

IBM Cloud Paks

An open, faster, more
reliable way to build, deploy
and run applications on
hybrid cloud

Andre Tost
Distinguished Engineer,
IBM Hybrid Cloud
andretost@us.ibm.com

Kyle Miller
Principal Offering Manager,
IBM Hybrid Cloud
millerkc@us.ibm.com

Introduction

Enterprises employ cloud technologies to deliver innovation at scale and at lower cost. New services are often built natively on cloud, but can also come with risk of “vendor lock-in” and escalating cost. Existing applications can be rewritten, but rewriting thousands (if not tens of thousands) of applications from the ground up is both cost and time prohibitive, so taking steps to modernize existing applications can be an attractive approach with faster time to value. Both strategies -- building new cloud-native applications and modernizing existing applications to support cloud environments -- need to be done in an open, portable manner that helps clients improve time to value while avoiding lock-in. Containers and Kubernetes enable this by providing portability and consistency in development and operations, but developers and administrators are still required to continuously connect component layers and verify interoperability. In addition, collecting,

integrating and analyzing data enables data engineers and scientists to help application developers infuse AI into applications; but the trick is to do this without adding complexity and cost. And, then, once applications are built and connected to data, IT operations need them to run in an environment that is high performing, scalable and reliable. Today, around 80 percent of existing enterprise workloads have not yet moved to the cloud due to these challenges and enterprises struggle with movement, connectivity and management across clouds.

To help clients move more workloads, faster, to cloud and AI, IBM announces:

A family of IBM Cloud Paks that give developers, data managers and administrators an open environment to quickly build new cloud-native applications, modernize/extend existing applications, and deploy middleware in a consistent manner across multiple clouds. Today, IBM introduces five new Cloud Paks: IBM Cloud Pak for Applications, IBM Cloud Pak for Data, IBM Cloud Pak for Integration, IBM Cloud Pak for Multicloud Management, and IBM Cloud Pak for Automation that deliver IBM enterprise software and open source components in open and secure solutions that are easily consumable and can run anywhere.

IBM Cloud Paks provide:

- Containerized IBM middleware and open source components.
- Consistent added capabilities for deployment, lifecycle management, and production quality of service – logging, monitoring, version upgrade and roll-back, vulnerability assessment and testing
- Certification by IBM to run on IBM Cloud Private (ICP), IBM Cloud Kubernetes Service (IKS), and Red Hat OpenShift Container Platform (OCP), providing full software stack support, and regular security, compliance and version compatibility updates by IBM

The IBM Cloud Pak for Applications reduces development time to market by up to 84 percent by reducing the compute required and by accelerating throughput of the continuous integration continuous delivery (CICD) pipeline, and reduces operational expenses by up to 75 percent through increasing IT admin efficiency and reducing related labor costs.

(Source: <http://ibm.biz/Ovum-WP>)

IBM is committed to delivering enterprise software from across its portfolio for modern cloud environments. IBM Cloud Paks provide enterprise container software that is pre-packaged in production-ready configurations and can be quickly

and easily deployed to Kubernetes-based container orchestration platforms. In addition, these Cloud Paks provide resiliency, scalability, and integration with core platform services, like monitoring or identity management.

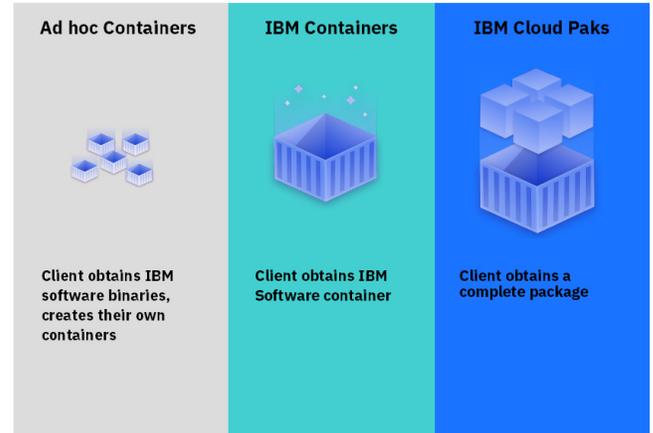


Figure 1. IBM software is supported and consumed as containers in 3 ways

IBM Cloud Paks enable you to easily deploy modern enterprise software either on-premises or in the cloud, and quickly bring workloads to production by seamlessly leveraging Kubernetes as the management framework supporting production-level qualities of service and end-to-end lifecycle management. This improves time-to-market, enterprise readiness and reduces the effort required to manage related applications, as shown in Figure 2.

