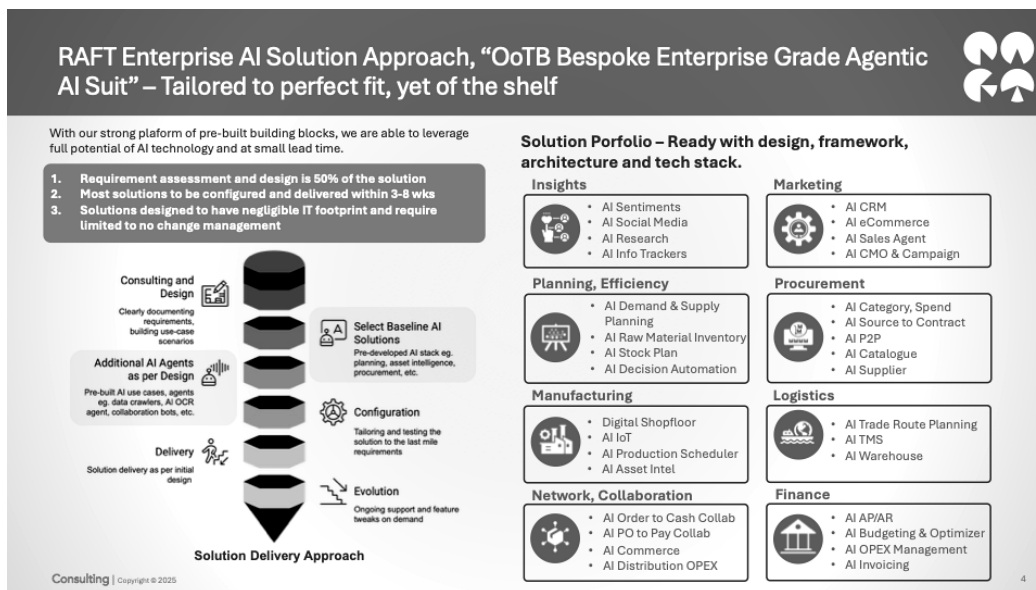


Enterprise-grade AI, Configured Fast

RAFT's engagement model is designed for customers who need enterprise-grade AI outcomes without the timeline and organizational weight of a traditional enterprise SaaS program. The model is out-of-the-box bespoke: pre-built building blocks, tailored to the exact operating requirement.



3-8 weeks: typical focused delivery window

- Faster time to first value
- Lower change management burden
- Reduced IT dependency
- Configurable path to scale

Assessment and design are treated as half the solution

RAFT begins with a focused requirement assessment because the best AI solution is rarely the one with the most features. It is the one that fits the customer's process, data maturity, governance needs, system landscape and adoption capacity. The design phase documents the use case, success metrics, human-in-the-loop points, data sources, exception paths, integration needs and target operating model before configuration starts.

This matters because many AI projects fail through ambiguity. Business users expect one outcome, technology teams build another, and leadership cannot see value clearly enough to scale. RAFT reduces that risk by making design explicit and measurable. The customer knows what the agent will do, what it will not do, how it will be evaluated, and how the process changes after deployment.

Why delivery can be faster than enterprise SaaS

Enterprise SaaS systems are powerful, but implementation timelines are often long because they involve broad process redesign, configuration across multiple modules, integration development, master data cleanup, testing cycles, security design, end-user training and change management. That is appropriate when the customer is replacing the system of record. It is often too heavy when the immediate need is a targeted AI capability.

RAFT shortens timelines by using a portfolio of ready solution patterns across insights, marketing, planning, procurement, manufacturing, logistics, finance and network collaboration. Instead of starting from a blank page, RAFT configures pre-built frameworks, architecture, agent patterns and technical components around the customer's specific requirement. For focused use cases, this can move delivery into a three-to-eight-week window, with negligible IT footprint and limited change management.

Customer benefit

The customer gets speed without losing enterprise discipline. Pre-built components reduce engineering time; bespoke configuration preserves process fit. Limited IT footprint means the solution can often work around existing platforms rather than forcing a major release cycle. Smaller change impact means business teams can adopt the capability faster, learn from real usage and scale the model where value is proven.

RAFT's model is especially useful for customers that want to move from AI experimentation to operational AI. It creates a structured path from assessment to baseline solution selection, configuration, delivery and evolution. The result is faster value realization, clearer accountability and an AI portfolio that can keep improving after the first release.