QUESTODES | Powered by Tech Mahindra

RAPID

Responsible, Automated, Production-grade, Integrated Development

A governed agentic SDLC acceleration proposition for enterprises that need AI-enabled software delivery to be faster, safer, traceable, and production-ready.

CEO-grade consulting brochure | Market-facing draft v3

MARKET PROPOSITION

Executive Summary

Questodes RAPID is a productised consulting proposition for organisations that want to use agentic AI to accelerate the software development lifecycle without diluting governance, security, engineering discipline or acceptance confidence.

RAPID stands for Responsible, Automated, Production-grade, Integrated Development. It combines Tech Mahindra engineering and transformation capability with the Questodes product suite to create a governed operating model for AI-augmented software modernisation.

> **Positioning statement**
>
> RAPID converts AI-assisted development from disconnected experimentation into a governed, evidence-ready delivery capability.

For CEOs Earlier business validation, a clearer route from modernisation investment to visible outcomes, and a risk posture that can be explained to the board.	**For CIOs and CTOs** A practical way to use agentic engineering while preserving architecture discipline, SDLC controls, security review, CI/CD, testing and traceability.
For transformation sponsors A contained entry point through readiness and pilot activity, followed by a repeatable model that can scale across domains.	**For delivery leaders** An operating rhythm that connects intent, backlog, agent execution, pull requests, testing, assurance evidence and acceptance decisions.

WHY NOW

The Client Problem

Enterprise leaders are under pressure to modernise legacy systems faster, release new digital capabilities sooner and reduce the cost of technology change. Agentic AI can materially improve delivery throughput, but speed alone does not solve the executive problem.

- Legacy estates often contain undocumented rules, brittle integrations, hidden data dependencies and business processes that cannot be casually regenerated.
- Regulated and semi-regulated environments require evidence, review, test coverage, secure engineering and clear accountability before changes can be accepted.
- Uncontrolled AI-assisted delivery can create code volume without corresponding confidence in quality, maintainability, security or production readiness.
- Executives need a model that accelerates delivery while making risk, evidence, ownership and decision quality more visible.

Why Traditional SDLC Modernisation Struggles

Conventional modernisation approaches remain essential, but they often struggle with the pace and scale now expected by business sponsors. Discovery takes too long, documentation is incomplete, engineering capacity is constrained, and governance evidence is assembled late in the delivery cycle.

Slow discovery Teams spend weeks interpreting screens, rules, data flows and dependencies before stakeholders can see tangible progress.	**Late assurance** Testing, security, evidence and acceptance packs are frequently treated as downstream activities rather than delivery assets created from the start.
Fragmented traceability Requirements, architecture decisions, code changes, tests, pull requests and approvals often live in disconnected tools.	**Weak executive visibility** Leaders see status summaries but not always the confidence chain behind scope, risk, quality and readiness decisions.

Why Unguided Agentic AI Is Not Enough

AI agents can analyse, generate, test and document at a pace that changes delivery economics. But without an operating model, agentic development can increase variability and assurance burden. RAPID is designed around the principle that agents accelerate production while humans govern intent, risk and acceptance.

> **The control question**
> Can every material AI-assisted output be traced to a valid intent, reviewed by accountable humans, tested through the delivery pipeline, assessed for security and used as acceptance evidence?

The RAPID Proposition

RAPID is a responsible agentic SDLC acceleration framework and governed engineering operating model. It packages discovery, agentic execution, assurance workflows, control-plane tooling and consulting delivery into a marketable offer that can be sold, piloted, scaled and defended.

Proposition layer	What it provides	Executive value
Responsible AI governance	Human oversight, role-based review, security checks, exception management and evidence readiness.	Acceleration remains explainable, auditable and aligned to enterprise risk appetite.
Agentic engineering workflows	Backlog execution, code analysis, screen replication, documentation, test generation, PR support and quality checks.	Delivery capacity increases without treating AI as an unmanaged productivity shortcut.
Production-grade delivery controls	CI/CD, maintainable architecture, security scanning, test coverage, observability hooks and acceptance workflows.	Outputs are suitable for controlled progression toward production, not only demonstration.
Integrated development evidence	Connected intent, requirements, stories, code, PRs, tests, exceptions, decisions and acceptance artefacts.	Sponsors can see the confidence chain behind delivery progress.

