

1 Data acquisition Edge computing

Data from sensors and other sources at the device edge are streamed using Red Hat AMQ streams (based on Apache Kafka on Kubernetes) to the public cloud (A, C) or datacenters (B), for analytics and machine learning (ML) model developments or updates, and for real time inferencing at the edge—to make data-driven predictions.

Containerized AMQ streams, inferencing applications, and other associated software are hosted at the edge using Red Hat OpenShift.

2 ML model and application developments (AppDev) or updates Datacenter or public cloud

Data generated at the edge streamed from step one to the datacenters (B) or public cloud (A, C) are stored in a data lake (based on Red Hat Ceph® Storage), and then used for analytics, ML models, and software application development. All the data engineering, models, software development, and delivery tools run on Red Hat OpenShift, including Red Hat Application Foundations, provide key capabilities like Red Hat Runtimes, application programming interface (API) management, and AMQ streams.

Once these models are tuned and ready for production, the artificial intelligence (AI)-powered intelligent application is deployed and automatically updated as needed, at the edge, using Red Hat OpenShift Pipelines and GitOps capabilities.

3 Inferencing to make data-driven decisions Edge computing

The intelligent AI-powered application running on Red Hat OpenShift at the edge, helps make real-time decisions based on new data streamed through Red Hat AMQ streams.

To speed up time to insight, the servers running Red Hat OpenShift may be powered by NVIDIA graphic processing units (GPUs).

Combined with Red Hat OpenShift, Red Hat Application Foundations creates a platform that streamlines execution across the entire MLOps life cycle by providing key capabilities like Red Hat Runtimes, application programming interface (API) management and AMQ streams.

