White Paper

Building the Open Hybrid Cloud: Red Hat OpenShift, Enterprise Linux, and Ansible on IBM Z, LinuxONE, and Storage

Sponsored by: IBM

Gary Chen   Al Gillen   Lucas Mearian
August 2021

IDC OPINION

In an industry that is accustomed to change, we have entered a period of rapid evolution driven by unprecedented disruption. The result is waves of new technology that fundamentally redefine what modern applications are, and as a consequence of the use of new technology, the creation of a competitive environment where new threats from many vectors can emerge suddenly. No longer are incremental improvements to applications and IT systems enough to retain a competitive advantage. It is now necessary to "disrupt or be disrupted," both internally and externally.

Enterprises are waking up to this new reality: their existing application portfolio, which has successfully empowered the organization to this point, can also become a maintenance liability due to the inability of these applications to evolve and change as quickly as they may need to.

To achieve a more agile and responsive IT environment — one that also offers longer-term investment protection — many organizations are now looking to modern development paradigms. These new paradigms use efficient modern languages to build modular applications and package the applications in containers so the deployment options become more flexible and far more numerous in number.

Of course, change never happens in isolation. In reality, organizations need to make a smooth transition from historical investments in their datacenters to a flexible, cloud-native environment that supports their new application needs.

For most enterprise customers, the ability to take advantage of modern development and deployment solutions while still using existing technology investments is an optimal outcome, and today, it has become practical reality. The emergence of native Kubernetes orchestration systems for virtually every architecture in use across the industry has cleared the way for containerized applications to benefit from deployment flexibility. In effect, such applications can be deployed on both generic cloud infrastructure and uniquely capable hardware architectures leveraging specialized processors and storage systems.

The industry has given enterprise customers the platforms as well as the tools and resources to build solutions that can span on-premises deployments with off-premises deployments, bridge IoT devices with diverse architectures into back-end systems collecting data from those devices, and abstract multiple architectural environments to a common set of applications.

SITUATION OVERVIEW

It would have been hard to imagine how important open source software (OSS) would become 20 years ago, when Linux was just beginning to establish enterprise credibility. However, today, open source software has become not only essential but also the model where new technologies are invented and developed, often without being constrained by backward compatibility. It is now surprisingly rare to see a new proprietary technology emerge from behind closed doors and rock the industry the way that OSS projects such as Kubernetes and the Hyperledger project have in recent years.

OSS has reached a level of extreme credibility and wide adoption, equal to nearly any tightly controlled proprietary technology. This means that systems at all levels of compute — from personal devices to deeply rooted systems of record — have embraced open source solutions at practically every level of the software stack. Today, the compelling combination of Linux, Kubernetes, and OSS management tools in conjunction with a well-defined set of DevOps practices makes it possible for organizations to modernize in place, with the platforms they already use.

There are two basic forms of open source software: community-supported technologies and commercially supported products. The commercial products are built from the community-supported technologies, but the process of creating a commercial distribution gains the benefit of additional enterprise-focused testing, security hardening, tooling, certification, and integration along with enterprise-grade support services.

Enterprise Computing and Storage Systems

Enterprise customers typically are among the largest and most sophisticated organizations. They usually have clear expectations and requirements for their computing platforms and systems. Enterprises typically are willing to pay for high-quality solutions that offer scalability, reliability, predictability, and a measurable return on investment. Accordingly, enterprise customers hold their suppliers to high standards for product support, product life cycles, performance, reliability, and security.

Red Hat's Commercial OSS Portfolio

Red Hat is arguably the best-known company that curates open source software and builds commercially supported products that deliver the attributes that enterprises demand. Some of those attributes are:

- Subscription-based contracts offering fixed pricing and robust support services
- Cloud-based subscription models offering flexible consumption-based pricing
- Commercial incident support available in varying degrees of responsiveness at different price points
- Predictable product support life cycles with published dates, including termination of support for aging products to guide deployment planning (Many products are supported for extended periods of times that enterprises require, often 10+ years in length.)
- Certification process for application software qualified to operate properly on Red Hat's infrastructure software stack
- Certification process for supporting hardware products
- Certification process for supporting public clouds
- Professional certifications for IT staff
- Consulting services, training, and jump start lab services

Red Hat products are typically available on the most popular platforms in the industry, with the x86 servers being the most widely deployed hardware architecture. However, Red Hat also supports IBM Z, IBM LinuxONE, IBM Power, and IBM Storage. In addition, Red Hat began offering support for Arm-based servers in 2017.

