Solution Brief

Data Center Operators, Service Providers

Intel and Red Hat Certify Kubernetes Operators for OpenShift

Operator Developed by Red Hat and Intel Delivers Agile and Scalable Cloud Native Solutions

Cloud native technologies are increasingly becoming an integral part of data center infrastructure and application modernization, as well as the basis for enterprise digital platform strategies.

Kubernetes is an industry-leading cloud native open-source container orchestration system. In a 2020 survey[1] from the Cloud Native Computing Foundation (CNCF), the use of Kubernetes increased to 83% of production data centers compared to 78% in 2019. In the survey, users ranked scalability, very fast deployment, and improved availability as the top benefits of Kubernetes in cloud-native projects. Developing and running applications using Kubernetes results in increased resilience, automation, and immutability.

Kubernetes is designed to place workloads on the least-used node in a cluster and automatically scale the application depending on application use. It is perfect for automating the deployment, scaling and management of microservices-based applications. Due to reduced overhead, application delivery is more agile and new application deployment can occur quickly. Additionally, Kubernetes encourages creativity through the facilitation of a DevOps culture of collaboration and unique design approaches.

One of the benefits of using cloud native containerization is that only a single control plane can be used to manage different deployments spanning various architectures. Kubernetes helps make this a possibility, as it is a proven orchestrator for hybrid and multi-cloud deployments. It offers established standards that enable data centers to realize the benefits of new computing paradigms such as edge computing. Containers improve portability and agility across different cloud infrastructures. Some of the biggest cloud providers are offering Kubernetes solutions because they make the hybrid cloud more cost effective and improve performance.

Built into the architecture of Kubernetes is its Operator functionality which enables expanded application functionality to be added to an application without a revision of the OS. Operators also allow Kubernetes deployments to take advantage of the latest compute hardware. To this end, Intel has developed a number of Operators to help data centers and other customers to build, migrate, and modernize applications using its accelerators, memory, artificial intelligence (AI) and analytics products. The company is working with Red Hat, an Intel® Network Builder ecosystem partner, to certify its Operators on Red Hat® OpenShift®.

Kubernetes Operators Extend Functionality

A Kubernetes Operator is a method of packaging, deploying and managing a Kubernetes-native application. Kubernetes applications are enterprise-ready containerized solutions with prebuilt deployment templates that manage applications built out of multiple container runtimes.

An Operator adds additional functionality and automation to complex applications and enables software vendors to build their applications to run on top of Kubernetes while accessing functionality not built into the OS. Operators provide a means to control and manage assets or containerized applications.

Kubernetes has allowed the deployment and management of distributed applications to be heavily automated, but Kubernetes wasn't designed to know about all application types. In these cases, its required to extend the understanding of a specific type of Kubernetes application into the OS. Without Operators, data center managers would have to manage these applications manually, reducing Kubernetes promised automation benefits. With Operators the automation functionality within Kubernetes is expanded into these new applications allowing data center service providers to capture and automate new tasks or functions.

Operators build on the controller concept that is built into Kubernetes and can be considered a custom controller for the functionality added by the new software. Controllers are software loops that runs continuously on the Kubernetes master nodes comparing and reconciling, if needed, the expressed, desired states of an object with the current state of that object. Objects are well known resources like pods, services, ConfigMaps, or PersistentVolumes. Operators apply this model at the level of a specific application.

Operators run in a Pod on the server cluster, interacting with the Kubernetes API server. It introduces new object types through custom resource definitions. These definitions are built into Kubernetes allowing for these extensions. These custom objects are the primary interface for a user, consistent with the resource-based interaction model on the Kubernetes cluster.

An Operator watches for these custom resource types and is notified about their presence or modification. When the Operator receives this notification, it will start running a loop to ensure that all the required connections for the application service represented by these objects are actually available and configured in the way the user expressed in the object's specification.

Red Hat OpenShift

Red Hat OpenShift is a hybrid cloud, enterprise Kubernetes application platform, trusted by more than 2,000 organizations. OpenShift is both containerization software and a platform-as-a-service (PaaS). OpenShift offers consistent security, built-in monitoring, centralized policy management, and compatibility with Kubernetes container workloads. OpenShift is engineered to be fast, with self-service provisioning, and integration with a range of tools.

Red Hat OpenShift allows data center Operators to operate consistently and innovate continuously - regardless of where applications reside. It uses the power of Operators to extend the functionality of the software and run the entire platform in an autonomous fashion. Configuration is exposed natively through Kubernetes objects, allowing for quick installation and ongoing updates.

Through its Operator partnership program, Red Hat OpenShift features 30 Operators which run the major parts of the platform, including version control, ingress, cluster autoscaling and others. Red Hat has worked with Intel to certify several Operators for use with OpenShift.

