opentext™

WHITE PAPER

Is your service management ready for AI?

A practical guide to creating AI-powered service management grounded in privacy and enterprise knowledge

Contents

AI in all forms, including predictive, generative, and agentic, is reshaping service management with new levels of efficiency, responsiveness, and personalization. But successful adoption demands trust, knowledge readiness, and strategic execution. This white paper explores how organizations can prepare their service management environments for AI with OpenText™ Service Management and OpenText™ Service Management Aviator™.

OpenText Service Management is a comprehensive IT service management (ITSM), enterprise service management (ESM), and IT asset management (ITAM) solution, designed to simplify and automate service delivery across the enterprise. Enhanced by OpenText Service Management Aviator—a generative AI service—the solution delivers smarter service experiences through AI-driven workflows and contextually relevant assistance.

Understanding Aviator

OpenText Service Management Aviator meets the rigorous demands of enterprise environments. It keeps data, users, and operations secure and compliant while delivering relevant, accurate answers and solutions.

Our approach to AI is grounded in two core pillars:

1. **Privacy and security.** Is the AI service private? Who can access your data? What internal controls are in place to safeguard it?
2. **Enterprise intelligence.** Does the AI service leverage your organization's own information? Does it provide context-aware responses? Does it know which sources to draw from when generating answers?

Let's take a closer look at each of these pillars.

A private, secure LLM

OpenText provides both private and public large language model (LLM) options to meet diverse enterprise needs. Choosing the private LLM option enables Aviator to run within a secure, controlled environment that is fully hosted and operated by OpenText.

Aviator accesses only the data that is indexed on your OpenText Service Management tenant, which includes enterprise knowledge you've chosen to share with OpenText Service Management (for example, internal websites and documents from your content management systems).

All data exchanged between Aviator and OpenText Service Management is encrypted—both at rest and in transit.

High-level overview of OpenText Service Management with Aviator.