Red Hat Solutions

Red Hat software is widely used in the industry. Red Hat first offered a Linux operating system called Red Hat Linux during the 1990s. The company evolved that product into an enterprise-grade solution and renamed it Red Hat Enterprise Linux (RHEL) in 2002. Since that time, Red Hat Enterprise Linux has become increasingly sophisticated, scalable, and available on additional architectures and underpinning IBM Storage solutions.

Major product additions that have expanded the value of Red Hat's overall software story include:

- Red Hat Application Services portfolio of middleware products (Red Hat Integration, Red Hat Runtimes, and Red Hat Process Automation)
- Middleware (Red Hat JBoss, Red Hat Middleware, and Red Hat Integration)
- Virtualization (Red Hat Virtualization)
- Private IaaS cloud (Red Hat OpenStack Platform)
- Container platform (Red Hat OpenShift)
- Application development (Red Hat CodeReady Studio and Red Hat CodeReady Workspaces)
- Storage (Red Hat Ceph Storage and Red Hat OpenShift Data Foundation)
- Container security (Red Hat/StackRox)
- Automation and management (Red Hat Ansible Automation Platform, Red Hat Insights, Red Hat Satellite, and Red Hat Advanced Cluster Management)

In recent years, Red Hat has been investing heavily in expanding its OpenShift product family; embracing Kubernetes container orchestration; and adding cloud deployment and support options, Kubernetes Operators, service mesh, and serverless. In addition, with Red Hat OpenShift 4, Red Hat has refined the deployment strategy through the inclusion of a lightweight Linux Kernel (Red Hat Enterprise Linux CoreOS) as part of the OpenShift installation, deployment, and updating experiences. Red Hat recently expanded its platform services designed to extend the OpenShift value proposition to include data streaming, more artificial intelligence (AI) capabilities, and leveraging its API management offering.

IBM and Red Hat

IBM's 2019 acquisition of Red Hat highlighted how important it was for IBM to have a leadership position in the hybrid cloud market and to have a fully mature support model, including cross-platform capabilities in critical open source technologies such as Linux and Kubernetes. While there were overlaps in Red Hat's portfolio, IBM did not have a top-to-bottom solution for a Linux container platform, including the underlying operating system layer. Linux plays a key foundational role in containers and Kubernetes, and thus, Red Hat OpenShift heavily leverages the existing and future work that Red Hat has invested into and continues to invest into Red Hat Enterprise Linux.

Containers and Kubernetes

The notion of a containerized application is not new and has been available in the industry in various forms for over 20 years – although the use cases were narrow. However, once Docker (the company and the technology) emerged, it created a compelling and user-friendly way to package software and run modern applications.

The Open Container Initiative (OCI) defined the formatting for containers that are delivered in a consistent manner, making it practical to package an application and all its dependencies into a single portable container image. These images are then shared on centralized container registries where they can be version tracked and iterated on by other developers or pushed to production. When containers are executed, they run in their own sandbox so that each container is isolated from other containers and the host OS. Developer-friendly tools and APIs made building, sharing, and running containers easy, and thus, containers quickly became popular.

To solve the challenges of running complex distributed applications at scale, Google took the concepts and learnings from its internal "Borg" container orchestration technology and created the open source software Kubernetes project with active participation from IBM, Red Hat, Microsoft, and many others. Governance of the Kubernetes project was handed over to the newly formed Cloud Native Computing Foundation (CNCF), and the industry quickly embraced Kubernetes as the preferred orchestration platform for containerized applications.

Along with the standardization of container formats and runtimes made possible by OCI, the industry now had a common container stack that would create consistency, interoperability, and portability between the various container platforms and container services. Other emerging areas in the container ecosystem, such as Istio for service mesh and Knative for serverless, are all open source and they are rapidly building communities and may become as ubiquitous as Kubernetes.

