

Automating the Move to Kubernetes - the Open Way

Nov 10, 2020



A recent [IBM Institute for Business Value \(IBV\) study on Hybrid Cloud](#) reveals that the value derived from hybrid, multi-cloud platform technology and operating model at scale is 2.5 times the value derived from a single platform, single cloud. Moreover, globally, 64% of advanced cloud companies recognize the need for enterprise transformation and application modernization to go hand-in-hand. So, as enterprises look to leverage the hybrid cloud for bringing substantial agility into their IT and to drive new cloud-enabled business models, the choice of platform becomes critical.

Kubernetes is increasingly becoming the standard platform for highly scalable and available cloud-native applications that are built and managed via agile DevOps practices. It provides one of the most compelling set of platform capabilities and tremendous community support towards a future-proof platform.


While many born-on-the-cloud applications are now running on Kubernetes, the large majority of enterprise applications still run on more traditional technologies such as J2EE and older cloud technologies such as Cloud Foundry and Docker Swarm. A substantial barrier to adoption of Kubernetes has been the fact that re-platforming existing applications from other platforms to Kubernetes can be very complex, costly and risky if not done right. At the same time, there is a lack of tools that help automate this re-platforming activity. As a consequence, this activity has to be largely done manually with the involvement of deep experts in the source platform, the target platform (Kubernetes) and the application being re-platformed. This combination of skills is often scarce and expensive. Additionally, doing this activity manually means the process of re-platforming can take several months with iterations and lack of predictability.

To address this challenge, about a year ago, researchers at IBM Research India started working on developing a tool that would be easy to use and would largely automate the translation of an application's deployment specification from source platforms such as Cloud Foundry and Docker Swarm to a Kubernetes-native deployment specification. This work has culminated in the creation of a tool called [Move2Kube](#). In our experiments with enterprise applications, we have seen substantial productivity gains. For instance, in a scenario for re-platforming a 100-container application from Docker Swarm to Red Hat OpenShift, we estimated a greater than 8x reduction in man days from 61 days to 7 days when compared with a manual approach.

A Community driven development

To enable the open community to adopt and enhance this capability further, **Move2Kube** has now been released into the open-source community by IBM. Given the lack of such tools and the community driven

development of Kubernetes, it is important to have the community drive the definition and development of such a migration tool for Kubernetes. [Move2Kube](#) has been open sourced as part of the [Konveyor](#) organization on github.com. Konveyor is a community of people passionate about helping teams and engineers modernize and migrate their applications to the hybrid cloud by building tools, best practices, and providing advice on how to break down monoliths, adopt containers, and embrace Kubernetes. The Konveyor community provides additional tools in various areas relevant to migration to Kubernetes, for instance mig-operator for migrating namespaces between Kubernetes clusters, and virt-operator for migrating virtual machines from traditional virtualization platforms to KubeVirt.



Move2Kube is architected to allow extensibility for supporting additional source platforms that may not already be supported. The modular design allows developers to easily extend Move2Kube even if their source platform is not supported out of the box, and is a foundational architectural principle to enable broad-based community adoption of Move2Kube for replatforming to Kubernetes.

Our team at IBM, as a part of this Konveyor community, looks forward to engaging with the broader open source community on Move2Kube and Konveyor. Several engagement channels have already been activated, including issues on github.com repositories for Move2Kube, a [Slack channel](#) dedicated to Konveyor in the Kubernetes workspace, a [Google group](#) and regular open [meetup calls](#). At the upcoming [KubeCon North America Virtual - Openshift Commons Gathering 2020](#), we will be sharing more insights on Konveyor and Move2Kube.

We are working to make Move2Kube more usable, [watch this space](#) for latest updates.

This blog is authored by -

[Amith Singhee](#) - Senior Manager, Hybrid Cloud, IBM Research, India

Blog Categories

[Modernize](#)