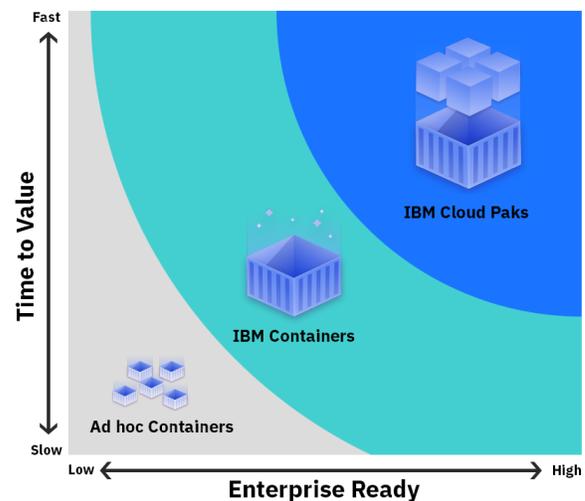


Figure 2. Time to Value and Enterprise Readiness of software

This paper describes IBM Cloud Paks in more detail, highlighting the additional value that this delivery model offers, with some background details on the underlying open technologies, for those who may be unfamiliar.

IBM Cloud Paks Simplify Enterprise-grade Deployment and Management for Software in Containers

IBM’s container platforms, IBM Cloud Private (ICP) and IBM Cloud Kubernetes Service (IKS), build on top of the open source [Kubernetes](#) orchestration technology. IBM is committed to delivering enterprise software designed for these modern container orchestration platforms and Red Hat OpenShift Container Platform.

Deploying complex software workloads in optimized and highly-available configurations can involve collecting or creating large numbers of disparate components, including the workload container images, configuration files, and assets for integrating with your chosen platforms or management tools.

IBM Cloud Paks bring together thoroughly-tested enterprise software container images using, [Helm](#) charts with intelligent defaults for simplified configuration and management and can include additional assets, such as [Operators](#) that intelligently manage software during runtime, in a single archive from a trusted source. As a result, you can quickly load software into your catalog, walk through a simple deployment experience, guided by logical defaults and helper text and easily deploy production-ready enterprise software onto IBM’s container platforms, in the cloud or in your own data center.

Core Services

IBM Cloud Private and IBM Kubernetes Service provide a set of core services by default, including security and identity services, logging, monitoring, auditing and more. For example, workloads that expose metrics using [Prometheus](#) can be monitored out of the box using the monitoring services provided by IBM’s container platforms. Similarly, logs that are generated by each workload container are collected and correlated by a platform-provided logging service that includes a collection, search and dashboarding capabilities.

Containers Revisited

Containers give you the ability to run multiple software elements, isolated from each other, within the same operating system instance. Unlike a virtual machine, a container shares the operating system kernel with its underlying host and since system calls can be made directly, a container can be run more efficiently and be instantiated faster, as shown in Figure 3.

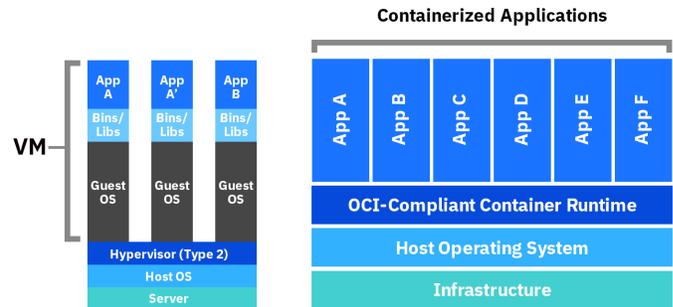


Figure 3. Virtual Machines compared to Containers

While containers are available in many forms and implementations, the Open Container Initiative (OCI) has emerged as the leading standard in the industry, defining open specifications for container images and container runtimes.

The fact that containers are lightweight and start quickly makes them ideal for hosting microservices, which are a key element of cloud-native application architectures. Traditional, more monolithic applications can also be run inside containers, but will benefit less from this technology. As always, keep in mind that a poorly architected and designed application is still a poorly architected and designed application when ran in a container.

Takeaway: Containers enable running software that is more lightweight and efficient than past runtime environments have provided. IBM’s software offerings increasingly support containers as the standard runtime model and IBM Cloud Paks bring that software to market in an easily consumable and enterprise ready package.