Containers and Kubernetes in Use

Containers are a flexible packaging technology that today are seeing broad use to deploy cloud-native and microservices-based applications on Kubernetes. For developers, containers are a perfect solution to efficiently encapsulate new cloud-native microservices and to push this software down increasingly automated software build pipelines using continuous integration/continuous delivery (CI/CD).

However, containers are not used just for new applications. IDC data shows that roughly half of enterprise containers are applications that are migrated from virtual machines (VMs) or bare metal servers, often with some level of refactoring. Developer-friendly APIs make working with complex software faster and more convenient, improving developer workflows. Containers can also improve code quality by helping enable both more automated testing and better environmental control as containers typically include all dependencies. Ultimately, containers enable faster development of software, faster deployment of changes, and more developer productivity.

The Benefits of Containers and Kubernetes

For IT operators, containers and Kubernetes offer a modern, highly scalable, and automated way of running large web-scale applications. Kubernetes and Kubernetes Operators patterns embed much of the knowledge and experience of web-based companies necessary to run fast-changing apps reliably at a very large scale. The combination of Kubernetes and Kubernetes Operators has established deployment patterns such as blue/green upgrades, A/B testing for new app features, and multiple automated scaling options.

In addition, container-style deployments help solve configuration headaches by leveraging the concept of immutable infrastructure. This means that a container's configuration is defined in its image, and it never changes during runtime. If a change is desired, the old container instances are brought down and a new image is started instead of patching or making configuration changes to a running instance. Container repositories also help centralize container images and maintain versioning. The lightweight and reactive nature of containers, combined with a modern control plane, enables IT to efficiently deploy and manage modern applications.

Modern applications have embraced the concept of statelessness, where no data persists between sessions. This simplifies the design and deployment of the application, particularly in clouds where servers are treated as commodity and any particular instance is not meant to be highly available. Stateless applications are also much simpler to scale. While modern applications strive to be stateless where possible, nearly all applications store state at some point. These stateful parts of applications, such as databases, were often not containerized in the early days of containers, but this has changed considerably as Kubernetes and other projects have addressed the need for stateful container functionality. For example, Kubernetes has added support for persistent storage with the Container Storage Interface (CSI) and automated the complexities of stateful applications with Kubernetes Operators. Today, enterprises are moving to containerize all parts of their applications, stateless and stateful, and they will require container platforms to support both types.

As customers move to hybrid cloud and multicloud deployments, containers can play a key role in portability and consistency across different environments. As previously discussed, the broad adoption of the OCI format and the Kubernetes control plane by the IT industry inherently makes container platforms similar at their core. The CNCF provides a Kubernetes conformance test and certification, which means that all Kubernetes installations must behave the same for core functions. As a result, customers can use virtually any stable Kubernetes product and expect a reasonable level of compatibility.

Developers can largely work with containers and Kubernetes using the exact same APIs and tooling that they wish, no matter the distribution or cloud service or infrastructure supporting the deployment. This can help provide a consistent developer environment both on premises and across various public clouds.

In addition, containers can help abstract differences in system architectures as the container interfaces remain the same regardless of the hardware underneath. As a result, the developer does not have to know the ins and outs of the platform or the operating system to develop applications. For IT operators who manage Kubernetes, there are some system-specific knowledge and deployment tools they will have to install and integrate, but operating Kubernetes and the apps managed by it remains largely consistent across any distribution.

Red Hat OpenShift and IBM Platforms

One of the benefits of a Red Hat OpenShift deployment is that it offers a portable environment, such that customers can deploy on OpenShift consistently on premises, off premises in public cloud, or on multiple hardware, virtualized, or cloud platforms.

Platform independence is enhanced by the use of interpreted languages such as JavaScript, Perl, and Python, as well as through bytecode-compiled languages such as Java. Interpreters abstract away architecture and endian concerns via bytecode representation of the application. The bytecodes then run on an interpreter that understands how to translate the platform-agnostic bytecode down to the specific machine architecture that the interpreter is executing on. Interpreters also use just-in-time compilers to further optimize the application performance.