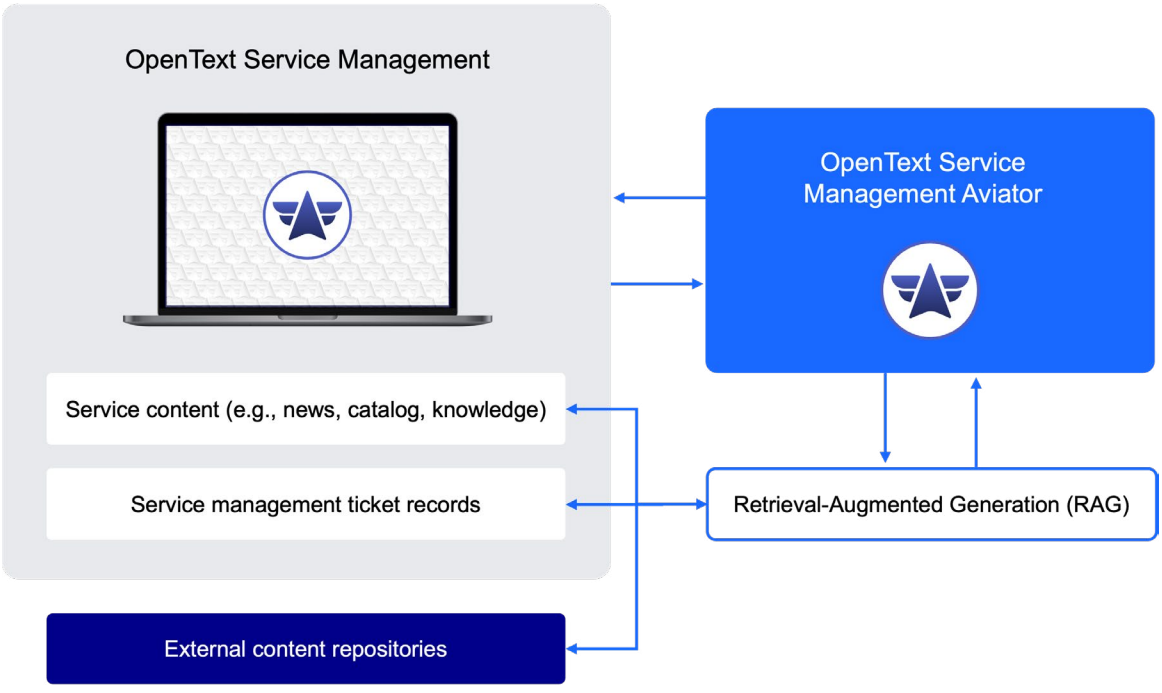

Protection of sensitive information

Aviator's access to data is governed by user entitlements and data domain permissions configured in OpenText Service Management, giving you precise control over the protection of sensitive and proprietary information. Permissions can be assigned at the user, group, or role level. For example, an IT agent may not have visibility into the same data as an HR agent.

During human–Aviator interactions, Aviator's access to information is always limited by the permissions of the human user. If a user or agent is restricted from viewing certain data, Aviator will be similarly restricted.

For agentic flows, business rules can be configured to operate under the access permissions of specific users. You can choose to run Aviator actions under the restrictions of the logged-in user, the ticket submitter, or a designated user group with customized permissions.

OpenText Service Management also includes out-of-the-box Aviator capabilities that automatically identify and flag sensitive information as private. When flagged, Aviator moves the ticket to a restricted group, ensuring that unauthorized users cannot view the data.

Most importantly, Aviator does not train on your data.

Enterprise intelligence

If Aviator does not train on your data, how does it understand your business? And how does it know which sources of enterprise knowledge to draw on? The answer lies in retrieval-augmented generation (RAG) and the configuration of Aviator prompt models.

RAG

Aviator uses RAG to improve the relevance and accuracy of its responses through two steps: retrieval and generation. RAG enables Aviator to access knowledge from your organization's trusted content sources. Aviator then uses that knowledge to generate a response—grounding its answers in your actual policies, procedures, and documentation rather than in generic internet data.

Aviator models

An Aviator model defines the prompt template used in user–agent interactions and AI-enriched workflows. Within each model, you can specify which RAG sources Aviator should search when generating responses.

For example, in an incident management assistance model, you can configure Aviator to search only resolved and relevant incident records. OpenText Service Management includes out-of-the-box Aviator models that are fully customizable, allowing you to tailor them to your specific business needs.

Unlock AI success with impeccable knowledge

Should you optimize your enterprise knowledge before using AI?

Absolutely.

AI is only as good as the information it can access, which is why preparing your sources of knowledge is a critical step in enabling Aviator to deliver precise, context-aware responses.

This preparation involves identifying and curating accurate, relevant knowledge required for RAG. Sources may include internal content native to OpenText Service Management and external content from repositories, both of which are governed by user entitlements and data domain permissions.

Knowledge management health check

Before integrating AI into your service management, assess the health of your knowledge environment:

- What are the primary sources of your enterprise knowledge?
- Do all sources meet a consistent standard of quality?
- Are service and support tickets fully documented at closure?
- Are knowledge articles actively used in service management, and are they current?
- Where do gaps or redundancies exist in your knowledge base?
- Do you have designated knowledge managers who actively curate enterprise knowledge?

