Guidance for Agentic AI Research Platform on AWS

This architecture showcases how Amazon Bedrock AgentCore enables the deployment and operation of specialized AI agents, automatically routing research queries to the right tool whether it's for data analysis, knowledge search, web research, or project management. The platform enables these agents to work together seamlessly, sharing context and collaborating on complex research tasks that would typically require entire analyst teams. This approach delivers intelligent, scalable research assistance by automatically matching each question or task with the most capable AI specialist to provide comprehensive, accurate results.

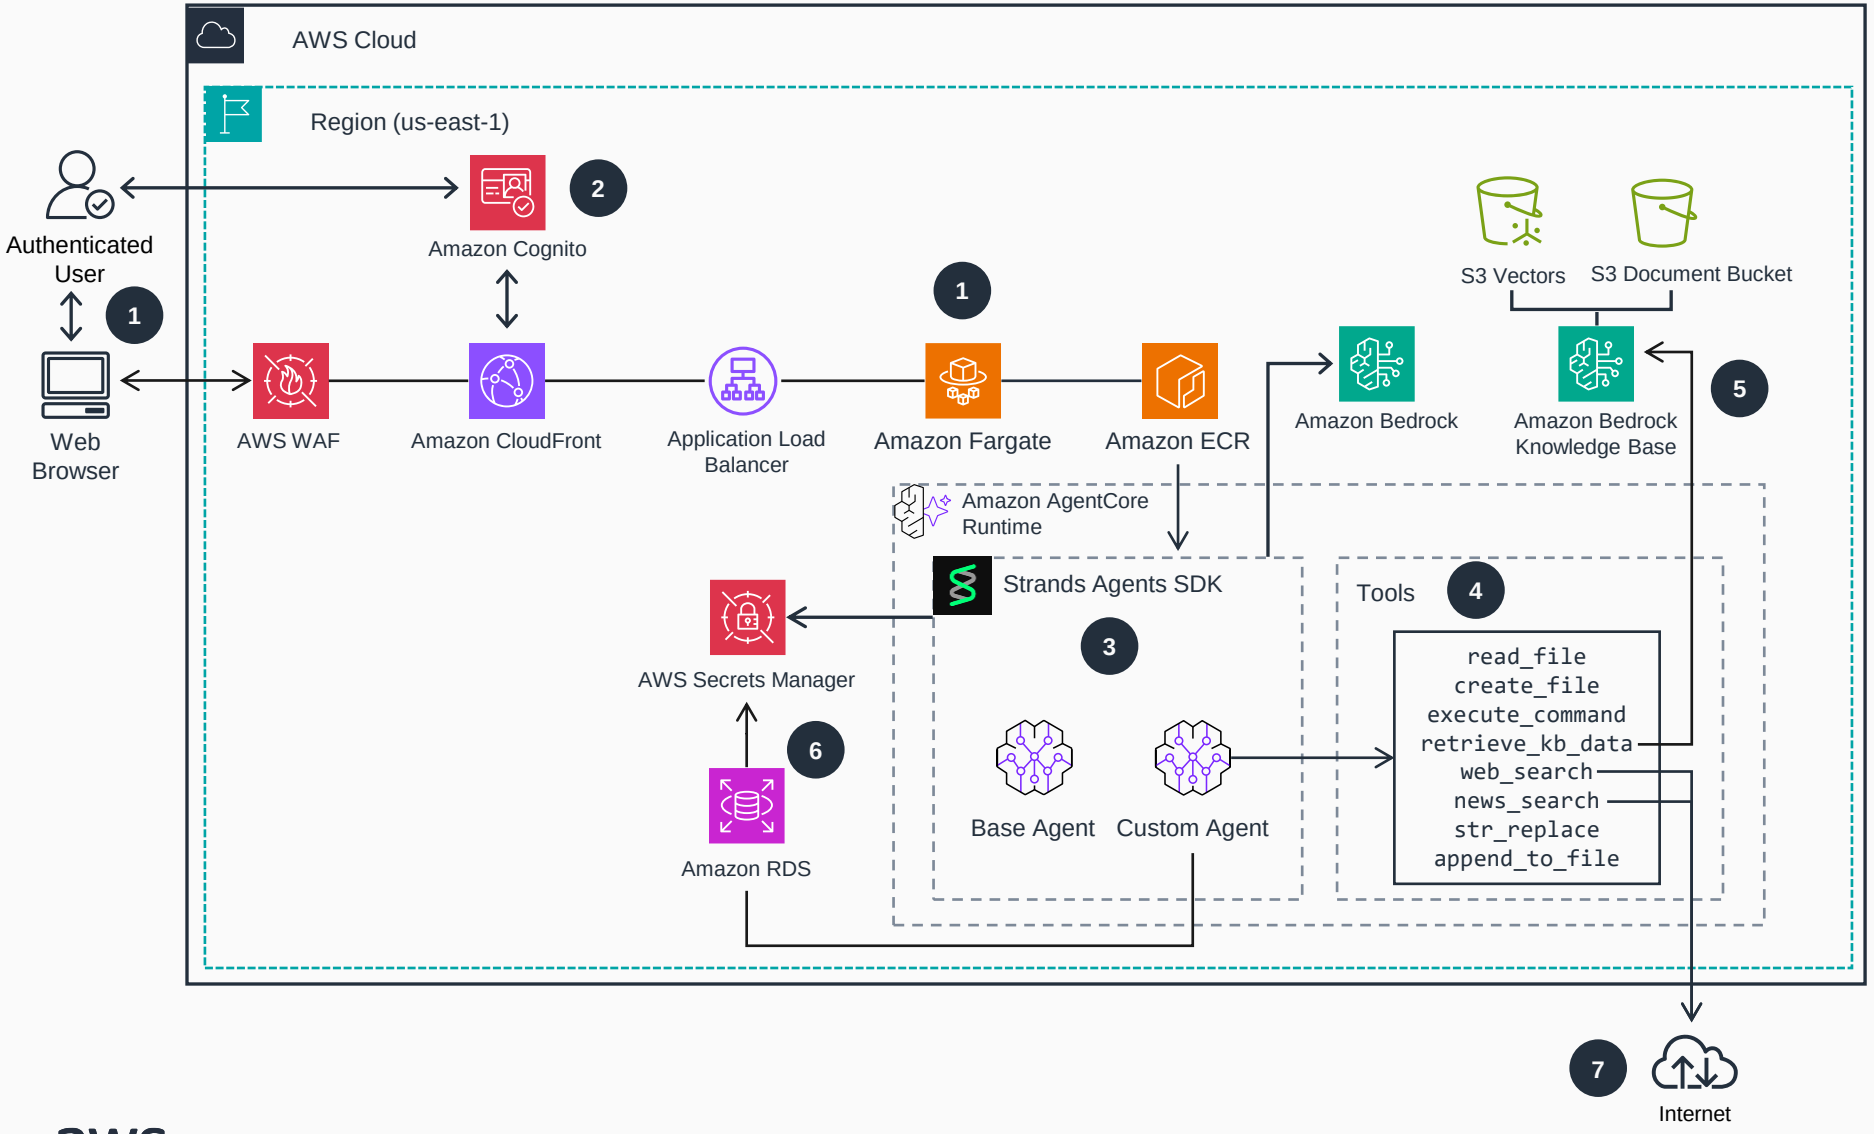

1. The authenticated user accesses the web application through their browser, with requests first passing through **AWS WAF** for security filtering. **Amazon CloudFront** then delivers the static web content, while **Amazon Cognito** handles user authentication and authorization for secure access to the platform.

2. User requests are routed through the **Application Load Balancer** to **Amazon Fargate**, which hosts the containerized application running on **Amazon ECS**. The container images for the application are stored and managed in **Amazon ECR**, ensuring secure and reliable deployment of the **Amazon Bedrock AgentCore Runtime** and associated components.

3. The application leverages **Amazon AgentCore Runtime**, built with the **Strands Agents SDK**, to orchestrate and manage the specialized AI agents that intelligently route research queries to the appropriate handler. When a research query is received, the **Amazon Bedrock AgentCore Runtime** routes it to either the Base Agent or Custom Agent depending on the task complexity and requirements.