By comparison, languages that undergo a static compilation prior to execution, such as C, C++, Go, and Haskell, will require that compiler to create binary code for all platform architecture choices where the application might be deployed. Container technology itself does not natively provide architectural abstraction, so binary applications in containers compiled for x86 Linux platforms may not be executable by Linux on IBM Z systems.

However, OCI container images can contain multiple binaries targeted for multiple architectures within a single image, allowing the creation of universal images that can be deployable on multiple hardware platforms, with that deployment managed by the CI/CD environment.

There are benefits associated with every platform. In the case of x86 platforms, there is ubiquity and availability in public cloud environments. IBM's traditional enterprise platforms, including IBM Z, IBM LinuxONE, IBM Power, and IBM Storage, offer benefits that extend to a Red Hat OpenShift deployment in many ways as a secure, resilient, scalable, and elastic cloud in a box for cloud-native workloads. Benefits of IBM platforms include:

- **High levels of scalability:** IBM's Z and LinuxONE platforms scale out to 85 isolated logical partitions (LPARs) that offer Common Criteria for Evaluation Assurance Level 5+ (EAL5+) security, hosting thousands of Linux guests and millions of containers within a single physical footprint.
- **Reliability:**
  - IBM platforms are designed to offer 99.999% uptime and beyond.
  - IBM Z and IBM LinuxONE are developed to adapt to planned or unplanned events while keeping services and operations running continuously.
  - The IBM Geographically Dispersed Parallel Sysplex (GDPS) offerings deliver resiliency capabilities for customers that run Linux guests on IBM z/VM, providing high availability and disaster recovery benefits in case of system, application, or network failure.
  - In addition, IBM Z uses redundant array of independent memory (RAIM) technology (similar to RAID storage used for disk), which goes beyond error checking to deliver high levels of resiliency by recovering from even the most catastrophic memory failures.
- **Security:**
  - **Access to FIPS 140-2 Level 4–compliant hardware security module and industry-leading trusted execution environment (TEE).**
  - **Highly isolated and secure virtualization options.** Logical partitions essentially carve up a Z system into separate servers using a firmware-based hypervisor, providing resource isolation at EAL 5+ certification levels, which means that partitions can be configured in a way that they match the air-gap isolation of physical servers. The IBM z/VM or a KVM hypervisor can be run within partitions to further subdivide the system.
  - **Cyber-resilience.** IBM Z Cyber Vault provides air-gapped data corruption protection and tools delivered by combining security and protection capabilities of IBM Z and IBM DS8900F. This solution keeps hundreds of immutable, safeguarded copies of data in a logical partition that is isolated and separated from the production system. The air-gapped data vault prevents reloading corrupt or compromised copies of data, which is especially important in mission-critical containerized environments to reduce the negative impact of adverse cyberevents.
  - **Security.** IBM z15 supports pervasive encryption with the ability to encrypt data at rest and in flight, including volumes for z/VM and KVM virtual machines. In addition, IBM Fibre Channel Endpoint Security extends the value of pervasive encryption by ensuring all data flowing on FICON and Fibre Channel Protocol (FCP) links from IBM Z to DS8900F is encrypted and protected. Hardware acceleration of encryption keeps the performance overhead of encryption low.

- **Performance:** IBM Z and IBM LinuxONE support high I/O bandwidth workloads well with a dedicated I/O subsystem that provides high throughput, and z15 integrates file compression capabilities with an on-chip compression coprocessor. In addition, IBM Z's scheduling and workload management enable efficient use of all computing resources even at high utilization levels. For those environments where the speed to access data is the most important, IBM Z, in combination with IBM DS8900F storage systems, provides fast response time and low latency.
- **Expands the developer base that can write to IBM traditional systems:** The availability of the Red Hat OpenShift on IBM systems means developers don't need unique skills specific to the IBM environment. Red Hat OpenShift provides abstraction and a consistent user experience such that developers only need to be comfortable with deploying containers on Red Hat OpenShift.
- **Extends IBM Z and IBM LinuxONE participation in the open hybrid cloud world.** Today, most platforms will need to support hybrid cloud environments if intended to be a deployment option for modern cloud-native, containerized applications. Having Red Hat OpenShift on IBM systems means that the portability, agility, and scalability attributes of a Kubernetes platform apply to the IBM systems as well. Because OpenShift runs wherever RHEL is certified to run, OpenShift helps extend the reach of IBM Z and IBM LinuxONE in hybrid and multicloud environments. For other customers, the use of Red Hat OpenShift on IBM Z and IBM LinuxONE using the IBM z/OS Cloud Broker and IBM z/OS Connect Enterprise Edition brings fast and easy access to existing data stores for modern applications.

