, labelled and activated across systems, models and channels, helping agents interpret context more accurately.
- Governed autonomy: Consent, lineage, filtering, security and compliance controls help ensure that agentic workflows remain traceable and regulated.

The strategic transformation capability

Technology alone does not change a business model; people and processes do. Diconium helps enterprises translate Agentic AI into secure, adopted and measurable execution.

- Target operating models: Designing the roles, responsibilities, governance structures and monitoring models required to manage autonomous workflows.
- Use case prioritisation: Identifying where Agentic AI can create measurable business value and where simpler automation or AI assistance is more appropriate.
- AI solution development and implementation: Developing and deploying Agentic AI workflows, AI assistants and value-oriented AI products, combining strategy, software and data science expertise.
- AI governance and enablement: Building human oversight models, decision accountability structures and enablement programmes that support responsible adoption.

Diconium brings strategy and implementation together, helping organisations move from readiness assessment to prioritised use case design and governed pilot deployment.

Together, Tealium and Diconium help enterprises answer the practical question at the heart of Agentic AI: where autonomy is valuable, how it should be governed and what must be built before it can scale.

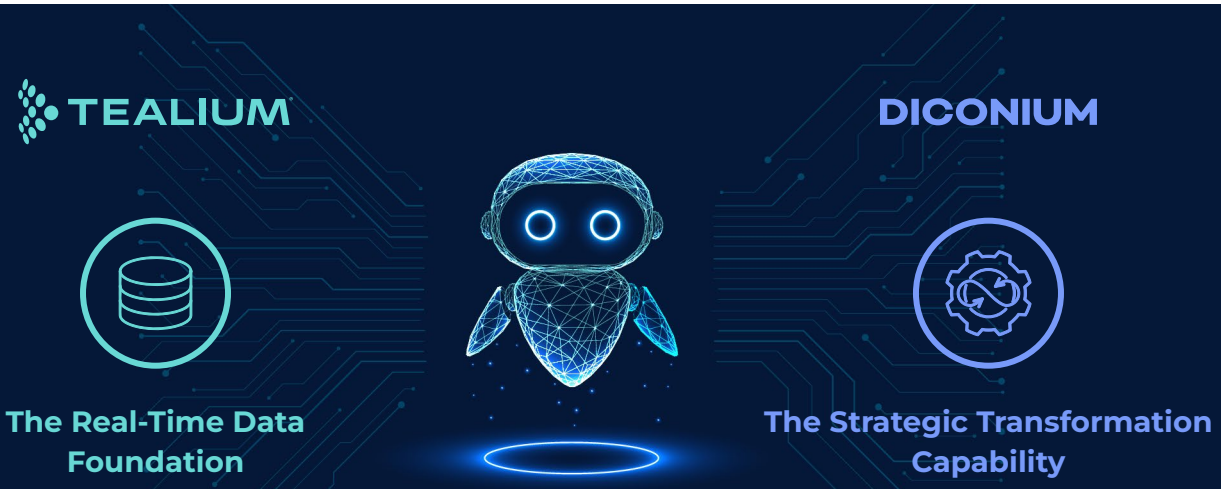

Ready to bridge the gap from experiments to execution?

The enterprises that succeed with Agentic AI will not be the ones that delegate broad autonomy first. They will be the ones that make AI a governed part of the operating model, supported by trusted data, clear accountability and carefully selected use cases.

Joint Agentic AI Readiness Assessment

A Tealium and Diconium readiness assessment helps you identify where your organisation stands today and what is blocking the path to trusted autonomous execution.

The assessment covers:

Data readiness for agents: Including structure, consent, context, accessibility and activation.

AI maturity and use-case suitability: Mapping your organisation against the four waves and identifying where autonomy can be introduced safely.

Governance, security and operating-model priorities: Clarifying the controls, roles and escalation paths required before agentic workflows move into production.

Request your Agentic AI Readiness Assessment:

Visit Diconium

Visit Tealium

You can also get in touch by emailing contact@diconium.com or by visiting Tealium.

References

1. Gartner, “Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027”, 25 June 2025.
2. European Commission, “AI Act enters into force”, 1 August 2024.
3. Tealium, “Tealium releases new AI features to accelerate real-time customer engagement”, 28 October 2025.
4. OWASP, “OWASP Top 10 for Large Language Model Applications.”


TEALIUM
DICONIUM