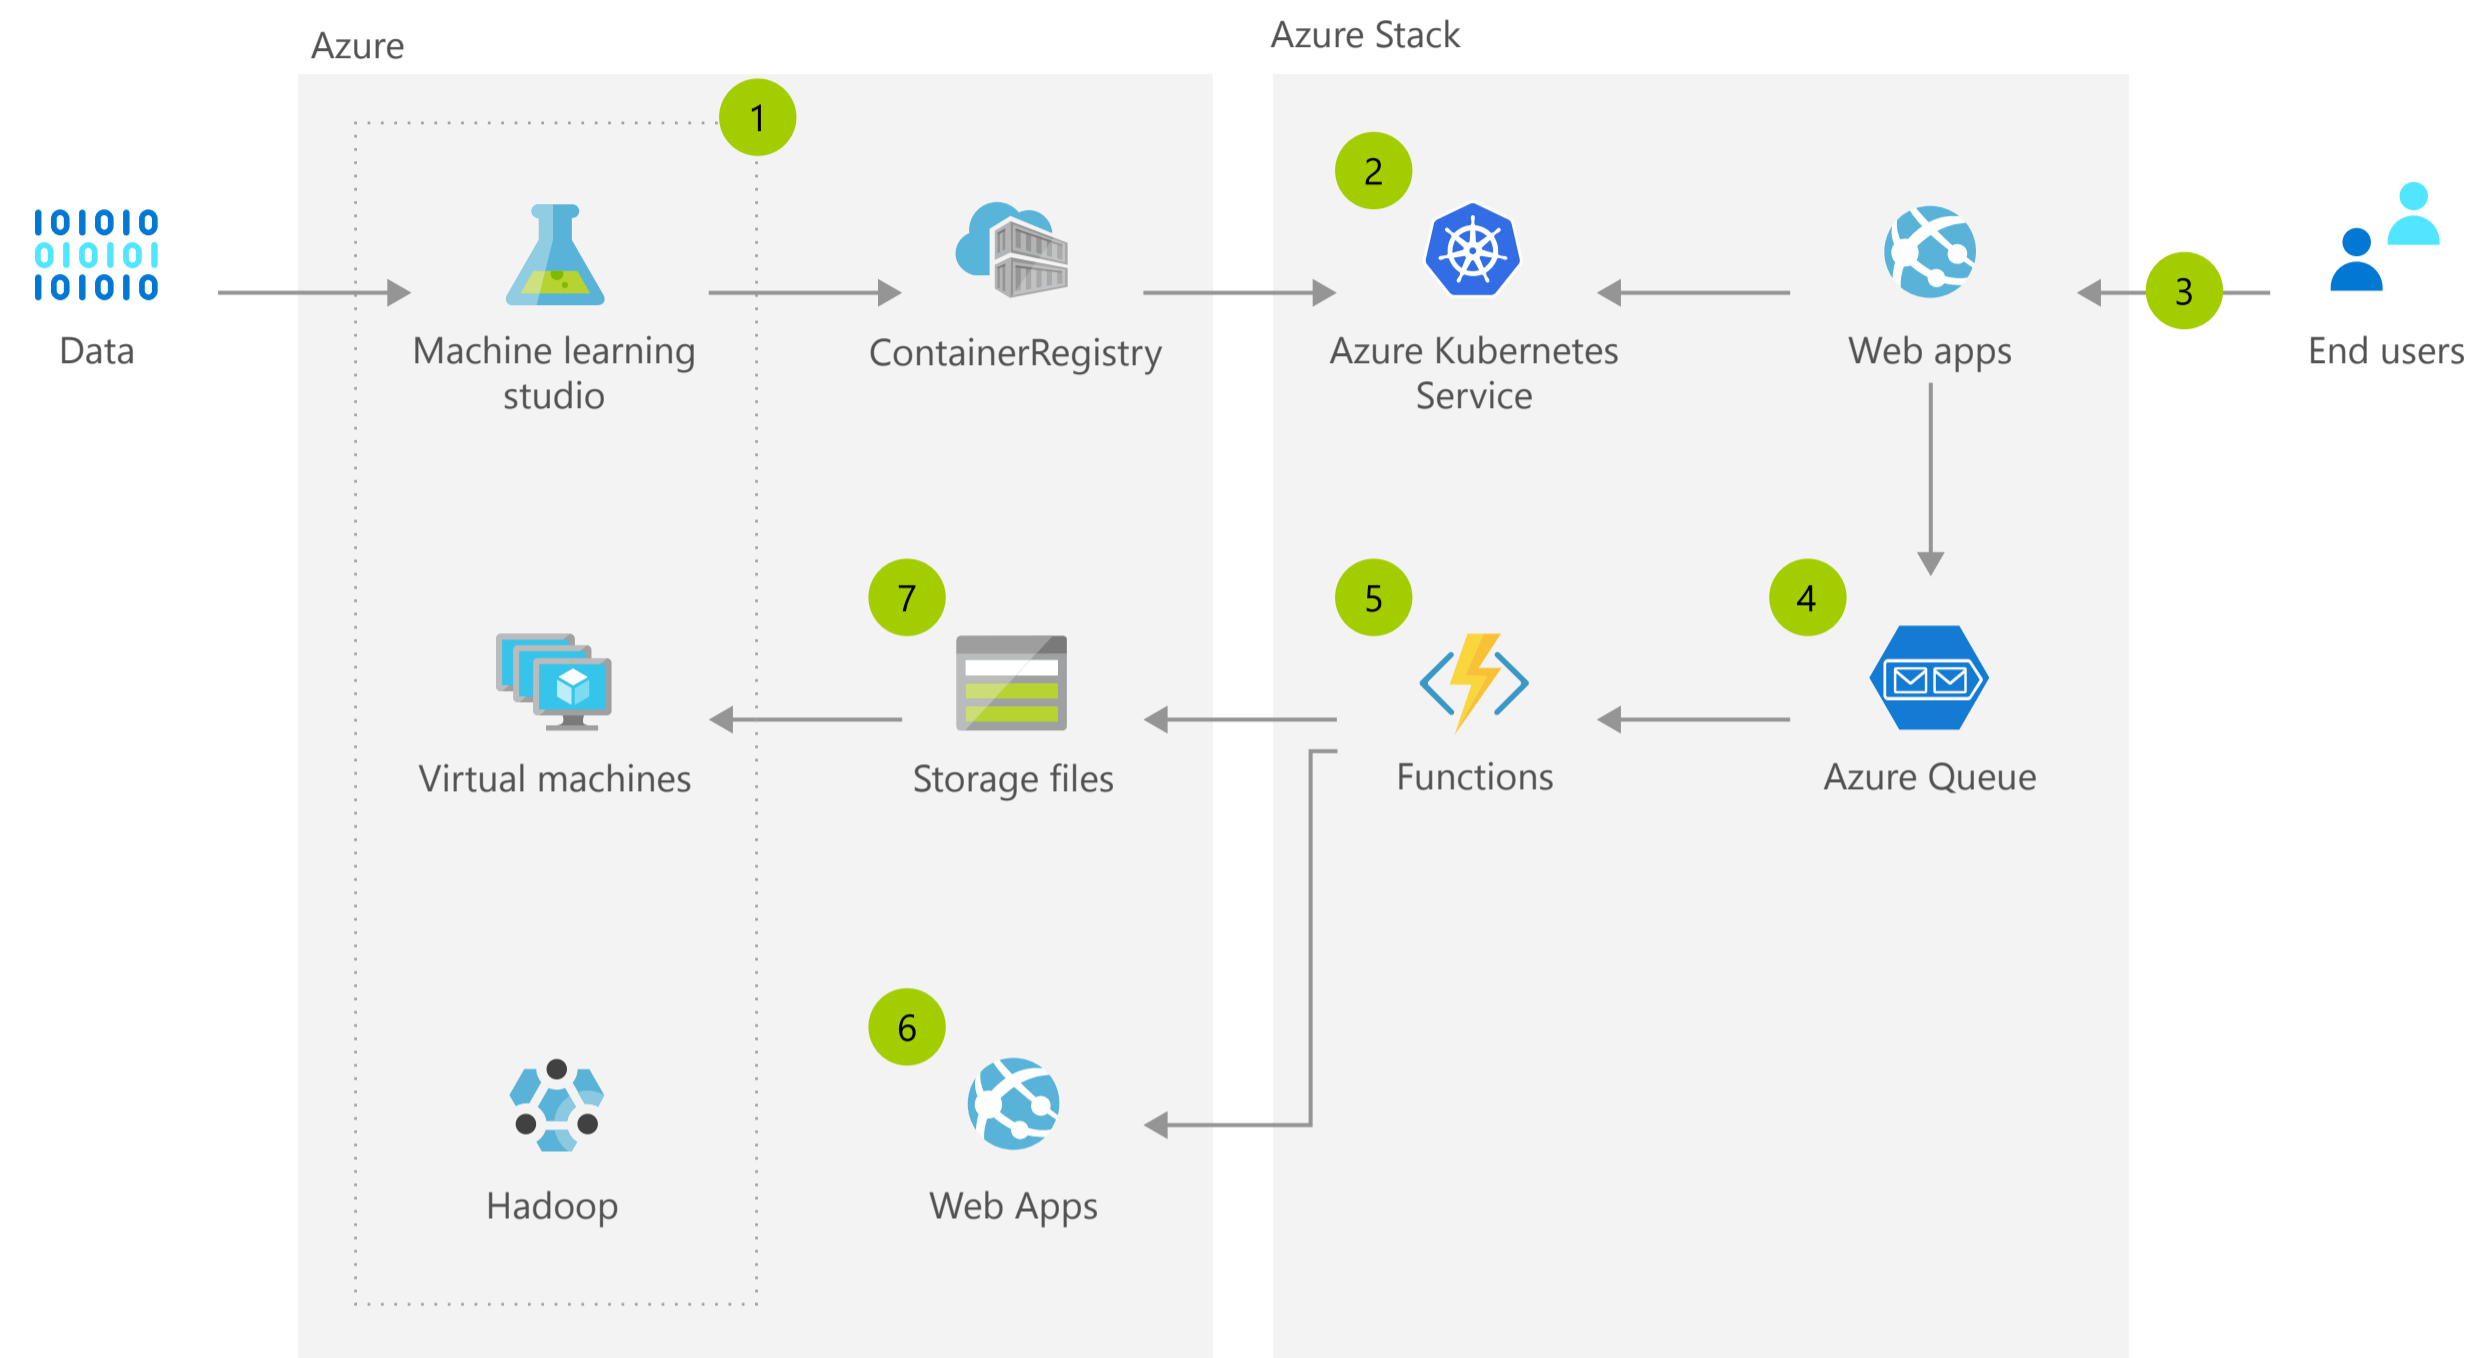


AI at the Edge



Architecture overview

With the Azure AI tools and cloud platform, the next generation of AI-enabled hybrid applications can run where your data lives. With Azure Stack, bring a trained AI model to the edge and integrate it with your applications for low-latency intelligence, with no tool or process changes for local applications.

- 1 Data scientists train a model using Azure ML workbench and an HDInsight cluster. The model is containerized and put in to an Azure Container Registry.
- 2 The model is deployed to a Kubernetes cluster on Azure Stack.
- 3 End users provide data that is scored against the model.
- 4 Insights and anomalies from scoring are placed into a queue.
- 5 A Function sends compliant data and anomalies to Azure Storage.
- 6 Globally-relevant and compliant insights are available in the global app.
- 7 Data from edge scoring is used to improve the model.

Azure products used in this solution

- Virtual Machines
- Machine Learning studio
- Storage files
- Functions
- Hadoop
- Azure Container Registry
- Azure Kubernetes Services
- Web Apps
- Azure Queue