z/OS Self-Service Access and Integration with OpenShift

Customers are increasingly confronted by a hybrid cloud reality, making the ability to integrate all core systems of record, along with their data and applications, imperative for success. This includes the mission-critical z/OS assets in customer shops. IBM provides the ability to connect and drive on-demand access and management to z/OS from OpenShift with the IBM z/OS Cloud Broker. This operator-based framework provides OpenShift users catalog access to z/OS services for application modernization. IBM z/OS Cloud Broker automates this process, enabling clients to be more responsive and provision complex environments perfectly each time.

The IBM z/OS Cloud Broker also provides an extensible automation framework to z/OS applications and systems. Users can self-provision the services they require and can drive additional actions to manage the development life cycle of services and applications deployed with the IBM z/OS Cloud Broker. This capability supports enterprisewide automation via the Red Hat Ansible Automation platform on OpenShift.

Development for Red Hat OpenShift on IBM Systems

Red Hat offers development environments including IBM Red Hat CodeReady Workspaces, an Eclipse Che–based integrated development environment that is optimized for development and deployment onto the Red Hat OpenShift platform.

IBM recently added IBM Wazi Developer for Red Hat CodeReady Workspaces to the DevOps toolchain for IBM Z. IBM Wazi Developer empowers developers by providing a consistent and familiar development experience for IBM z/OS that is optimized to run on Red Hat OpenShift. Because it is based on Red Hat CodeReady Workspaces, this approach removes barriers associated with z/OS development through key capabilities, including IBM Wazi Code, IBM Wazi Sandbox, and IBM Wazi Analyze.

IBM Wazi Code provides developers with a choice of IDEs for development with intelligent editors and build tools and a debugger with rich language support for REXX, COBOL, PL1, and Assembler through Eclipse, Visual Studio Code (VS Code), and Red Hat CodeReady Workspaces. Developers can use modern software configuration management (SCM) tools such as Git to embrace true parallel development. Developers can then build with IBM Dependency Based Build and accelerate delivery time.

With Wazi Sandbox, developers are provided a containerized, personal z/OS sandbox environment running on Red Hat OpenShift. This provides agile, flexible sandbox environments for development, test, and education that can be spun up and down on demand. IBM Wazi Analyze enables developers to increase application understanding and help assess the impact of code changes before any changes are made, thus reducing risk.

As a result of these three aforementioned key components in IBM Wazi Developer for Red Hat CodeReady Workspaces, developers can accelerate hybrid application development and delivery through intelligent application analysis, choices in IDE, and personal z/OS sandbox. This acceleration of development and delivery increases agility, improves reliability, and shortens releases.

IBM Cloud Paks

In addition, IBM ported and enhanced its Cloud Paks product suite so that these capabilities can be installed and utilized directly on Red Hat OpenShift. IBM has been expanding its Cloud Paks portfolio with products designed to take advantage of common services (called the Bedrock layer) and the IBM Automation Framework. These two foundational layers work in concert to provide consistent services and automation to the following products:

- **IBM Cloud Pak for Data.** Unifies and simplifies the collection, organization, and analysis of data. In the future, IBM Cloud Pak for Data will add capabilities to build, deploy, and manage analytic and AI models with trust and transparency.
- **IBM Cloud Pak for Business Automation.** Automates business operations to achieve better performance.
- **IBM Cloud Pak for Watson AIOps.** Provides application impact avoidance and hybrid application management and observability; brings automation and AI into IT management to proactively prevent incidents.
- **IBM Cloud Pak for Integration.** Automates application and data flows to improve client experiences.
- **IBM Cloud Pak for Network Automation.** Automates networks to deliver zero-touch operations.
- **IBM Cloud Pak for Security.** Generates deeper insights into threats, orchestrates actions, and automates responses.