4. These agents are powered by **Amazon Bedrock** and have access to a comprehensive toolset including file operations (read_file, create_file, append_to_file), command execution (execute_command), knowledge retrieval (retrieve_kb_data), search capabilities (web_search, news_search), and text manipulation (str_replace).

5. The agents retrieve contextual information and domain knowledge from **Amazon Bedrock Knowledge Base**, which is populated with relevant data stored in S3 Vectors for semantic search and **S3 Document Bucket** for raw document storage.

6. Sensitive configuration data, API keys, and database credentials required by the agents are securely stored and retrieved from **AWS Secrets Manager**, which maintains a connection to **Amazon RDS** for persistent storage of secrets and configuration management. This ensures that all agent operations maintain security best practices while accessing necessary resources to complete research tasks and deliver intelligent, scalable research assistance.

Reviewed for technical accuracy February 5, 2026

AWS Reference Architecture

Guidance for Agentic AI Research Platform on AWS

This architecture showcases how Amazon Bedrock AgentCore enables the deployment and operation of specialized AI agents, automatically routing research queries to the right tool whether it's for data analysis, knowledge search, web research, or project management. The platform enables these agents to work together seamlessly, sharing context and collaborating on complex research tasks that would typically require entire analyst teams. This approach delivers intelligent, scalable research assistance by automatically matching each question or task with the most capable AI specialist to provide comprehensive, accurate results.

