White Paper

Building the Open Enterprise: Leveraging Red Hat Enterprise Linux, OpenShift and Ansible on IBM Z

Sponsored by: IBM

Gary Chen Al Gillen Mary Johnston Turner
April 2020

IDC OPINION

In an industry that is used to change, we have entered a period of rapid evolution driven by unprecedented disruption. The result is waves of new technology that fundamentally re-define 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 a 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 now need them 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, using efficient modern languages to build modular applications, and packaging those 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, customers 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 onto uniquely capable hardware architectures and leverage specialty processors.

The industry has given enterprise customers the tools and resources to build solutions that can span on-premise deployments with off-premise 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 not only become essential; it has become 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.

That OSS has reached a level of extreme credibility, equal to nearly any tightly-controlled proprietary technology 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, and integration along with enterprise-grade support services.

Enterprise Computing

Enterprise customers typically are among the largest and most sophisticated organizations, and usually have clear expectations and requirements for their computing platforms and systems. Enterprises typically are willing to pay for high quality solutions that offer reliability, predictability, and a measurable return on investment. Accordingly, enterprise customers hold their suppliers to high standards for product support, product lifecycles, 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 a commercially supported product that delivers the attributes that enterprises demand. Some of those attributes include:

- Subscription-based contracts offering fixed pricing and robust support services
- Commercial incident support available in varying degrees of responsiveness at different price points
- Predictable product support lifecycles with published dates for termination of support for aging products. Many products are supported for extended 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 jumpstart 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 the IBM Power, IBM Z and IBM LinuxONE platforms. 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 in 2002. Since that time, Red Hat Enterprise Linux has become increasingly sophisticated, scalable and available on more platforms including IBM Z, as Red Hat continued to expand its surrounding software portfolio.

Major product additions that have expanded the value of Red Hat's overall software story include middleware (Red Hat JBoss, Red Hat Middleware, Red Hat Integrations), virtualization (Red Hat Virtualization), private cloud (Red Hat OpenStack Platform), container support (Red Hat OpenShift Container Platform), application development (Red Hat CodeReady Studio, Red Hat CodeReady Workspaces), Storage (Red Hat Ceph Storage, Red Hat OpenShift Container Storage); and automation and management (Red Hat Ansible Platform Automation, Red Hat CloudForms, Red Hat Insights, Red Hat Satellite).

In recent years, Red Hat has been investing heavily in expanding its OpenShift product, embracing Kubernetes container orchestration, adding cloud deployment and support options, Kubernetes Operators, service mesh, serverless, and most recently, refining the deployment strategy with OpenShift Container Platform 4 through the inclusion of a lightweight Linux Kernel (Red Hat Enterprise Linux CoreOS) as part of the OpenShift installation.

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 open source market and to have a fully mature support model, including in critical areas like Linux and Kubernetes. While there were overlaps in Red Hat's portfolio and IBM's portfolio, including a pre-existing Kubernetes offering with IBM Cloud prior to the acquisition, 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 Container Platform heavily leverages the existing and future work that Red Hat has invested and will 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 ship software and run modern applications.

Docker and 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 a centralized container registry 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, and others. Governance of the Kubernetes project was handed over to the newly formed Cloud Native Compute Foundation (CNCF) to make the project truly open and available to the IT industry. The industry subsequently embraced Kubernetes as the preferred orchestration platform for containerized applications.

Along with the standardization of container formats and runtimes with the Open Container Initiative (OCI), the industry now had a common container stack which 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 eventually become as ubiquitous as 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). 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 test systems 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.

For IT operators, containers and Kubernetes offer a modern, highly scalable, and automated way of running large web-scale applications. Kubernetes and Kubernetes Operators embed much of the knowledge and experience of web-based companies necessary to run fast-changing apps reliably at 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 stateless and immutable infrastructure. This means that a container state 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.

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 must behave the same for core functions. As a result, customers can largely use any Kubernetes product and expect a certain 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 underneath. This can help provide a consistent developer environment across on-premise and various public clouds. In addition, containers can help abstract across different system architectures as the container interfaces remain the same regardless of the hardware underneath, and 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 the Red Hat OpenShift Container Platform deployment is that it offers a portable environment, such that customers can deploy on OpenShift consistently on-premise, off-premise 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. Code compiled for x86 Linux platforms will not be executable on Linux on IBM Z systems, or on IBM Power systems, as a container platform does not provide architectural abstraction. 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 offer benefits that extend to a Red Hat OpenShift Container Platform deployment including:

- **High levels of scalability.** The many virtualization options for Z provide numerous ways to allocate resources to OpenShift. In addition, the dynamic allocation options and fast provisioning of Z hypervisors can be leveraged to augment Kubernetes' scale up and scale out capabilities.
- **Security.**
  - IBM offers 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 KVM hypervisor can be run within partitions to further subdivide the system.
  - 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. Hardware acceleration of encryption keeps performance overhead of encryption low.
  - IBM Hyper Protect Virtual Servers (the evolution of IBM secure service container technology) provides a type of secure enclave to enable what is commonly called trusted or confidential computing. Container workloads can be run within the enclave and inherit the security and privacy without requiring any code changes. Hyper Protect Virtual Servers

provide fully encrypted and isolated areas within the CPU and memory, guarding even against a malicious admin. This complements OpenShift's default use of SELinux policies for every container at time of creation to provide resource and "nearest neighbor isolation."