IBM Storage for Red Hat OpenShift

Containers are natively stateless. Containerized applications as a whole, however, need persistent storage in order to be stateful. Those containerized applications also produce data that must be stored, and for that, persistent storage is needed.

For a container to be stateful, it must map to a persistent storage resource that saves the data related to workloads. There are nuances involved when talking about stateful and stateless, but the basic premise is this: Stateless containers are like volatile memory — they are useful only if their processes remain alive in the runtime. Once they are destroyed, the state of that application and the data it has created are lost. Stateful containers are capable of saving data to primary (server) and external storage, whether it's block, file, or object. Stateful container environments, however, need a way to connect to storage, and that involves either a cloud-native storage application or a plug-in known as the Container Storage Interface (CSI).

Thanks to the CSI, storage providers can now write just one driver that will work with Kubernetes or any other orchestration platform. CSI, however, only works for localized storage, and while that allows containerized applications a stateful environment, it doesn't maximize application and data mobility in a hybrid cloud environment.

Conversely, a cloud-native storage application is packaged in containers and deployed as microservices. As a containerized platform, these cloud-native storage applications are high performing and extremely scalable and portable, meaning they can move between clusters onsite and between a private cloud or a hybrid cloud environment.

Container clusters also must be managed, and that's where container orchestration engines come in. While there are several container orchestration engines, Kubernetes has emerged as the de facto platform. Kubernetes will move containers around a cluster to ensure they have the infrastructure resources they need (i.e., CPUs, memory, storage) as workloads change. DevOps teams also most often create containerized applications in a public cloud, test them there, and then bring them back on premises to run in production.

OpenShift is a container orchestrator based on Kubernetes that has a native CSI plug-in and also offers critical data services, such as backup, data migration, and disaster recovery capabilities. OpenShift is one of two paths IBM has created to enable persistent storage for Kubernetes environments.

Red Hat OpenShift Data Foundation

Red Hat OpenShift Data Foundation (ODF), previously Red Hat OpenShift Container Storage (OCS), delivers data persistence for containerized applications. Red Hat ODF is based upon Red Hat Ceph, which is a full-featured, unified, software-defined storage platform that delivers massive scalability and performance for demanding data workloads. Red Hat ODF enables a hybrid cloud experience where applications and data can be deployed across clouds and across existing legacy storage platforms.

IBM DS8000 Storage Systems

Through OpenShift and its CSI interface, containerized applications can access the IBM DS8000 storage system. This enterprise storage solution combines a Kubernetes orchestrator with a more traditional platform, which provides the performance, data resiliency, and security to manage container-based workloads in mission-critical production environments. Built with IBM's most advanced processor technology, the DS8000 is designed to match the capabilities of the IBM Z and LinuxONE server platforms and enables a storage foundation to consolidate traditional transitional processes and containerized workloads into a single, highly available storage offering.

IBM FlashSystem Family

Built with IBM Spectrum Virtualize, the IBM FlashSystem family is a portfolio of cloud-enabled storage systems designed to be easily deployed and quickly scaled to help optimize storage configurations, streamline issue resolution, and lower storage costs. These offerings support Red Hat OpenShift and Kubernetes container environments, accelerating the deployment of persistent volumes with the IBM storage CSI driver, certified by Red Hat and IBM.

IBM Spectrum Scale

With IBM Spectrum Scale container-native storage, DevOps or storage admins can more easily configure storage for containers. Application performance and storage resources can be optimized as a Spectrum Scale container. Native storage can be deployed for many containers or a container can have a dedicated storage interface to control the parallel access to the data.