Learn more, visit:
youtube.com/OpenShift

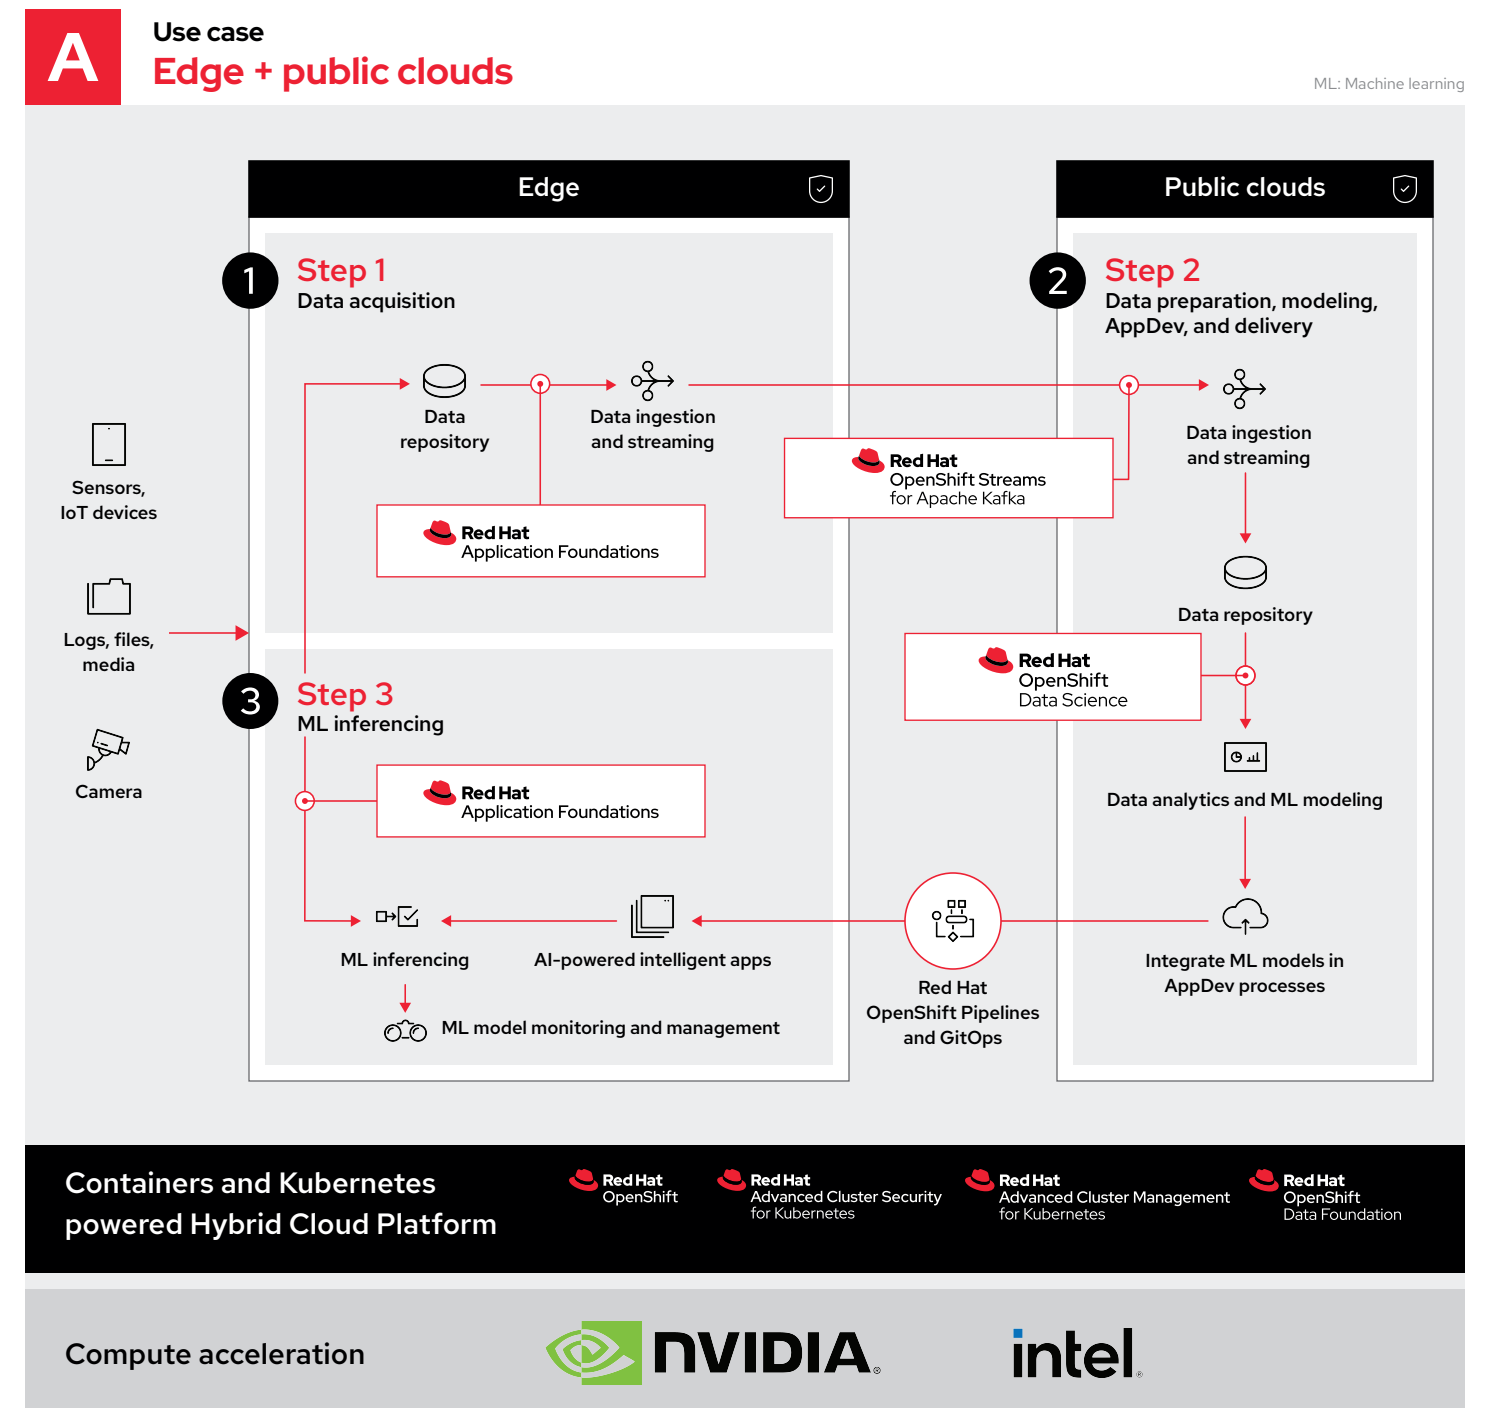
Watch our video:
AI/ML at the edge with Red Hat OpenShift

