# OpenText AI Service Tiers

OpenText AI engineers work inside your team to take agents from first proof to always-on operation




## Benefits

- Prove agentic value on your own content in four weeks
- Go live with production agents, not demos
- Scale to a suite of agents with a dedicated AI team
- Adapt as you learn, without new scope or change orders

Most organizations know that AI can improve their business. Far fewer have an agent running in production, on their own information, inside their own controls. The gap is rarely the model. It is engineering capacity that understands your content, your permissions and governance, and that stays with you long enough to get something live.

OpenText AI Service Tiers give you that capacity. Each tier places forward deployed engineers, AI specialists, inside your operation for a fixed term at a fixed price. They run discovery with your business, build agents against your own information, evaluate them with your people and take them through your own release and governance processes. You choose the tier by how far you want to go, not by how many hours you want to buy.

## Prove value on your own content in four weeks

OpenText Agent Prototype is a four-week engagement with a dedicated forward deployed engineer. It opens with a discovery workshop to select a use case with real volume and a measurable outcome, then moves straight into building a working agent on content from your own estate. You finish with a demonstrated agent, the evaluation evidence behind it and a clear view of what production would take. The work runs in a suitable non-production environment of yours, or in a demonstration environment OpenText provides for the four weeks.

**OpenText Advisory Services ›**

Modernize, automate, and de-risk operations using AI to turn governed data into business outcomes.

## Go live with production agents, not demos

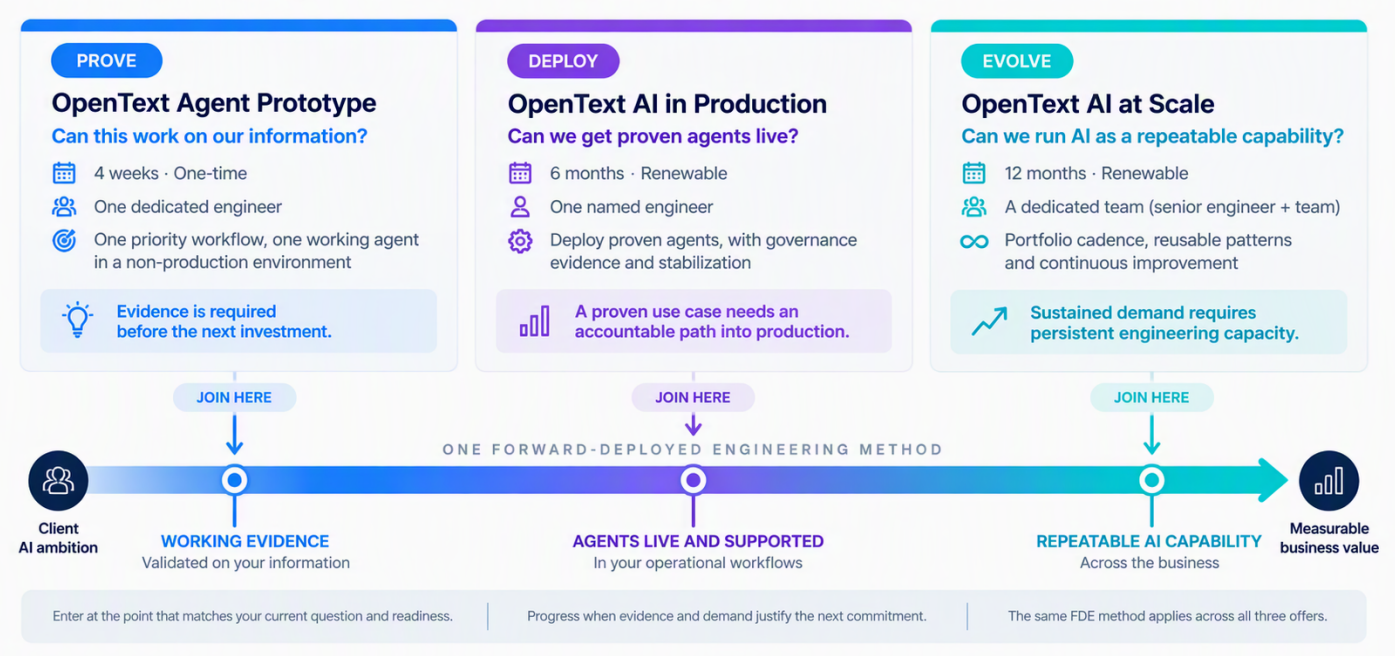
OpenText AI in Production is a six-month subscription with a dedicated forward deployed engineer working alongside your team. Across the term you run two discovery and design cycles and advance up to two agents into your production environment, at your pace and through your own change, security and release processes. Each agent is followed by 30 days of hypercare, so the engineer who built it is still with you when real users meet it for the first time.

