



Red Hat and Partners Launch OperatorHub.io to Expand Kubernetes Automation Ecosystem

March 11, 2019

By: [Mary Johnston Turner](#), [Jim Mercer](#), [Gary Chen](#)

IDC's Quick Take

Red Hat's recent [launch of OperatorHub.io](#), in collaboration with a range of partners (Google Cloud, AWS, and Microsoft), offers enterprises a single open source destination to locate and download native full-stack Kubernetes Operator automations designed to configure, deploy, and manage complex containerized applications at scale on immutable infrastructure. Operators are designed to treat the full stack of custom application code, middleware services, Kubernetes configurations, storage and data management configurations, and cloud integrations as a single immutable unit controlled by highly prescriptive and documented automation code. This represents a potentially disruptive approach to configuration and provisioning that could drive I&O teams to restructure core operational processes and skills as the use of containers, Kubernetes, and DevOps proliferates across operational IT environments.

News Highlights

The Kubernetes Operator concept has been nurtured by CoreOS (now Red Hat) and upstream Kubernetes community members for the past two years. A number of cloud providers and ISVs have introduced public operators and/or use them internally. Google Cloud has recently published Operators for such use cases as Apache Spark on Google Cloud, which has been available in the Google Cloud Marketplace since January 2019. Other vendors such as Couchbase and Dynatrace have also been early implementers, and Red Hat has promoted the concept during the beta program for OpenShift 4.0, which incorporates significant automation capabilities, including Operators, that were acquired from CoreOS.

OperatorHub.io is designed to provide a unified, public registry for all vendors to publish and share Operators in a way that will simplify distribution and make it easier for customers to find curated Operator-backed services and basic documentation. It also identifies active Operator communities and vendor-backed initiatives including maintenance commitments and basic testing results.

IDC's Point of View

Kubernetes is a widely supported container orchestration technology used to dynamically automate the scaling and workload balancing of container-based infrastructure. Stateless cloud-native applications are well suited to the horizontal scaling, automated self-healing restarts, and progressive rollout of new containers that are the hallmarks of Kubernetes orchestration. Stateful applications, however, are not always well suited to this type of rapid, automated infrastructure behavior as they require persistence and predictability to remain in a stable state. While Kubernetes has been developing more features to accommodate stateful applications, these types of applications are inherently more complex to deploy, manage, and operate. It's often said there are no stateless apps, only stateless parts of apps. Statefulness will remain a key element of cloud-native apps, and management of these types of workloads will need to be incorporated into containerization strategies.

Initially developed by CoreOS (acquired by Red Hat), the Operator model defines an approach to full-stack automated configuration, deployment, and management of containerized applications, enabling services, OS, storage/data, and cloud dependencies as immutable units using Kubernetes-native automations. Operators can be developed for any type of containerized application, but are particularly useful for stateful apps, which are more complex and require more elaborate automation. The [Operator Framework](#) is an open source toolkit that provides an SDK, life-cycle management, metering, and monitoring capabilities that enable developers to build, test, and publish Operators. Examples currently available in the OperatorHub.io registry include:

- **AWS Service Operator** manages AWS resources using Kubernetes Custom Resource Definitions (CRD). Using the AWS Service Operator enables a GitOps workflow to drive infrastructure to the desired state leveraging Kubernetes CRD, the Kubernetes internal control loop, and AWS CloudFormation orchestration.
- **Couchbase Autonomous Operator** allows users to deploy Couchbase Server provision nodes and set up clusters with a single command. The operator can detect and auto-recover from node failures, replicate data across geo-diverse datacenters, and define persistent attached storage for each node.
- **Dynatrace OneAgent** installs full-stack monitoring on Kubernetes clusters and connects back to Dynatrace's hosted monitoring tools.

Operators follow a [maturity model](#) that ranges from basic functionality to including specific operational logic for an application. Operators' capabilities differ in sophistication depending on how much intelligence has been added into the Operator itself. Advanced Operators are designed to handle upgrades more seamlessly and to react to failures automatically.

Advice for Technology Buyers

The Operator approach to Kubernetes application and infrastructure automation is a significant departure from traditional infrastructure as code and legacy patching and life-cycle configuration control approaches. Built for immutable architectures, the Operator specifies the full suite of apps to infrastructure configurations and dependencies needed to set up and manage nodes and clusters to support even complex, stateful applications. Early users are quoted as saying the use of Operators can shave weeks off the process of setting up Kubernetes clusters to support their applications.

For many enterprise I&O teams, the management of Kubernetes-based applications at production scale is just becoming a concern. Many have found the setup and management of Kubernetes clusters to be too complex to handle internally and have opted to use cloud services that abstract the complexity away and allow developers to focus on code and functions. Optimizing I&O processes, skills, and toolsets for large-scale operation of stateful and stateless applications on shared Kubernetes infrastructure will become an increasing priority for many I&O leaders.

This ability to consistently automate the operation of applications on Kubernetes running at production scale will be particularly important for enterprises that have embraced containers, at least in part, with a goal of ensuring application portability across multicloud environments. To fully enable that vision, I&O teams will need to be able to automate the entire Kubernetes life cycle and integrate those activities with monitoring, analytics, security, and compliance processes. By abstracting away the intricacies of Kubernetes management, Operators can potentially lower the barriers for moving applications to Kubernetes.

Operators may provide DevOps teams with enhanced capabilities for more predictability with deployment and runtime of applications without the need for deep Kubernetes knowledge. Developers can use the Operator SDK to create Operators for their applications or take advantage of Operators available on OperatorHub.io. I&O organizations can reap the benefit of opinionated deployments that ensure optimal configuration and tuning is in place. Operators can assist by providing telemetric data for improved monitoring and health checks as well as the potential for automated scaling. These new capabilities could help to increase both DevOps velocity and runtime reliability.

The initial set of Operators found on OperatorHub.io is focused on cloud and infrastructure configuration and agent deployments. These serve to illustrate the potential power of Operators in defining and maintaining complex operational environments at large scale. I&O leaders should monitor the evolution of the Operator framework as the community matures and expands.

Key Takeaways

- Operators have the potential to significantly enhance the deployment and running of applications on Kubernetes. While there is support from Red Hat and key cloud vendors, it is still early days. Some early success coupled with wider-scale adoption by the community and industry standardization is the key to the long-term success of Operators.
- Organizations have been hesitant to run stateful applications that require persistence and predictability on Kubernetes. Operators have the potential to address these concerns. The Operator model defines an approach to full stack automated configuration, deployment, and management of stateful application code, enabling services, and infrastructure dependencies as immutable units using Kubernetes-native automations.
- The launch of the OperatorHub.io registry is an important step in the maturity of the Operator approach to automation. It provides a central location to find an array of curated and validated Operators and affords contributors of Operators the chance to publicize their work.

Subscriptions Covered:

[DevOps Solutions: Software, Infrastructure and Transformation](#), [Software-Defined Compute](#)

Please contact the IDC Hotline at 800.343.4952, ext.7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC or Industry Insights service or for information on additional copies or Web rights. Visit us on the Web at www.idc.com. To view a list of IDC offices worldwide, visit www.idc.com/offices. Copyright 2019 IDC. Reproduction is forbidden unless authorized. All rights reserved.