This provides higher performance with less server resources and leverages the high parallel throughput of IBM Spectrum Scale and Elastic Storage System nodes. There are no capacity limitations on the local system, and performance is not affected by the application server but is optimized by external storage resources. IBM Spectrum Scale provides the benefits of container-native storage without the limits of using only local disk and server resources. IBM Spectrum Scale increases business agility for applications and helps eliminate data silos with a global data mesh that enables applications to access both file and object resources on the edge, in the core, or even in the cloud.

IBM Storage Suite for IBM Cloud Paks

IBM Storage Suite for Cloud Paks is a set of flexible software-defined storage solutions from the IBM and Red Hat storage portfolios specifically designed to support container environments while delivering the data services the Cloud Paks need. With simplified access to data, it provides a consistent experience for the OpenShift environments. What's more, these data services are available to IBM Z and LinuxONE.

IBM Storage Suite for Cloud Paks is integrated with the automation capabilities of Kubernetes, Red Hat OpenShift, and IBM Cloud Paks. As a result, it enables IT infrastructure and operations to improve developer speed and productivity.

Red Hat Ansible Automation Platform for IBM Systems

IT automation has become a strategic imperative for enterprises across DevOps, hybrid cloud, and edge. Ansible is a popular open source automation language that allows systems, DevOps, and cloud administrators to standardize and scale configuration, deployment, and operational orchestration activities with a human-readable, agentless approach to infrastructure as code. The Red Hat Ansible Automation Platform is a fully supported, integrated foundation for developing, managing, and scaling Ansible's core automation code building blocks, generally known as playbooks, modules, and roles.

The Red Hat Ansible Automation Platform includes all of the tooling needed for building, deploying, and managing end-to-end automation at scale. Ansible Automation Platform makes it possible for users across an organization to share, vet, and manage automation content by means of a simple, powerful, and agentless technical implementation. IT managers can provide guidelines on how automation is applied to individual teams. Meanwhile, automation developers retain the freedom to write tasks that use existing knowledge, without the operational overhead of conforming to complex tools and frameworks. It is a more secure and stable foundation for deploying end-to-end automation solutions, from hybrid cloud to the edge.

The Red Hat Ansible Automation Platform includes a graphical user interface (GUI) to manage complex automation workflows and role-based access controls. The Automation Platform also provides API enabling Ansible automation with a wide variety of product integrations via Red Hat Ansible Certified Collections.

Red Hat Ansible Certified Content for IBM Z was made available as part of the Ansible Certified Content Collections on Ansible Automation Hub in 1Q20. This certified and supported content provides the core automation artifacts necessary to quickly construct, test, and deploy standardized and reusable automation procedures for common tasks on z/OS such as:

- Job submissions and queries
- Data set management
- Application deployment with CICS, IMS, and UrbanCode Deploy
- Orchestration of existing automation with z/OSMF and IBM Z System Automation
- Management of Hardware Management Console resources with Ansible
- Automation of IPL or LPAR provisioning processes

Today, customers can run and use the base capabilities of Ansible to manage Linux on IBM Z. This means that they can use most existing and off-the-shelf Ansible automation with the platform to support a wide variety of automation use cases from day one deployment to day two operations.

Red Hat Ansible Automation is widely adopted across many operational domains including server, storage, network, cloud, DevOps, and security configuration, orchestration, and operations. IDC's research shows that productivity improvements for IT staff can be as high as 68% when Ansible Automation is introduced to specific operational domains. When applied across multiple domains, the benefits are even more dramatic as IT staff is able to seamlessly integrate complex workflows and more easily share and reuse automation resources.

Many Red Hat and IBM customers rely on Ansible Automation Platform to create OpenShift operators and to integrate OpenShift operations into CI/CD tool chains. With the z/OS Cloud Broker running on OpenShift, users will have an operator framework leveraging Ansible Automation Platform to perform a wide number of repeatable operations on z/OS from an OpenShift-based environment.

In addition, network and security operations teams rely on Ansible to manage and configure a wide range of network devices and remediate security threats and incidents. Systems, storage, and cloud managers build complex configuration workflows, updates, and application deployment automations. By implementing Ansible Automation for z/OS, large enterprises can create more consistent, integrated, and efficient automation environments that reduce human error and help the organization rapidly scale and adapt to digital business innovations that rely on data and resources running on traditional and cloud-native platforms.

