HYBRID CLOUD SOLUTIONS FOR THE ENTERPRISE

– Sponsored by IBM® –

As the marketplace continues to evolve and new challenges present themselves, the industry is seeing continued advancements as it relates to the options available for cloud-based solutions in enterprise environments. Looking back to the advent of virtualization, which helped reduce the number of servers an enterprise needed to maintain and in turn reduce overall capital and operating expenses, cloud providers have since taken this model further, creating and running data centers with massive server farms, which is allowing organizations to outsource their IT infrastructures.

The containerization of applications being a natural evolution of virtualization has helped to further reduce resource requirements and increase savings amongst other benefits. Further, Kubernetes, an open source container orchestration platform, has now emerged as the de facto standard for coordinating, automating deployment, scheduling, and scaling containerized applications. The issue however, is that many enterprises cannot easily outsource or transition to these models due to regulatory restrictions for sensitive data, low latency requirements, and/or a preference for the reliability a mainframe offers their mission critical applications.

Today, several cloud vendors are now starting to provide on-premises options to address these restrictions while allowing customers to reap the benefits provided by delivering applications via a container-based cloud model. As such, this analysis, sponsored by IBM, looks at some of the key players from across the industry (i.e., AWS, Google, IBM/Red Hat®, and Microsoft®) from the perspective of on-premises enterprise application modernization.

While AWS, Google, IBM/Red Hat and Microsoft each provide their own Kubernetes Platform as a Service (PaaS) solutions through the public cloud, the focus here is the on-premises implementations of these offerings, which allow data regulatory restrictions to be more effectively addressed across the enterprise.

That said, each vendor has its own unique approach in how it is going to market and the variability of how these vendors' on-premises Kubernetes solutions are delivered will inevitably require some discussion around Infrastructure as a Service (IaaS), and hybrid and multi-cloud implementations.

The following provides a brief introduction to the on-premises solutions provided by each:

- **Amazon Web Services Outposts (AWS Outposts)** provides an IT as a Service solution (i.e., completely provided and managed by AWS) that extends AWS infrastructure and select services (e.g., Elastic Kubernetes Service (EKS), Elastic Compute Cloud (EC2), Elastic Block Storage (EBS), etc.), to a customer site using the same hardware found in AWS data centers.
- **Google Anthos** provides a software focused solution that allows customers to implement a private cloud using the Google Kubernetes Engine (GKE) from their own or service provider's datacenter.
- **IBM/Red Hat OpenShift® Container Platform** (Red Hat OpenShift) offers Platform as a Service and full-stack automated operations to ease the management of private, hybrid, and multi-cloud application container deployments across multiple hardware architectures including x86, IBM Z® and IBM LinuxONE.
- **Microsoft Azure Stack Hub** provides similar services as the Azure public cloud offering for more disconnected environments and includes the Azure Kubernetes Service (AKS)—replacing the former Azure Container Service offering—for deploying and managing containerized applications.

While implemented differently, these offerings are comparable in a few areas such as user management, monitoring and alerts management, maintenance, and more. Since individual customer requirements can vary significantly, the following is meant to focus on some of the core areas of differentiation that may play a role in the decision making process for prospective customers.

1 Supported Options for On-Premises Deployment

Prescriptive Hardware

Many organizations will already have processes in place for the management of their IT environments including hardware acquisition and management. When implementing an on-premises Kubernetes solution, these vendors provide a wide range of options from entirely managed hardware to roll-your-own, enabling customers to capitalize on existing hardware facilities if desired.

On one end of the spectrum, AWS Outposts is designed and provided exclusively by AWS. Unlike Microsoft and other providers, AWS owns the entire supply chain of AWS Outposts, advertising the hardware rack to be exactly the same as that used in the AWS public cloud, therewith providing a uniform experience across its public cloud and AWS Outposts.

Hyper-Converged Infrastructure (HCI) Partners

With its Azure Stack Hub, Microsoft has chosen to go with certified integrated systems hardware through partners (e.g., HPE, Dell, Lenovo, etc.) over DIY infrastructures to ease sourcing, implementations, updates, and more. At the time of writing, Nodes within an Azure Stack cloud must be the same, including the CPUs, the amount of memory, storage, firmware, etc. Even though customers must pay for the hardware, they do not have full access to Azure Stack hosts (e.g., to run agents for backup or anti-virus). Administrators must configure the system through the Azure portal and Azure Resource Manager APIs. In short, Azure Stack is meant to complement existing on-premises and cloud-based virtualization solutions, not replace them.