Building production-ready images

All IBM container images provided in IBM Cloud Paks follow a set of well-defined best practices and guidelines, ensuring support for production use cases, and consistency across the IBM software portfolio. In addition to support for IBM Cloud Private and IBM Cloud Kubernetes service, IBM Cloud Paks support deployment to Red Hat OpenShift Container Platform using Red Hat Certified Containers.

One element that is especially important to IBM is support for multiple hardware architectures, including Linux on IBM Power and Linux on IBM LinuxOne, and providing images for the hardware platforms the respective IBM products support.

Management of security vulnerabilities is also critically important. IBM Cloud Paks are scanned regularly for known image vulnerabilities as part of the standard build procedures. All IBM Cloud Paks must have a documented process for managing newly identified vulnerabilities. Additionally, IBM follows [Secure Engineering Practices](#) for development of software and maintains a Security Vulnerability Management process (PSIRT) for commercial software supported by IBM. IBM Software delivered as an IBM Cloud Pak inherently follows those corporate standards. IBM Cloud Paks delivered by partners must have a documented process for addressing security image vulnerabilities.

Kubernetes – a management environment for containers

Up to this point, we have discussed the basics of building, running and maintaining container images, which can be used to run containers in a standalone fashion. But containers alone do not provide a framework for implementing production-grade qualities of service like resilience, scalability or maintenance.

For example, software running inside a container may write data to a file. If the file exists within the container, deleting the container will also delete the file. If the software's state must be maintained, that state data should be written to a volume outside of the container. If the state needs to be consistent even with the failure of a host, then that volume should exist on storage that is accessible by multiple hosts, most likely over a network. To maintain availability of the application during the failure of a host, you would also need to run multiple instances of the container

on multiple hosts and load balance incoming requests across those containers. This would require a reasonable amount of effort to manage manually, especially if you want to be able to seamlessly upgrade to newer versions of an application or build a continuous integration process.

Kubernetes is an open source orchestration platform for containers that solves these administrative challenges by providing a declarative framework for deploying, scaling, and managing container-based workloads. It is a popular choice for managing clusters of containers throughout the industry; IBM offers Kubernetes for customer-managed environments as part of IBM Cloud Private and as a managed service in the public cloud via the IBM Cloud Kubernetes Service.

The declarative definition of abstract resources that influence how the cluster behaves and manages workloads is a key feature of Kubernetes and will be covered briefly below. IBM Cloud Paks are built for Kubernetes-based environments and include all the configuration artifacts you need to easily customize and deploy an enterprise-grade Kubernetes workload.

Takeaway: Kubernetes is a popular framework for running containers in a scalable, resilient, highly available fashion, supporting production use cases for enterprise applications. IBM has chosen Kubernetes as its container orchestration platform both on-premises and in the cloud, and IBM Cloud Paks are designed specifically for deployment to IBM's Container Platforms

Kubernetes Resources

Kubernetes provides users with a set of defined resources including a way to describe how containers should run in the cluster, how the system reacts to events like failures, how to make containers accessible over the network and how and where to store data.

You can describe the provisioning and management of your application workload by defining the desired state of these resources using a YAML file and Kubernetes will manage the cluster environment accordingly.

Internally, Kubernetes delegates the management of the resource to its associated controller.

A few of the most common Kubernetes resources are described briefly below.

Deployment

Describes the desired state of one or more Pods, which are collections of running containers

StatefulSet

Similar to the Deployment resource mentioned above but describes containers that maintain state.

Service

Describes how pods that are part of a deployed workload (Deployment, StatefulSet, etc.) can be accessed from outside of the Kubernetes cluster. Gives clients a well-defined target address/port combination across multiple pods, including across restarts and recreations of these pods.

PersistentVolume / StorageClass

Enables you to define an allocation of storage that persists across the lifetime of the pods that use it. Pods can attach to a suitable volume by using a PersistentVolumeClaim. The StorageClass resource describes different qualities of service that are available for different types of storage that may be offered.

ConfigMap

Enables separating configuration data for a pod into a separate object.

Secret

Similar to ConfigMaps, Secrets contain sensitive data (for example, passwords or ssh keys) and are stored separately from container that use them.

This list barely scratches the surface of the resource types available in Kubernetes, which also supports defining custom resource types. For a more detailed description of Kubernetes resources, see the [official documentation](#).

The resource definitions mentioned above contain configuration metadata that is critical in ensuring enterprise-grade qualities of service of the workloads

running in Kubernetes. For example, you can define memory and CPU allocations for individual pods, ensuring that sufficient capacity is available when creating containers, while also ensuring that individual workloads cannot use more than their allocated resources, enabling effective sharing of hardware resources. As another example of the control afforded by Kubernetes, you can define affinity and anti-affinity rules that let you control which of your worker nodes certain pods run on.

