

# HEALTHCARE AI AGENTS

Amna khan  
GHUNCHAS

## Pricing, Billing & Usage Management

Understanding pricing and usage management is essential for organizations deploying Healthcare Agent Service at scale. This module explains how billing works, how usage is measured, and how organizations can control costs effectively while maintaining performance and safety.

### Billing Model Overview

Healthcare Agent Service follows an action-based billing model, meaning organizations are charged based on the actions performed by the agent rather than a fixed subscription fee. Actions may include processing user queries, accessing credible sources, executing triage workflows, or invoking plugins. This model ensures that organizations only pay for what they actually use. The Pay-as-you-go model is available under the C1 tier, which is designed for production use. This tier allows organizations to scale their usage dynamically as demand increases, making it suitable for hospitals, healthcare platforms, and enterprises. In addition, a Free (F0) tier is available for learning, experimentation, and early development. The free tier enables teams to explore the platform's capabilities with limited usage, making it ideal for training and proof-of-concept work.

### Action Consumption Metrics

Usage is measured through specific consumption metrics that track how the system is used. Customer source usage refers to actions where the agent accesses private organizational data, such as hospital documents or internal knowledge bases. These actions typically involve vector search and grounded response generation.

Credible source usage tracks interactions where the agent relies on trusted medical databases, such as government or research institutions, to answer user queries. Triage session costs are calculated based on the use of symptom checking and triage workflows, which rely on structured medical logic. By tracking these metrics, organizations gain clear visibility into how their usage translates into billing.

### Cost Optimization Strategies

Cost optimization begins with efficient scenario design. Well-structured scenarios reduce unnecessary steps and avoid repeated actions, which helps minimize consumption. Thoughtful plugin selection also plays a key role, as invoking complex or external plugins may increase action usage. Monitoring usage trends through analytics allows organizations to identify high-cost patterns and optimize them proactively. By reviewing usage reports regularly, teams can adjust workflows, refine scenarios, and ensure that costs remain predictable and controlled without compromising user experience.

## Responsible AI & Medical Disclaimer

Responsible AI is a foundational principle of Healthcare Agent Service. This module explains how ethical AI practices are applied and clarifies the legal and medical boundaries of AI-generated healthcare information.

### Responsible AI Principles

Transparency ensures that users clearly understand when they are interacting with an AI system. Healthcare Agent Service openly communicates that responses are AI-generated and provides evidence and citations to support its answers. This openness builds user trust and prevents misunderstanding. Accountability ensures that organizations deploying the agent remain responsible for how it is used. While the platform provides safeguards and controls, organizations must configure and monitor the system appropriately. Safety-first design prioritizes patient well-being by restricting unsafe behaviors, validating medical content, and preventing harmful outputs. These principles ensure that AI is used ethically and responsibly in sensitive healthcare environments.

### Medical Device Disclaimer

Healthcare Agent Service is explicitly designed as a non-diagnostic system. It does not provide medical diagnoses, prescribe treatments, or replace healthcare professionals. Instead, it offers informational and educational support to assist users in understanding health-related topics. Legal responsibility remains with the deploying organization to ensure that users are informed of these limitations. Customers are required to present appropriate disclaimers and ensure that the agent is used within approved use cases. This clarity protects both users and organizations while maintaining regulatory compliance.

## Updates, Monitoring & Future Enhancements

Healthcare Agent Service is an evolving platform that continuously improves through updates, monitoring, and user feedback. This module highlights how organizations can stay informed and maintain long-term system effectiveness.

### Service Updates & Announcements

Microsoft regularly publishes updates and announcements through official Tech Community resources. These updates may include new features, performance improvements, enhanced safeguards, or expanded integrations. Tracking the product roadmap helps organizations plan future enhancements and align their deployments with upcoming capabilities.

Staying informed about service updates ensures that organizations can take advantage of new features while maintaining compliance and system stability.

## Monitoring & Maintenance

Monitoring is critical for maintaining a reliable healthcare AI system. Usage analytics provide insights into how users interact with the agent, which scenarios are most active, and where improvements may be needed. Performance monitoring helps identify latency, errors, or configuration issues before they impact users. Feedback-driven improvement allows organizations to refine responses, scenarios, and workflows based on real user input. By continuously reviewing analytics and feedback, teams can ensure that Healthcare Agent Service remains accurate, efficient, and aligned with user needs over time.

## Final Summary & Enterprise Readiness

### Enterprise-Grade Capabilities

Healthcare Agent Service is designed as an enterprise-grade platform that can scale reliably across organizations of all sizes, from small clinics to large hospital networks. Built on Microsoft Azure, the service supports high availability, elastic scaling, and consistent performance even during peak usage. This ensures that healthcare agents remain responsive and dependable in real-world, high-demand clinical and administrative environments. Security is a core foundation of the Healthcare Agent Service. The platform follows Azure's security best practices, including secure authentication, encrypted data transmission, and protected data storage. Sensitive healthcare information is handled with strong safeguards to prevent unauthorized access, data leakage, or misuse. These protections are applied consistently across conversations, integrations, and external system connections.

Compliance readiness is another key enterprise capability. Healthcare Agent Service is designed to support regulatory requirements such as HIPAA and GDPR by incorporating audit trails, consent management, data retention controls, and privacy-aware processing. These built-in compliance features help organizations adopt AI responsibly while meeting legal, ethical, and organizational governance standards.

### Real-World Healthcare Adoption

Healthcare Agent Service is built for practical adoption in real healthcare environments, supporting both clinical and non-clinical workflows. In clinical settings, the service assists healthcare professionals by providing evidence-based medical information, symptom triage guidance, and structured data access. While it does not replace medical professionals, it acts as a supportive tool that improves efficiency and

information availability during care delivery. Administrative automation is another major area of impact. Healthcare agents can handle repetitive tasks such as appointment scheduling, policy inquiries, documentation support, and patient guidance. By automating these workflows, organizations can reduce administrative workload, minimize human error, and allow staff to focus more on patient care and high-value activities.

AI-assisted decision support further enhances healthcare operations by delivering grounded, evidence-backed insights drawn from credible medical sources and organizational data. With strong safeguards in place, the system provides consistent and trustworthy responses that support informed decision-making without compromising safety or compliance. This makes Healthcare Agent Service a practical, scalable, and trustworthy AI solution for modern healthcare organizations.