Custom Hardware

At the other end of the spectrum, both Google Anthos and IBM/Red Hat OpenShift focus primarily on software-only solutions where customers can simply download the respective software, installing it on their existing infrastructures. That said, for customers requiring additional infrastructure, Google Cloud and Red Hat have both partnered with hardware vendors to provide prepackaged hyper-converged reference implementations for a simpler deployment experience. Furthermore, Red Hat OpenShift supports IBM Z and LinuxONE platforms facilitating highly secure, flexible, scalable, and reliable implementations for mission-critical workloads.

Physical vs. Virtualized Infrastructures

While both Google Anthos and Red Hat OpenShift are essentially hardware agnostic allowing customers to use their existing infrastructures, the caveat is that Google Anthos requires a VMware infrastructure (VMware vSphere 6.5 and up) for its virtual appliances. Red Hat OpenShift's CoreOS can be deployed on bare-metal and/or within virtualized infrastructures (including IBM z/VM®) allowing Red Hat OpenShift to work with a variety of hardware—including customized hardware already available on customer sites, even directly on AWS Outposts and Microsoft Azure Stack Hub virtualization services—as long as the hardware is aligned with the recommended minimum requirements.

On-Premises vs. Hybrid

While the focus of this discussion is on-premises implementations, future expansion or offloading of processes and applications not restricted by regulation may eventually be desirable. In all cases, these vendor solutions allow expansion to their respective public cloud offerings (i.e., hybrid) in some way shape or form, but support for third parties to facilitate multi-cloud implementations are again variable.

First, AWS Outposts and Google Anthos deployed on-premises require an active connection to their respective cloud offerings for access to the control plane, monitoring, and/or management. In short, these solutions are not designed for offline on-premises implementations, essentially architected as hybrid-cloud solutions to address regulatory requirements.

While AWS focuses on AWS only, the tight integration of Google Anthos and GKE means that a customer's hybrid cloud strategy must be Google based, leaving little room for a multi-cloud strategy depending on the functionality that GKE Connect might eventually provide. With the Microsoft introduction of Azure Arc, Microsoft will eventually (i.e., in preview at the time of writing) support deployment of applications to "any" standard Kubernetes cluster.

Multi-Cloud Support

Following on its support for custom infrastructure components, Red Hat OpenShift is not specific to the IBM Cloud® offering, but is deployable on-premises and via third party cloud infrastructures (including AWS, Google Cloud Platform, IBM Cloud, and Microsoft Azure), facilitating true multi-cloud implementation scenarios.

Horizontal vs. Vertical Scalability

With respect to on-premises implementations, the ability to scale is ideally constricted only by the customer's available hardware infrastructure. In the case of AWS and Microsoft, the requirement for prescribed hardware can artificially introduce limitations even if the customer's datacenter has additional resources to spare. For example, Microsoft Azure Stack integrated systems can support 4 to 16 nodes in an individual cloud.

Each of the vendors support expansion into their respective public cloud offerings to scale as required. For example, one aspect of the incoming Microsoft Azure Arc offering is that in combination with Azure cloud platform, this can facilitate "unlimited scalability" by seamlessly spinning up additional Kubernetes clusters in Azure Kubernetes Services if the currently configured environments run out of capacity. In all cases however, customers would need to ensure that secure data and processes can be forced to remain behind the firewall to meet regulatory restrictions.

In addition to adding and removing compute machines manually, Red Hat OpenShift also provides an elastic, dynamically provisioned method on top of private (or public) cloud infrastructures automatically scaling machines and cluster sizes. Administrators can set the minimum and maximum scaling limits for resources such as cores, nodes, memory, GPU, etc., increasing (or decreasing) the physical size of a cluster to meet deployment needs.

When running on the IBM Z or IBM LinuxONE platforms, Red Hat OpenShift provides additional vertical scalability through a single physical footprint to meet the needs of even the most demanding enterprise environments.

Application Portability

One of the benefits of containerized applications is their portability, which allows customers to easily move applications between development, test, and production environments without issues surrounding different configurations. Due to the plethora of service options available from vendors, organizations might need to be conscientious to maintain this portability. For example, Microsoft, while also supporting Docker, provides customers with other architectural options for application development (e.g., Azure App Service, Azure Functions, etc.) that can tie applications to the Azure platform, reducing portability outside the Azure ecosystem.

AWS is still making significant investment into its Amazon Elastic Container Service (ECS) and is putting a lot of effort in promoting it as the leading container orchestration tool on AWS. Even when AWS does provide a managed Kubernetes service in the form

of Amazon Elastic Kubernetes Service (EKS), it is evident that AWS considers a managed Kubernetes experience a secondary choice to Amazon ECS.