Watch our video:
Edge architectures on Red Hat OpenShift

Copyright © 2022 Red Hat, Inc. Red Hat, the Red Hat logo, OpenShift, and Ceph are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Different stages of AI at the edge

Industries world-wide are using edge computing with data generated from either Internet of Things (IoT) devices, sensors at factories, city streets, hospitals, vehicles, retail stores, or oil wells that need to be processed and acted upon in near real-time. Yet, every industry has its own unique deployment and architecture use case.



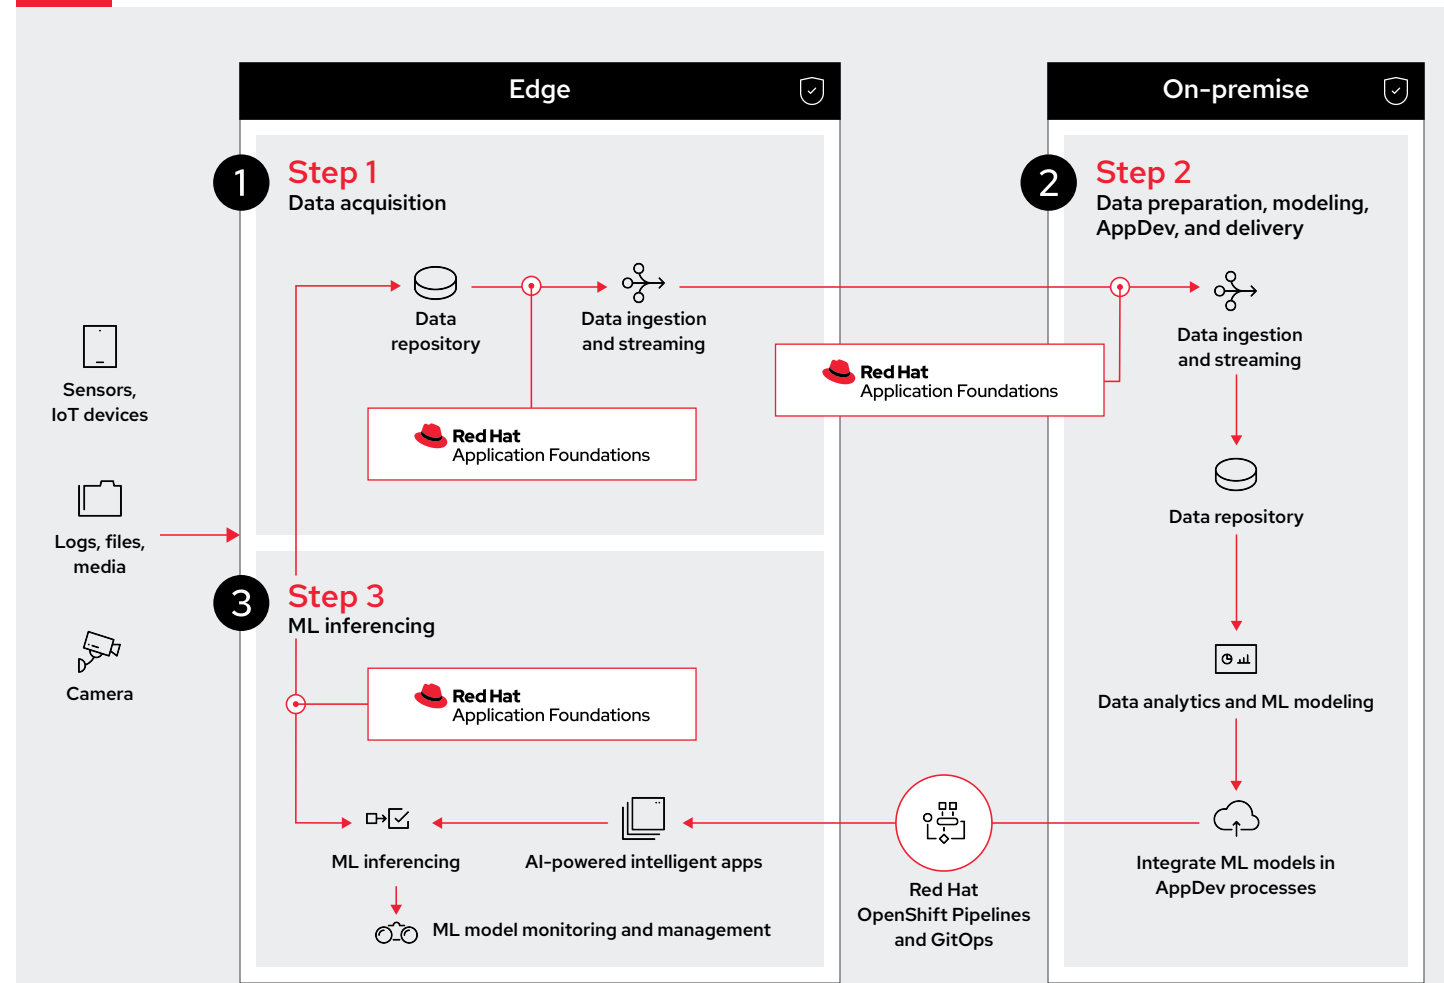
Red Hat® OpenShift® and Red Hat OpenShift Data Science allow you to design, deploy, and manage your intelligent applications consistently across cloud environments, datacenters, and edge.