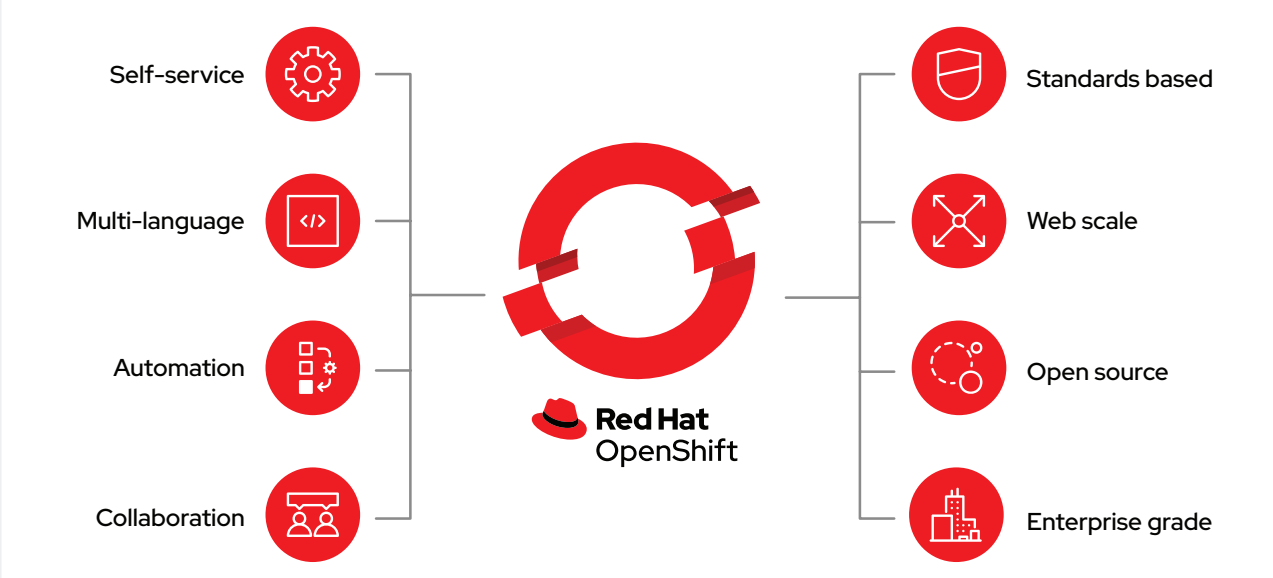

Figure 1. Key functionality of Red Hat OpenShift Kubernetes distribution.

Intel's Operator Strategy

Intel creates Operators that make it easy for its partners to deploy its silicon solutions in a Kubernetes application. The company has focused its early Operator developments on access edge networks, specifically 5G open and virtual radio access network (vRAN/O-RAN) applications.

Intel worked with Red Hat to certify these Operators for OpenShift users and they can be downloaded from the Red Hat Ecosystem Catalog. Right now, there are five validated Intel Operators for Red Hat OpenShift. Below is a short description of each Operator:

- **Intel Operator for Intel® FPGA PAC N3000 (Programming):** The Intel Operator for Intel® FPGA PAC N3000 (Programming) is a fully functional tool to manage the Intel FPGA PAC N3000 card and its resources autonomously in a Red Hat OpenShift environment. This Operator is a collection of container images for programming Intel FPGA PAC N3000 on OpenShift based on the Software Developer Kit (SDK) Operator framework.
- **Intel Operator for Wireless FEC Accelerators:** The Intel Operator for Wireless FEC Accelerators (including the Intel vRAN Dedicated Accelerator ACC100) is a collection of container images for management of SR-IOV forward error correction (FEC) resources of the Intel Wireless FEC Accelerators on OpenShift. This Operator is also a collection of container images for the Intel Wireless FEC Accelerators based on the SDK Operator Framework. It is a complete package for automating installation, the configuration of SR-IOV FEC IP, and resource allocation to the pods.
- **Intel® Optane™ PMem CSI Operator:** Intel® Optane™ PMem CSI Operator is a container storage interface (CSI) storage driver for Red Hat OpenShift. It makes local persistent memory (PMem) available as a filesystem volume to container applications.
- **Intel® oneAPI AI Analytics Tookit Operator:** The Intel® oneAPI AI Analytics Toolkit Operator gives data scientists, AI developers, and researchers familiar with Python® tools and frameworks to accelerate end-to-end data science and machine learning pipelines on Intel architectures. The components are built using oneAPI libraries for low-level compute optimizations. This toolkit maximizes performance from preprocessing through machine and deep learning phases and provides interoperability for efficient model development and deployment across single and multiple nodes.
- **OpenVINO™ Toolkit Operator:** The OpenVINO™ Toolkit Operator simplifies managing optimized deep learning inference at scale in OpenShift. This Operator includes OpenVINO notebooks for development and OpenVINO Model Server for deployment. Features include easy deployment of model server at scale, support for multiple model storage options, configurable resource restrictions and security context, and configurable service options.

Conclusion

Intel is committed to supporting its customers in their transition to cloud-native applications. Through its development of Kubernetes Operators for the Intel FPGA PAC N3000, Wireless FEC Accelerator (Intel® vRAN Dedicated Accelerator ACC100), OpenVINO and Intel oneAPI AI Analytics, the company is demonstrating support for high-speed edge network applications based on Kubernetes. Certifying these Operators on Red Hat OpenShift makes them trustworthy for OpenShift customers allowing them to extend the OS' functionality to deliver greater performance, agility and scalability.

Learn More

Red Hat OpenShift

Intel Red Hat Certified Operators

Intel® Network Builders

Notices & Disclaimers

[1] https://www.cloudflight.io/expert-views/cloud-native-kubernetes-trends-2021-44653/

Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

0522/TM/H09/PDF ♻ Please Recycle 349568-001US