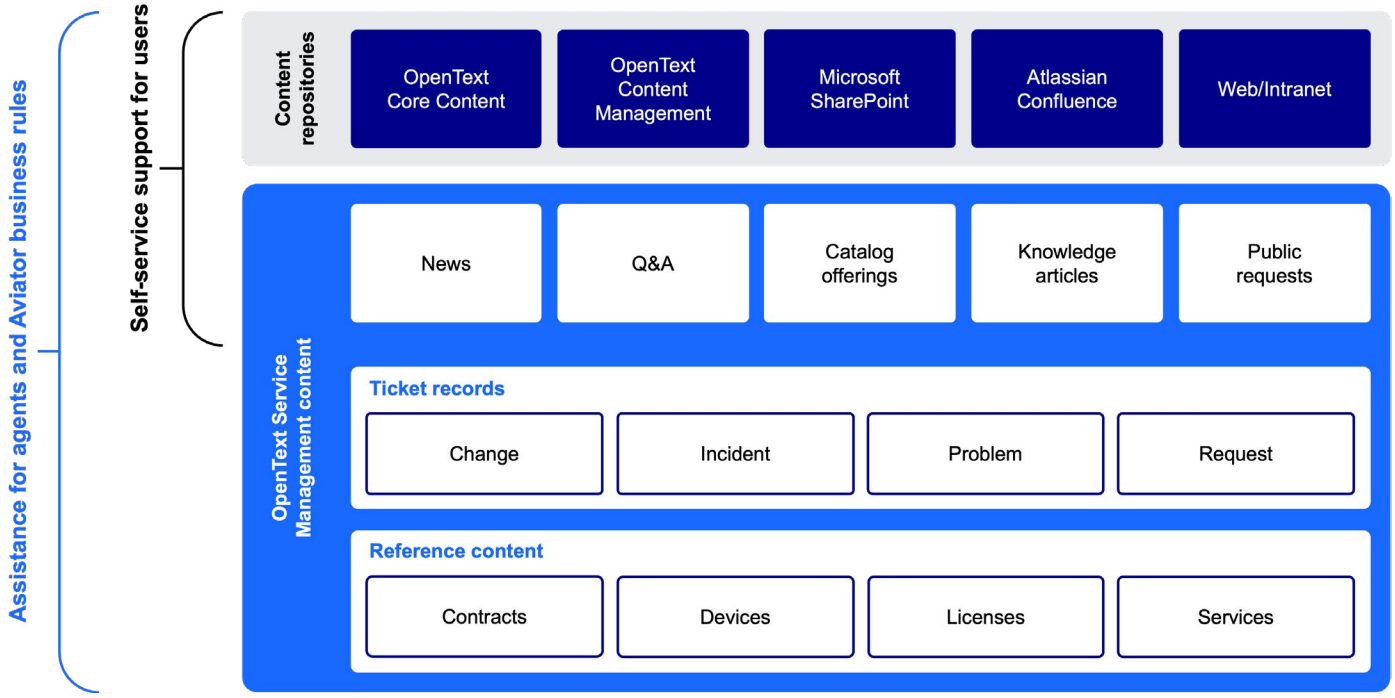

Examples of internal and external content sources used by Aviator in OpenText Service Management.

Knowledge sources for Aviator

Aviator draws on multiple knowledge sources within OpenText Service Management, which fall into two categories: external and internal content.

External content

This content lives in external repositories such as intranet sites and content management systems, including Microsoft SharePoint, Atlassian Confluence, OpenText™ Core Content, and OpenText™ Content Management (xECM).

It is indexed using OpenText™ Knowledge Discovery (IDOL), an advanced data indexing and analytics platform built into the core of OpenText Service Management. Indexing is enabled by configuring IDOL connectors to your OpenText Service Management tenant.

Internal content

This content consists of native content and records, such as:

- News, Q&A, catalog offerings, and knowledge articles.
- Ticket records across request, incident, problem, change management, and other ITIL processes.
- Reference content related to contracts, devices, licenses, and services.

Best practices for knowledge management with Aviator

Consider these key guidelines to enhance the quality and relevance of your enterprise content.

External content

- Start by mapping enterprise content sources to your AI-supported use cases. For instance, if Aviator handles HR inquiries, your HR intranet pages serve as a key content source.

Tips:

- Encourage users to provide feedback by giving a thumbs up or down and commenting on Aviator's responses within the self-service portal. This input helps administrators and knowledge managers identify opportunities for improvement—such as addressing knowledge gaps, updating outdated content, or refining Aviator's prompt templates. Feedback and ratings apply to responses generated from both internal and external content sources.
- Consider structuring your content with clear headings, subheadings, lists, and smooth transitions between paragraphs to optimize AI-driven content retrieval and generation.