The Red Hat OpenShift Ecosystem

As enterprises accelerate their application modernization efforts, they are increasingly leveraging best-in-class third-party solutions for specific workload needs with the goal of increasing operational efficiencies and reducing time to market. Today, most ISVs are building new applications (and repackaging existing applications) on the most common deployment platform, which is often a Kubernetes-based environment. Accordingly, these modern applications can usually be deployed on OpenShift, including on IBM Z or IBM LinuxONE, given availability of the right binaries.

IBM has a diverse and evolving ecosystem of ISV applications across different industries that are based on OpenShift for cloud-native development. Examples of IBM's growing ecosystem includes:

- Core banking applications such as Temenos Transact, Technisys Cyberbank Core, and Infosys Finacle

- Digital banking engagement applications such as Backbase
- Relationship-based product, pricing, and billing applications such as SunTec Xelerate and Zafin
- Payment modernization applications such as Vestigo Issuing, Fiorano Open Banking, and ISO 20022 messaging solutions
- Database management applications such as Fujitsu Enterprise Postgres

FUTURE OUTLOOK

It is unusual for the industry to coalesce around a de facto standard as quickly as has happened for containers and Kubernetes runtime environments. However, the industry has clearly locked on to this platform as the deployment vehicle for modern applications and as a target for lift and shift modernization efforts.

The industry has settled on Kubernetes for the container control plane; hence, Kubernetes becomes the nexus for innovation and integration, leading to multiple benefits:

- First, Kubernetes itself will benefit from a massive community that is focused on one platform.
- Second, there is a growing collection of innovative projects that are not part of Kubernetes per se but integrate with and optimize for Kubernetes, extending the capabilities and benefits of Kubernetes orchestration environments.

Examples of closely related technologies include Knative (serverless), Istio (service mesh supporting Kubernetes), Prometheus (real-time monitoring), Helm (packaging manager for Kubernetes), and OpenTelemetry (instrumentation for distributed tracing on Kubernetes). It stands to reason that some of the more interesting system software developments that happen over the next few years will take place in or around the Kubernetes infrastructure layers.

CHALLENGES/OPPORTUNITIES

- **Challenge:** Default IT infrastructure biases favor x86 servers.

  Opportunity: The industry is shifting to focus on deployment to a Kubernetes environment, not on the underlying architecture that is supporting that environment. This approach has made it possible for customers to write code that deploys in cloud environments, on edge devices, and on traditional enterprise IT infrastructure such as the IBM Z. Red Hat OpenShift expands the attraction for all platforms where OpenShift is available. The portability that OpenShift enables reduces concerns for vendor lock-in and, at the same time, expands the opportunity for recruiting developers since the Kubernetes runtime environment is CNCF certified and behaves as other Kubernetes environments will.
- **Challenge:** Applications on IBM Z and IBM LinuxONE are not binary compatible with x86 systems.

  Opportunity: Today's modern development tools and CI/CD environments can manage a diverse set of deployment platform options. When combined with Red Hat OpenShift, application compatibility becomes relatively simple, since the deployment environment differs only by the binaries that are deployed. Having modern applications running on IBM Z means extremely low latency to access data sources also residing on those platforms.

CONCLUSION

The industry is deeply engaged in an era of accelerated innovation, and to support the increased demand, customers need to solve deployment and portability challenges that abstract the underlying environment — whether on premises, off premises/cloud, or a multi-architecture scenario — so developers can focus on building compelling and valuable applications.

The standardization of infrastructure and deployment software layers means that customers can — and should — focus more of their resources on creating differentiation for their applications, including optimizing user experience, providing support for multiple user devices, and enhancing the functionality/completeness of the application itself.

The combination of a Red Hat software stack along with IBM Z and IBM LinuxONE brings the best of both worlds, where:

- Red Hat delivers a cross-platform deployment and runtime environment that allows customers portability and flexibility.
- IBM's platforms offer a clear benefit in terms of scale, flexibility, reliability, security, and low-latency access to massive corporate data sets for applications that benefit and thrive in this deployment environment.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com