- **Performance.** IBM Z supports 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.
- **High reliability levels.** IBM Z is designed to adapt to planned or unplanned events while keeping services and operations running continuously. The IBM GDPS offerings deliver resiliency capabilities for customers who 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.
- **Expands the developer base** that can write to IBM traditional systems. The availability of the Red Hat OpenShift Container Platform on IBM systems means developers need no unique skills specific to the IBM environment; Red Hat OpenShift provides abstraction such that developers only need be comfortable with the Red Hat OpenShift Container Platform. IBM Cloud Pak for Applications then builds on OpenShift to help bridge applications written for IBM Z with modern application development tools, extending the capabilities of all developers looking to modernize existing IBM Z applications.
- **Extends IBM Z participation** in the open hybrid cloud world. Today most platforms will need to support hybrid cloud environments, if they will be a deployment option for modern cloud-native, containerized applications. Having Red Hat OpenShift on IBM systems means that portability, agility and scalability attributes of a Kubernetes platform apply to the IBM systems. Because OpenShift runs wherever RHEL is certified to run, OpenShift helps to extend the reach of IBM Z in hybrid and multicloud environments.

Development for Red Hat OpenShift Container Platform on IBM Systems

Developers deploying on Red Hat OpenShift gain the benefit of a collection of development environments. 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 has also stated its intention to add IBM Wazi for Red Hat CodeReady Workspaces to the suite of IBM Z solutions, with the intent of further strengthening an enterprise-wide standard DevOps toolchain that empowers developers with the capability to develop and test hybrid applications in a containerized, virtual Z environment optimized for Red Hat OpenShift.

IBM intends to base this offering on Red Hat CodeReady Workspaces and is currently planning to offer enterprise development teams a collaborative and scalable platform that can enable developers to more efficiently and effectively deliver hybrid applications.

Developers will have the flexibility to choose an integrated development environment with options that include 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. They can build with IBM Dependency Based Build (DBB) and automatically unit test within their standard integrated development environment to improve quality, reduce cost, and accelerate delivery time.

A driving factor for standard tools across the stack is enabling developers to get started quickly with IBM Z and z/OS development by working with familiar cloud-native tools that improve productivity.

IBM Cloud Paks

In addition, IBM ported and enhanced its Cloud Paks product suite so the suite can be installed and utilized directly on Red Hat OpenShift. IBM has been expanding its Cloud Pak portfolio with products designed to support the following needs:

- **IBM Cloud Pak for Applications.** Includes tools for modernizing existing applications and building new cloud native applications.
- **IBM Cloud Pak for Data.** Includes tools to help customers collect, organize and analyze data.
- **IBM Cloud Pak for Integration.** Offers application integration tools such as API management and lifecycle, data integration, high speed data transfer and integration security.
- **IBM Cloud Pak for Automation.** Offers business process automation tools and capabilities to analyze enterprise processes.
- **IBM Cloud Pak for Multicloud Management.** Tools to support visibility, automation and governance across multiple clusters and cloud platforms.
- **IBM Cloud Pak for Security.** Includes tools to find hidden threats, enabling response to those threats, and empowering customers to integrate tools and workflows across multicloud environments.

Red Hat Ansible Automation for IBM Z

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, validating and reusing Ansible's core automation code building blocks, generally known as playbooks, modules and roles.

The Red Hat Ansible Automation Platform includes Red Hat Ansible Tower to manage complex automation workflows and role-based access controls, as well as to enable API integrations with source control libraries, including certified Red Hat Ansible Content 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 Q1 2020. This certified and supported content provides the core automation artifacts necessary to quickly construct, test and deploy standardized and reusable automation procedures for such common tasks as z/OS job queries, data set management, and command executions.

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 to create OpenShift operators and to integrate OpenShift operations into CI/CD tool chains. Network operations teams rely on Ansible to manage and configure a wide range of network devices. 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 to rapidly scale and adapt digital business innovations that rely on data and resources running on traditional and cloud native platforms.

FUTURE OUTLOOK

It is unusual for the industry to coalesce around a de facto standard as quickly as has happened for containers and Kubernetes run time 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 a part of Kubernetes per se, but integrate with and optimize for Kubernetes, extending the capability 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, as well as on traditional enterprise IT infrastructure such as the IBM Z.

Challenge: Applications on IBM Z 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 Container Platform, 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.

Opportunity: 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 developer opportunity for recruiting developers since the Kubernetes run-time environment is CNCF-certified and behaves as other Kubernetes environments will.

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 — be it on-premises off-premise/cloud, or a multi-architecture scenario — from developers so they 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 and IBM platforms such as Z brings the best of both worlds, where Red Hat delivers a cross-environment deployment and run-time environment that allows customers portability and flexibility, while IBM's platforms offer a clear benefit in terms of scale, flexibility, reliability, security and low latency access to massive corporate datasets 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

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com