Takeaway: Individual workloads, including IBM software content that runs in IBM Cloud Kubernetes Service or IBM Cloud Private, are described using predefined Kubernetes resources. IBM Cloud Paks define Kubernetes resources for your workloads using intelligent defaults, and provide for easy customization during deployment.

Using Helm charts to orchestrate containerized workloads

As mentioned above, Kubernetes uses abstract resources to allow describing the desired target state of a workload, paired with controller implementations that enforce the defined target state.

Each application or service running in Kubernetes is represented by multiple resources, each of which is typically defined in its own YAML file. Each resource also carries several attributes with it, whose values may differ from deployment to deployment based on the specifics of the environment and the supported usage.

The Helm project aims to simplify the deployment and maintenance of complex workloads in Kubernetes environments. It provides a packaging format called a chart, which you can use to group together YAML templates that define related sets of Kubernetes resources. An instance of a Helm chart that has been installed into a target Kubernetes cluster is called a release. Helm not only simplifies orchestration of Kubernetes resources, it also simplifies the ongoing maintenance of your releases. This makes production-level operations like rolling upgrades more manageable and contributes to the overall availability and maintainability of your application.

IBM Cloud Paks use Helm charts to describe how IBM software should be deployed in a Kubernetes environment. These resource definitions can be easily customized during deployment, and upgrades can be easily rolled out or rolled back.

The IBM and Red Hat Certification processes are complementary, and a Certified IBM Cloud Pak built with Red Hat Certified Containers is certified by both IBM and Red Hat for use on IBM's Kubernetes platforms and Red Hat OpenShift Container Platform.

Kubernetes Operators

[Operators](#) are flexible and powerful custom Kubernetes resource definitions that can be used for deploying and managing containerized workloads in a Kubernetes environment. They can also be used for packaging applications, in a manner similar to Helm charts, or they can be used together with Helm, in a complementary manner. By building specific knowledge and best practices about deploying and managing a software product directly into an operator, a software provider can capture domain-specific expertise about operating the product, giving end-users powerful automated runtime and lifecycle management capabilities without requiring that same level of expertise from the end user.

For example, IBM Cloud Paks can utilize operators to deliver IBM's expert knowledge about deploying and managing IBM enterprise software products in modern container orchestration environments as part of the software offering itself, transferring some of IBM's expertise to the customer automatically.

Takeaway: IBM Cloud Paks include Helm charts, which assemble all of the Kubernetes resource definitions related to a piece of IBM software, and provide for easy customization, deployment, and maintenance using IBM's Kubernetes platforms, on-premises or in the cloud, and can include Operators, which capture product-specific deployment and management expertise

IBM Cloud Paks

IBM Cloud Pak for Applications

To remain competitively relevant, enterprises much consistently update their software applications to meet the demands of their customers and users. Meeting this demand requires an application platform that allows for the quick building, testing and deployment in a modern, microservice-based architecture. To satisfy this crucial need, IBM is introducing IBM Cloud Pak for Applications.

IBM Cloud Pak for Applications supports your enterprise's application runtimes, and offer instrumental developer tools and modernization toolkits, DevOps, Apps/Ops Management and a self-service portal. IBM Cloud Paks for Applications can accelerate the ability to build cloud-native apps by leveraging built-in developer tools and processes, including support for microservices functions and serverless computing. Customers can leverage this Cloud Pak to quickly build apps on any

cloud, while also providing the most straightforward modernization path to the cloud for existing IBM WebSphere clients, with security, resiliency and scalability.

IBM Cloud Pak for Automation

Companies in nearly every industry are digitizing and automating their business operations. They're freeing employees from low-value tasks and assisting them with high-value work to drive a new wave of productivity, and customer and employee experiences. However, it can be challenging to effectively automate work at the pace of customer and internal expectations.

To address these challenges, IBM is introducing, IBM Cloud Pak for Automation is a pre-integrated set of essential software that enables you to easily design,

build and run intelligent automation applications at scale. With Cloud Pak for Automation, you deploy on your choice of clouds, anywhere Kubernetes is supported – with low-code tools for business users and real-time performance visibility for business managers. It's one flexible package with simple, consistent licensing. No vendor lock-in. And existing customers can migrate their automation runtimes without application changes or data migration.