The RAPID Acronym

Responsible Governed AI usage, human oversight, ethical	**Automated** Agentic workflows for analysis, backlog execution,

guardrails, risk management, security review and assurance checkpoints.	code generation, documentation, test generation, evidence creation and continuous quality checks.
Production-grade Secure code, maintainable architecture, CI/CD, observability, role-based controls, quality gates and enterprise deployment discipline.	**Integrated Development** A connected SDLC across architecture, requirements, backlog, code, PRs, tests, security, evidence and acceptance workflows.

How RAPID Works as an Operating Model

RAPID is organised as a governed flow from business intent to production-grade delivery. It is not a tool overlay. It is a delivery model that defines how agent work is scoped, executed, reviewed, tested, traced and accepted.

Stage	Core activities	Governance artefact
1. Discover	Legacy screens, code, data behaviour, business rules, integrations, risks and domain boundaries.	Discovery evidence, system map and modernisation backlog.
2. Structure	Target domains, bounded contexts, architecture decisions, service boundaries and delivery increments.	Reference architecture, ADR log and scoped execution plan.
3. Execute	AI-assisted code, tests, documentation, API/service layers, UI replication and environment artefacts.	Agent execution record, changed files, PR links and generated verification assets.
4. Govern	Qontrol Intent review, specification alignment, exception handling, risk flags and PR assurance.	Governance status, traceability record and exception ledger.
5. Assure	Human review, QA, security scanning, sandbox validation, UAT support and acceptance evidence.	Evidence pack, readiness decision and acceptance recommendation.

The Role of Qontrol Intent

Qontrol Intent is positioned as the governance and assurance layer for agentic development execution. It provides post-execution governance so developers can move at pace while material outputs are still brought under review, traceability and acceptance control before merge or evidence use.

- Ingests agent execution data from Augment Intent, CLI push, GitHub PR events or manual sources.
- Links agent intent and specification sections to changed files, pull requests, verifier output and evidence contributions.
- Supports technical review, assurance sign-off, exception handling and risk flagging when outputs drift from specification or merge before governance.
- Provides leadership dashboards for governed rate, open exceptions, stale executions, requirement coverage and evidence readiness.

> **Assurance principle**
> No material agent output should become acceptance evidence without a governed execution record, review status, test alignment and sign-off path.

Questodes Control-Plane Ecosystem

RAPID sits within a broader Questodes operating model for governed transformation. The suite provides control planes for requirements and intent, agentic execution governance, traceability, evidence management, delivery assurance, stakeholder workflows and executive reporting.

Suite component	Role in RAPID	Control-plane contribution
Qontrol	Transformation governance and acceptance control.	SOW obligations, risks, decisions, evidence packs, programme health and executive reporting.
Qontrol Intent	Agentic execution governance.	Intent-to-spec-to-execution traceability, PR assurance, exceptions and evidence readiness.
Qlarity / Qlarion	Business analysis intelligence.	Stakeholder inputs, requirements, review workflows, quality scoring, traceability and reports.
Qonvey	Discovery intelligence.	Conversation/document capture, evidence discipline, insight validation and report-ready exports.
Qatalyst	Growth and strategy control plane.	Opportunity, proposition, evidence, strategy and market signal governance for commercial adoption.
Qudos	Adoption and recognition layer.	Reinforces new behaviours, approvals and delivery momentum where change adoption matters.

Delivery Model and Service Offerings

Offer	Typical duration	Outcome	Best fit
RAPID Readiness	2-4 weeks	Opportunity case, risk profile, agent-readiness view, target architecture and pilot roadmap.	Executive decision-makers.
RAPID Pilot	6-10 weeks	Contained domain build with Qontrol Intent governance, tests, security checks and business sandbox.	CIO/CTO sponsors.
RAPID Factory	12+ weeks	Scaled agentic delivery pod integrated with client SDLC, assurance and acceptance governance.	Complex modernisation programmes.
RAPID Control Plane	Parallel workstream	Questodes governance layer connecting intent, requirements, PRs, tests, evidence and acceptance packs.	Assurance and programme leaders.
RAPID Scale Playbook	Post-pilot	Reusable standards, agent catalogue, templates, delivery controls and enterprise adoption model.	Enterprise rollout sponsors.

