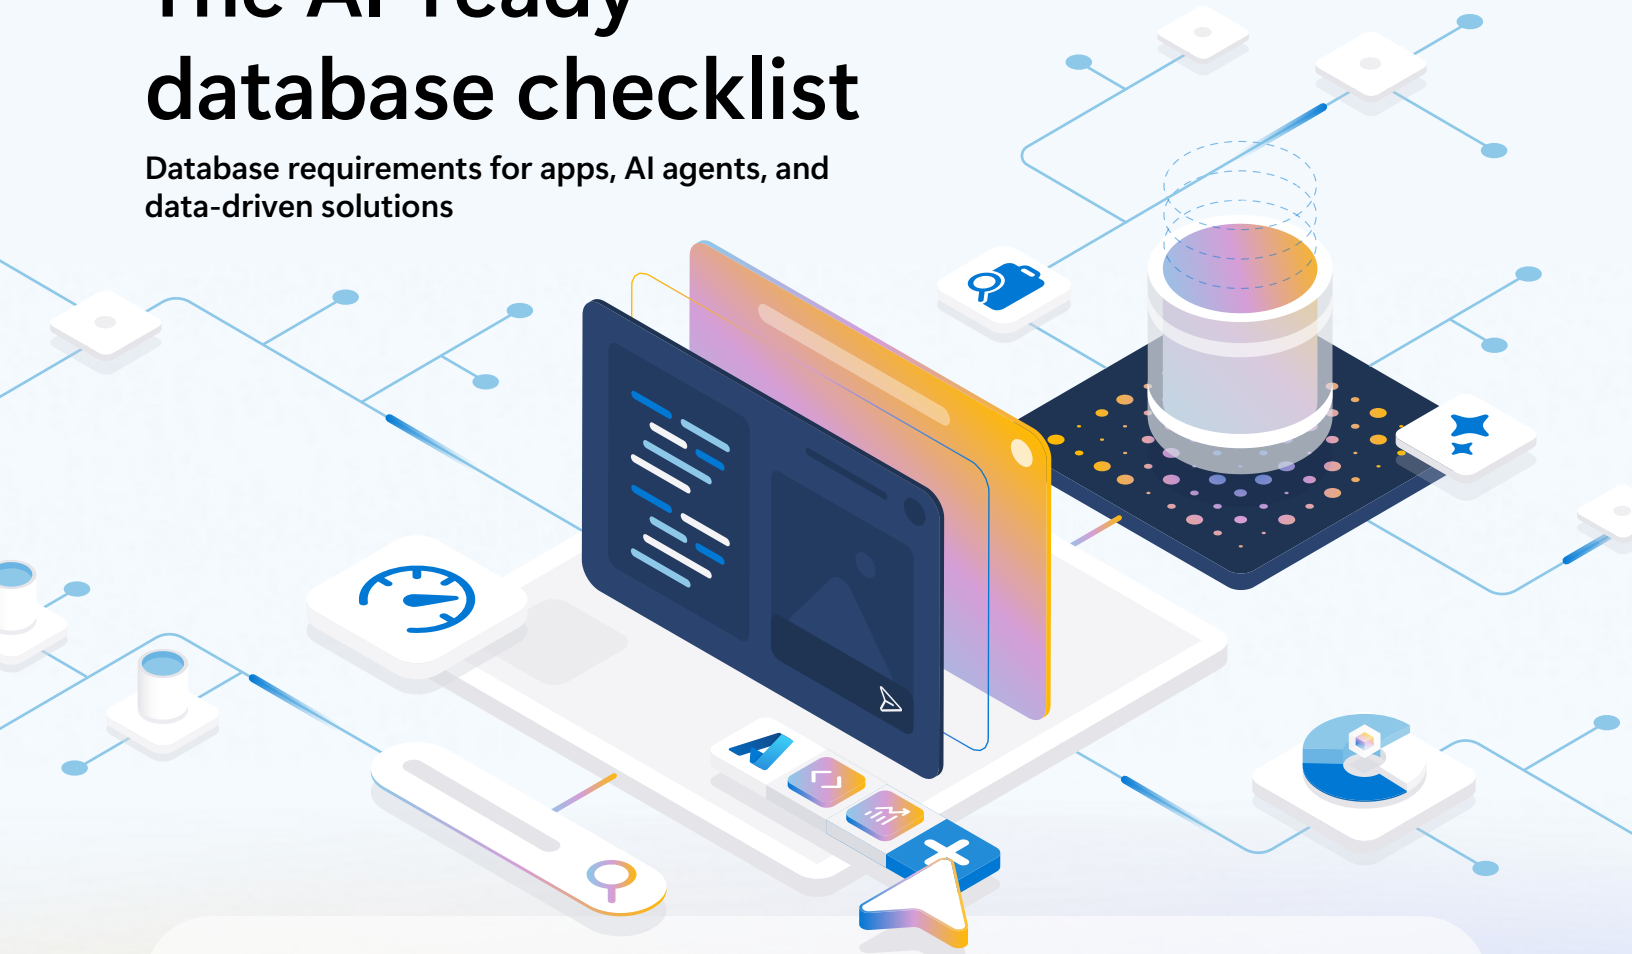






# The AI-ready database checklist

Database requirements for apps, AI agents, and data-driven solutions



## Is your database AI-optimized?

Look for these essential capabilities:

-  **Data at AI scale**  
Capacity and elasticity to handle changing needs
-  **Data in real-time**  
Quickly ingest, process, and make data available to apps and models
-  **Fast and efficient vector support**  
Powerful and cost-effective vector storage, semantic search, and approximate nearest neighbor
-  **Powerful integrations**  
Seamless connections to AI models, tools, and analytics systems

# Build AI agents and apps faster with Azure Databases

Azure Databases provide the scale, speed, and integrations to power it. Build with confidence, knowing your data is ready for AI.

## Azure Cosmos DB



## Azure Database for PostgreSQL



## Azure SQL Database



### Data at AI scale

- Support workloads of any size
- Flexibly adapt to changes and new requirements
- Elasticity to support dynamic workloads



### Data in real-time

- Provide AI models and analytics platforms with immediate access to data
- Low-latency reads and writes to apps and agents



### Fast and efficient vector support

- Ability to natively store and query vectors
- Advanced information discovery with vector search, reranking, and semantic search



### Powerful integrations

- Built-in support for AI models and tools
- Enable advanced analytics through real-time access for OLAP systems
- Easily integrate with Microsoft Foundry, Azure Databricks, and Microsoft Fabric

## AI use cases Azure Databases are built to power

[Learn more](#)



### Operational + Vector Database

Keeping operational and vector data together in the same database to reduce complexity and cost



### Retrieval Augmented Generation (RAG)

RAG (Retrieval Augmented Generation) Apps to personalize GenAI experiences



### Conversational History

Optimize application responses, manage threads, and enable auditing



### Semantic Caching

Reduce latency and cost of GenAI experiences



### AI agents

Multi-agent and AI agent solutions



### Fine-tune SLMs

Use SQL data to fine-tune low latency SLMs used in mobile apps