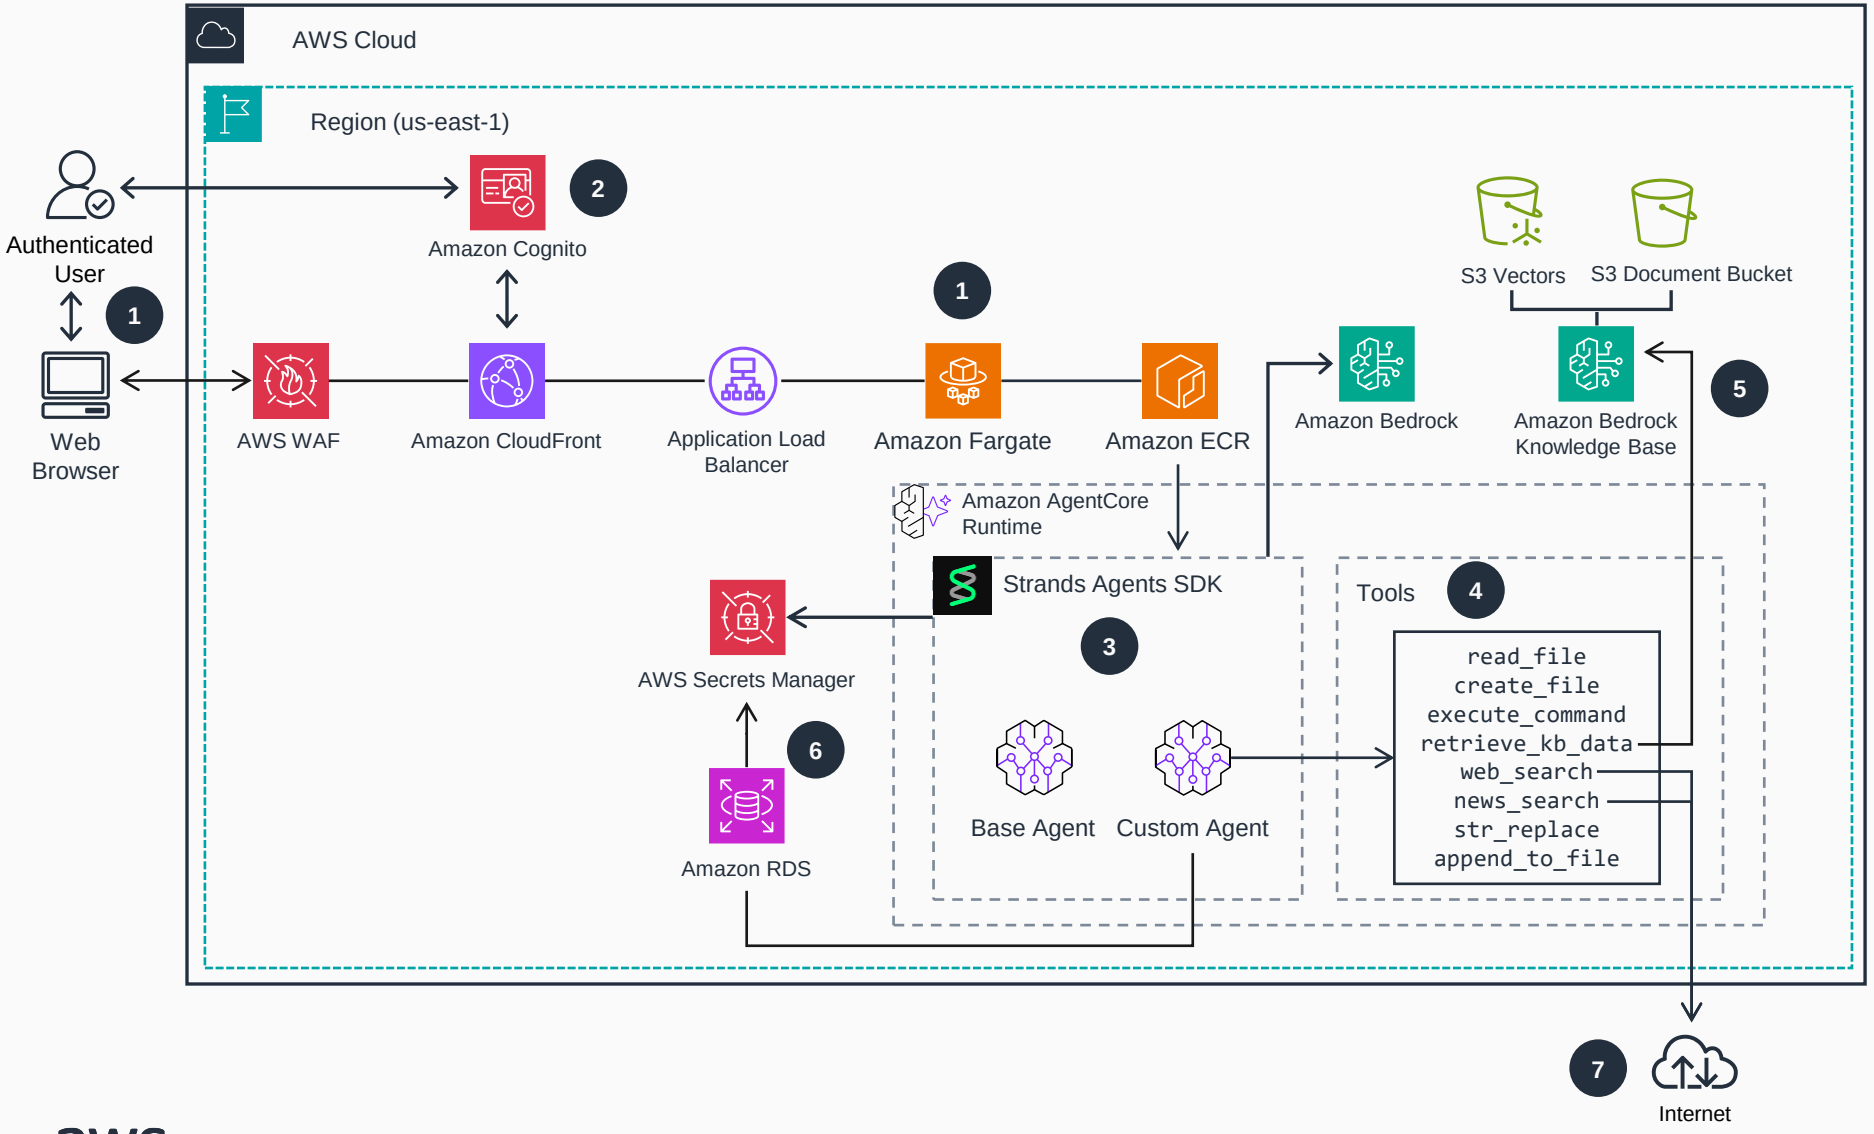

7 The web_search and news_search tools allow agents to search for websites and news from the public internet.

AWS Reference Architecture