Develop and run any workload in any environment or location

Red Hat edge helps organizations extend the open hybrid cloud across the edge, core, and cloud locations, allowing organizations to develop and run any workload in any environment or location.

B Use case
Edge + on-premise



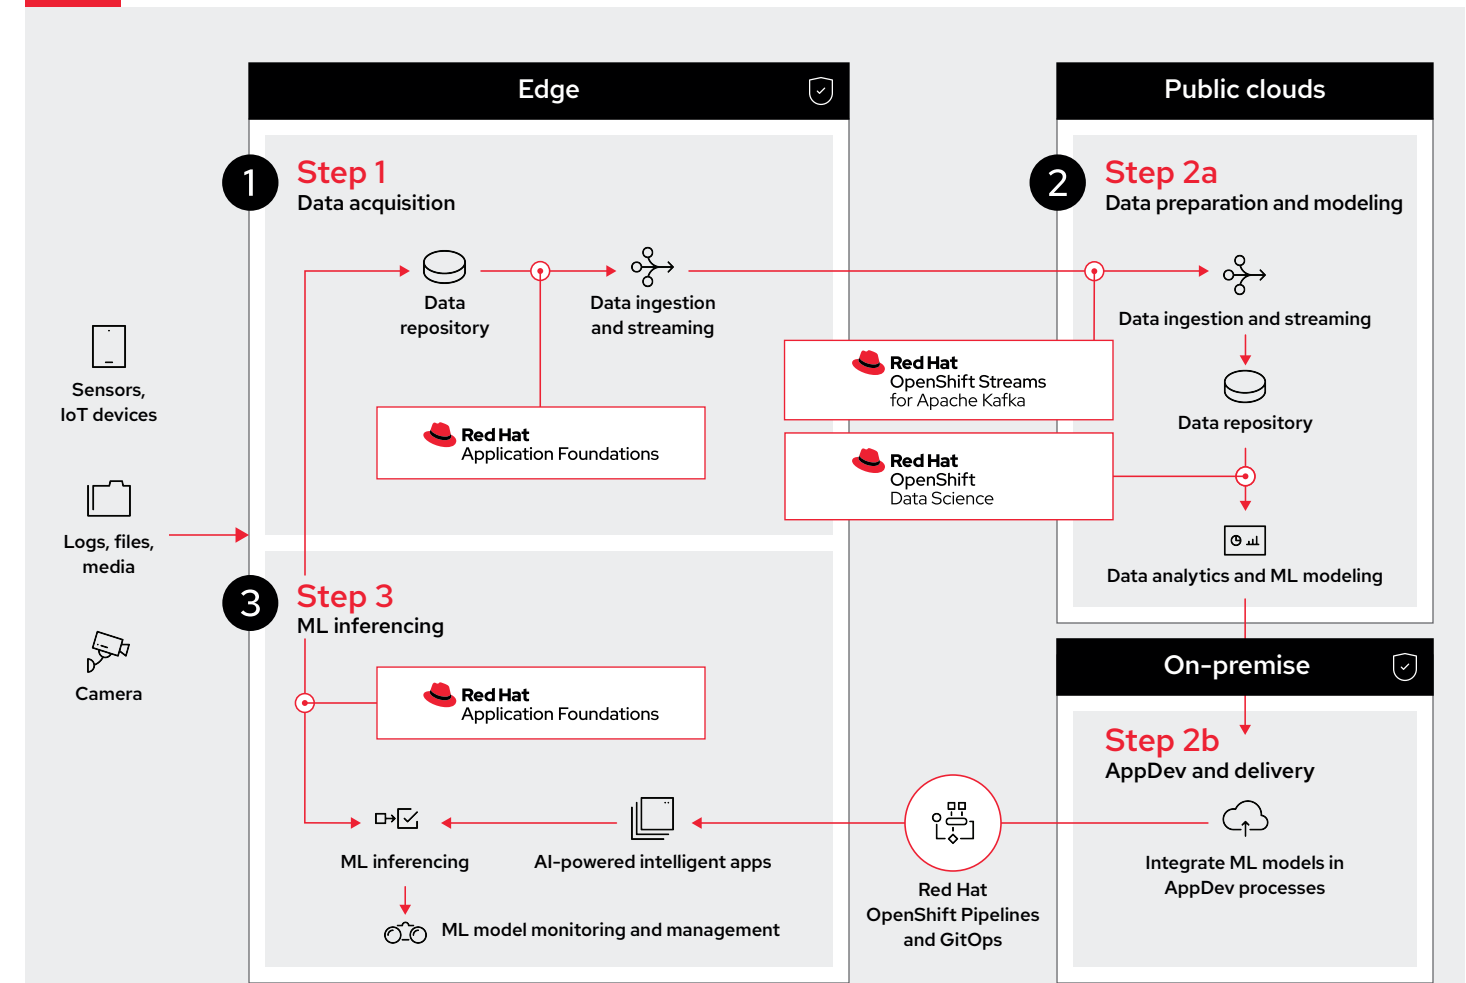
Containers and Kubernetes powered Hybrid Cloud Platform

Red Hat OpenShift, Red Hat Advanced Cluster Security for Kubernetes, Red Hat Advanced Cluster Management for Kubernetes, Red Hat OpenShift Data Foundation

Compute acceleration

nvidia intel

C Use case
Edge + hybrid cloud



Containers and Kubernetes powered Hybrid Cloud Platform

Red Hat OpenShift, Red Hat Advanced Cluster Security for Kubernetes, Red Hat Advanced Cluster Management for Kubernetes, Red Hat OpenShift Data Foundation

Compute acceleration

nvidia intel