IBM Cloud Pak for Data

As companies continue to harness the potential of AI, they need to use data from diverse sources, support best-in-class tools and frameworks, and run models across a variety of environments. However, 81% of business leaders do not understand the data required for AI. And even if they did, 80% of data is either inaccessible, untrusted, or unanalyzed. Simply put, there's no AI without an information architecture.

IBM recognizes this challenge our clients are facing. As a result, IBM is introducing IBM Cloud Pak for Data with the goal of creating a prescriptive approach to accelerate the journey to AI: the AI Ladder, developed to help a client drive digital transformation in their business, no matter where they are on their journey. IBM Cloud Pak brings together all the critical cloud, data and AI capabilities as containerized microservices to deliver the AI Ladder within one unified multicloud platform.

IBM Cloud Pak for Integration

Traditional integration approaches cannot cope with the volume and pace of business innovation. Digital transformation enables organizations to unlock the power of data to create personalized customer experiences, utilize artificial intelligence, and innovate faster to stay ahead of the competition. In order to keep up, businesses need the ability to integrate in hybrid environments outside the data center and drive speed and efficiency in integration development while lowering costs. To facilitate these new, evolving demands, IBM is introducing IBM Cloud Pak for Integration.

IBM Cloud Pak for Integration is designed to support the scale, security and flexibility required to empower your digital transformation. With the Cloud Pak, enterprises can integrate across multiple clouds with a container-based platform that can be deployed across any on-premise or Kubernetes cloud environment, and easily connect applications, services, and data with the right mix of integration styles, spanning API lifecycle management, application integration, enterprise messaging, event streams, and high-speed data transfer.

Enable your business to set up the appropriate organizational models and governance practices to support a modern agile approach to integration with IBM Cloud Pak for Integration.

IBM Cloud Pak for Multicloud Management

As application innovation accelerates, enterprises have increasingly adopted a hybrid, multicloud architecture to build, test and deploy applications. With this new hybrid, multicloud architecture, the volume and complexity of objects and metrics to manage has skyrocketed, making monitoring and securing the enterprise IT ecosystem more difficult. To mitigate some of this complexity, IBM is introducing IBM Cloud Pak for Multicloud Management.

IBM Cloud Pak for Multicloud Management provides consistent visibility, automation, and governance across a range of multicloud management capabilities such as cost and asset management, infrastructure management, application management, multi-cluster management, edge management, and integration with existing tools and processes. Customers can leverage IBM Cloud Pak for Multicloud Management to simplify their IT and application ops management, while increasing flexibility and cost savings with intelligent data analysis driven by predictive signals.

Summary

IBM Cloud Paks provide an easy and powerful way to run high-quality, container-based enterprise software on a modern Kubernetes-based orchestration platform that enables high availability, scalability, and ongoing maintenance for enterprise applications, from a source you know and trust. They include container images that are built and tested by product teams, capturing product expertise and best practices in a form factor that is easy to consume and deploy in a location of your choice, on-premises or in the cloud. Images provided by IBM are regularly scanned for known security vulnerabilities and follow a rigorous process for managing newly identified issues.

IBM Cloud Paks also include pre-configured Helm charts that describe runtime environments for IBM software products based on established best practices and can be easily customized during the deployment process. They may also include Operators that build product-specific deployment and lifecycle management expertise into the software. These capabilities combine to provide a first-class deployment experience, integration with core platform services, and production-ready qualities of service. Certified IBM Cloud Paks built with Red Hat Certified Containers build the combined expertise of IBM and Red Hat into trusted enterprise software solutions that combine fast, easy deployment with enterprise qualities of service and simplified, flexible pricing.

The new family of Cloud Paks from IBM, including IBM Cloud Pak for Applications, IBM Cloud Pak for Data, IBM Cloud Pak for Integration, IBM Cloud Pak for Multicloud Management, and IBM Cloud Pak for Automation give customers the capabilities they need to bring the next 80 percent of their workloads to modern, cloud-based environments.

Resources

IBM Cloud Private

<https://www.ibm.com/cloud/private>

IBM Cloud Kubernetes Service

<https://www.ibm.com/cloud/container-service>

What is Docker?

<https://www.ibm.com/cloud/blog/new-builders/kubernetes-vs-docker-its-not-an-either-or-question>

Microservices

<https://www.ibm.com/cloud/blog/new-builders/video-what-are-microservices>

Open Container Initiative

<https://www.opencontainers.org/>

Kubernetes Concepts

<https://www.ibm.com/cloud/blog/new-builders/video-kubernetes-explained>