## Scale to eight agents with a dedicated AI team

OpenText AI at Scale is a twelve-month subscription with a team rather than an individual: a senior forward deployed engineer guiding the program and two dedicated engineers building it. It supports up to eight agents in production across multiple business areas, with hypercare running for the whole term rather than a window per agent. It suits organizations that have stopped asking whether AI works and started asking how to run it as a capability, with an agent portfolio, reusable components and governance evidence produced as the work is done.

## Adapt as you learn, without new scope or change orders

AI programs change direction, usually because the first agent teaches you something you did not know when you scoped it. Every tier is bought as committed engineering capacity and an agreed rhythm of work, not as a fixed list of deliverables priced in days. When the priority moves, the next cycle moves with it. There is nothing to renegotiate, no change request to raise and no clock to argue over.



## Three tiers at a glance

| | OpenText Agent Prototype | OpenText AI in Production | OpenText AI at Scale |
|---|---|---|---|
| **In one line** | Prove a use case on your own content | Get your first agents live and supported | Run AI as a standing capability |
| **Term** | 4 weeks, one time | 6 months, renewable | 12 months, renewable |
| **Your OpenText team** | One dedicated forward deployed engineer | One dedicated forward deployed engineer | A senior forward deployed engineer and two dedicated engineers |
| **Discovery and design workshops** | 1 | 2 | 8 |
| **Agent build cycles** | 1 | 2 | 8 |
| **Agents advanced into production** | Not included. Demonstration environment only | Up to 2 | Up to 8 |
| **Hypercare after go-live** | Not applicable | 30 days per agent | For the full term |
| **Where the work runs** | Your non-production environment, or a demonstration environment | Your OpenText environment | Your OpenText environment |
| **Governance** | Closing review with your sponsor | Regular progress review with your sponsor | Regular progress reviews and executive reviews across the term |

## Have full confidence in OpenText

Your agents are only as good as the information they read and the controls around it. OpenText already manages that information: the content, the permission model, the retention rules, and the audit trail. Our engineers work with it every day, and they bring direct access to OpenText product engineering, cloud, and support. That is the difference between an AI pilot and an AI capability you can defend to your board, your auditors and your regulators.

**4,000+**
World's largest pool of information management experts

**50,000+**
Engagements over more than 25 years

**2,000+**
Managed Services customers

**One team**
Products, solutions, services, support

## Learn more

OpenText Professional Services ›

OpenText Aviator AI services ›

OpenText Aviator AI blogs ›

OpenText Advisory Services ›

## Before you start, or alongside

Not every organization is ready to put an engineer to work on day one, and not every requirement belongs inside a tier. **The OpenText AI Data Readiness Workshop** is a two-day workshop that assesses your data and leaves you with a readiness report and a roadmap. **OpenText Advisory Services** builds the AI governance, policy and maturity framework for organizations that need it in place first. **OpenText Consulting Services** delivers the platform implementation, upgrade and migration work, which sits outside the AI Service Tiers and is delivered as its own project.

## Ready to start?

To talk to an OpenText expert about the AI Service Tiers, please email profservices@opentext.com or visit opentext.com/services.

Copyright © 2026 Open Text • 10.26 | 242-000158-001

opentext™