Governance, Assurance and Trust by Design

Scoped execution Agent work is grounded in agreed requirements, specification sections, backlog items or bounded modernisation objectives.	**Human review** SMEs, engineers and assurance leads review material output before it is merged, accepted or used as formal evidence.
Security and quality gates Static checks, dependency review, vulnerability scanning, unit tests and build controls are part of the operating model.	**Evidence readiness** Delivery evidence is created alongside work, not reconstructed at the end of the programme.
Traceable decisions ADRs, exceptions, waivers and risk acceptances remain linked to the work they affect.	**Executive visibility** Leaders can see governed rate, open exceptions, coverage, evidence readiness and readiness decisions.

Early Proof Points and Practical Evidence

Early AI-assisted development work indicates practical feasibility and significant acceleration potential. These proof points should be positioned as directional evidence from early application, not as universal promises for every estate.

Signal	What it indicates	How to state externally
200+ screens accelerated	Agentic delivery can move beyond small demos into substantial business-application reconstruction.	Early work demonstrated rapid screen replication into a modern stack, subject to estate complexity.
3,000+ unit tests generated and executed	AI assistance can create verification assets as part of build activity.	Generated tests require engineering review and build-control validation before reliance.
10 domain service areas structured	Legacy scope can be shaped into modernisation boundaries with human validation.	Domain structuring supports maintainability and migration planning.
Security and CVE checks applied	AI-assisted engineering can be paired with objective security review.	Controls improve assurance posture but do not replace formal client security review.
Movement toward governed production delivery	The model can shift from experimentation to controlled SDLC adoption.	RAPID should begin with readiness, then pilot, then scale.

Example Client Use Cases

Legacy application modernisation Replicate existing screens and business flows into a modern architecture while preserving traceability to source behaviour.	**Regulated workflow digitisation** Accelerate backlog execution while maintaining review, evidence, security and acceptance controls.

Application estate rationalisation Use discovery and control-plane evidence to identify duplication, risk, reuse opportunities and modernisation sequencing.	**AI-augmented delivery factory** Stand up a governed engineering pod with approved agents, PR assurance, CI/CD and acceptance evidence creation.
Pre-UAT business familiarisation Create sandbox environments where SMEs and support teams can validate behaviour before formal UAT.	**Enterprise SDLC adoption** Convert isolated AI productivity gains into repeatable standards, controls and executive reporting.

Implementation Pathway

Phase	Leadership question	Output
1. Frame	Where can agentic acceleration create value without creating unacceptable risk?	Opportunity thesis, target domain and decision criteria.
2. Assess	Is the estate, SDLC and governance environment ready for RAPID?	Readiness score, risk profile and pilot backlog.
3. Pilot	Can a contained domain produce working output with defensible controls?	Validated proof, evidence chain and lessons learned.
4. Industrialise	What standards, tooling and roles are needed to repeat the model?	RAPID factory model, Qontrol Intent workflow and engineering playbook.
5. Scale	How should the capability extend across portfolios, vendors and delivery teams?	Enterprise control plane, reporting cadence and adoption roadmap.

Commercial Value Levers

- Earlier visibility of working software and user experience, reducing the time spent debating abstract requirements.
- Higher delivery throughput by shifting repeatable analysis, documentation, test generation and code-support activity to agentic workflows.
- Reduced assurance drag through evidence creation, traceability and review records generated as part of delivery.
- Lower delivery risk by pairing acceleration with architecture discipline, human review, security controls and clear exception handling.
- More scalable commercial model through packaged diagnostics, pilots, factories and control-plane services.

Why Tech Mahindra and Questodes

RAPID is designed to combine boutique product proposition intensity with credible global technology consulting depth. Questodes brings the control-plane product family, branded operating model and executive narrative. Tech Mahindra brings sector experience, delivery scale, engineering capability, security discipline and transformation reach.

Questodes-led	**Tech Mahindra powered**

Ownable product language, control-plane thinking, executive storytelling and governance-by-design product concepts.	Engineering scale, modernisation capability, delivery assurance, industry credibility and enterprise execution depth.
Commercially defendable Packaged services, clear entry points, repeatable delivery model and controls suitable for senior-client scrutiny.	**Client-ready posture** Ambitious enough to be distinctive, controlled enough to be credible, and practical enough to start with a contained pilot.

Call to Action

Recommended executive move

Launch RAPID with a 2-4 week readiness diagnostic for a selected modernisation domain, followed by a 6-10 week governed pilot using Qontrol Intent as the assurance layer.

The first engagement should validate value, risk, estate suitability, governance requirements, delivery controls and executive sponsorship. The second should prove a repeatable model: scoped agentic execution, governed PRs, generated tests, security checks, evidence packs, business familiarisation and a scale recommendation.