- Index only the content that matters. Since external content repositories can be large, focus on indexing only the data that is truly relevant to Aviator's use.
- Ensure content is current and accurate. Outdated or incorrect information will generate poor AI responses.

News, Q&A, catalog offerings, and knowledge articles

- Maintain up-to-date knowledge articles. Regularly review existing content to eliminate redundancies and create new articles that address gaps in your knowledge base.

Tips:

- Use Hot Topic Analytics to identify recurring service and support requests, and create targeted knowledge articles to close knowledge gaps.
- Leverage Aviator to generate knowledge articles. Based on ticket documentation, Aviator can draft detailed articles that agents and knowledge managers can edit and refine before publishing.

- Ensure catalog offerings include clear, well-written descriptions as Aviator uses these descriptions to suggest solutions.
- Maintain timely and consistent news related to service status, disruptions, and availability.

Ticket records

- Specify which record types (for example, changes, incidents, news, and service definitions) Aviator should reference in the Aviator agent interaction model. Enable only record types that contain reliable documentation.
- Ensure incident and request tickets are thoroughly documented with resolution details before they are closed.

Tips:

- Train support and service teams to maintain detailed, consistent ticket resolution.
- Frame performance goals around AI to encourage thorough documentation practices.

- Document root causes, known errors, and solutions in detail to support thorough problem management. Aviator can leverage this data when analyzing related incidents.
- Clearly document change descriptions and implementation details to help Aviator assess change risks and support planning for future changes.

Reference content

- Ensure records related to Service Asset and Configuration Management (SACM) entities (such as devices, actual services, and service definitions) include comprehensive and descriptive details. Aviator can leverage this information to support related queries.

  For example, Aviator can better assess the risk of a change to a service if it knows the business functions the service supports, its criticality to operations, and peak usage periods, all of which can be documented in the service CI description.

Adopt a phased rollout and human-in-the-loop approach

Because Aviator's responses rely on the quality of the content it references, begin your rollout with high-quality, trusted knowledge sources. Then expand as additional sources become AI-ready.

Additionally, strengthen human-in-the-loop processes with continuous reviews and pilot programs before scaling AI across the organization.

Expand your knowledge sources

You can enhance Aviator's capabilities by integrating additional enterprise data sources, once their quality and completeness have been validated. These sources may include:

- **Self-service Q&A forums:** User-generated discussions and answers can provide valuable insights that Aviator can reference when responding to similar queries.
- **Public requests:** Aviator can leverage requests that users make public, allowing others to view and follow them. This source is useful for issues affecting shared locations or teams—for example, printer problems in a specific office.
- **Reference content:** Documents like contracts, licenses, and service catalogs can be used by Aviator to support agent interactions and provide more informed responses.

Keep humans in the loop

Human feedback plays a critical role in continuously refining and expanding Aviator's knowledge. Involve your employees in key AI processes—such as rating the accuracy and relevancy of Aviator's responses.

Knowledge managers, too, play an important role in curating high-quality enterprise knowledge. Consider establishing domain-specific knowledge management roles to maintain the accuracy and relevance of your enterprise knowledge. These experts can identify knowledge gaps, improve content structure, and ensure that enterprise knowledge consistently meets user expectations and adapts to evolving AI-assisted use cases.

To complement these human-in-the-loop processes, OpenText Service Management's role-based access control can be used to deploy Aviator to select pilot groups. These groups can test the validity and acceptance of responses before expanding access to broader audiences.

Start your AI journey

We'd love to show you how Aviator works and help you succeed with it.

Reach out to your OpenText account manager or request a demo to connect with a local solutions consultant.

Be sure to ask about our **AI Readiness Workshop**—a collaborative engagement with our Professional Services team designed to assess your knowledge maturity and provide tailored recommendations to optimize your service management for AI-driven experiences

opentext™