Red Hat OpenShift ensures portability of services. In the case of IBM, this portability extends from distributed to mainframe platforms allowing customers to, for example, develop Linux applications in distributed environments and move to IBM Z (either Public or Private cloud implementations) to capitalize on its reliability and security benefits without application changes. Using containers and standard Kubernetes for container orchestration allows customers to deploy open and portable applications to multiple cloud environments without vendor lock-in.

2 Availability of On-Premises Services (Compared to Cloud Offerings)

IaaS vs. PaaS

The strategies offered by each vendor can be divided into two groups: Google and IBM/Red Hat focus on software-only PaaS on-premises implementations allowing customers to use their own infrastructures; while the hardware requirements associated with AWS and Microsoft solutions facilitate the delivery of both IaaS and PaaS services on-premises.

As AWS Outposts is setup similar to any other rack in AWS public cloud, both containerized workflows and virtual machine workflows are supported. Furthermore, given the initial positioning of AWS Outposts as an edge device, database and big data processing workloads are supported using both containerized and non-containerized solutions—containerized workflows can be scheduled using Amazon ECS and Amazon EKS. Similarly, through Microsoft Azure Stack customers can provision VMs to provide infrastructure services in addition to deployment of PaaS solutions.

While Google and IBM both provide IaaS offerings through their respective public cloud solutions (i.e., Google Compute Engine and IBM Cloud Compute), their focus on PaaS essentially leaves on-premises VM provisioning and management to established best of breed solutions; although Google Anthos deployed on-premises does require VMware specifically.

If required, IBM and Red Hat do provide dedicated IaaS functionality on-premises through solutions such as Red Hat OpenStack, Red Hat Virtualization, and IBM Cloud Infrastructure Center, which can provide on-demand self service provisioning of virtual machine instances on IBM Z and IBM LinuxONE platforms.

On-Premises vs. Hybrid

While each vendor provides a comprehensive set of services via their respective public cloud offering, they offer only a subset of these services for on-premises purposes. As mentioned, Google and IBM/Red Hat focus on PaaS, eliminating IaaS functionality from an on-premises implementation without the use of a dedicated solution. That said, Red Hat OpenShift does provide the same PaaS experience whether delivered on-premises or on a cloud infrastructure provided by AWS, Google, IBM or Microsoft.

Although Microsoft Azure Stack also provides the facilities to deploy and manage virtual machines and infrastructures on-premises, the VM sizes supported by Azure Stack (e.g., A, Av2, D, Dv2, and F series) are a subset of those supported by the public Azure offering. The Azure public offering itself imposes resource limits to avoid overconsumption. The same limitations are imposed on Azure Stack to match Azure. From the PaaS perspective, Microsoft continues to expand its support, adding highly requested capabilities while it continues to synchronize Azure Stack and Azure functionality. Similarly, only a subset of services are supported on AWS Outposts, with claims that more services will be provided in time.

3 Availability of Management Facilities to Minimize Operational Effort

Platform Maintenance & Upgrades

Each of these vendors provide facilities to ease the management of their respective solutions. For example, not only are Red Hat OpenShift Operators a means of packaging, deploying, and managing Kubernetes applications, but they also play a fundamental role to the OpenShift Container Platform code base in that they are also used to automate upgrades for the OpenShift operating systems (i.e., RHCOS) and control plane applications. The Operator Lifecycle Manager facilitates the installation, update (automated or manual), and management of all Operators and their associated services running across a cluster, ensuring that the OS, Kubernetes, and OpenShift Container Platform components are all updated.

Since the on-premises Google Kubernetes Engine is deployed as a virtual appliance on compatible on-premises hardware using VMware vSphere and ESXi, Google treats this as a logical extension to its managed GKE cluster to offload cluster upgrades and patching from the customer. Similarly, AWS Outposts being essentially an edge device simplifies management by connecting to a parent AWS region for management. As an entirely managed solution, AWS is also responsible for the monitoring and maintenance of the underlying appliances. If there is a hardware failure, AWS will send out replacement hardware, as well as a technician to perform the swap out or instruct the client IT teams to replace the faulty hardware.

Microsoft Azure Stack receives updates on a monthly basis, and customers must update it at least every three months to continue being supported. These updates not only contain security and bug fixes, but also new functionality in an effort to keep Azure Stack up to date with the public Azure cloud offering. Azure Stack Hub hardware partners are responsible for the servicing of hardware-related firmware and driver update packages. Additionally, Azure Stack Hub hardware partners own and maintain guidance for all software and hardware on the hardware lifecycle host.

Unified Hybrid Management

Today, enterprises are implementing multi-cloud architectures using multiple providers and a mix of public and private platforms as the preferred strategy. This allows the separation of different workloads into specific environments depending on requirements. An ideal management solution would provide a single view for the management of all these resources.

One of the reasons AWS and Google require an active connection to the respective core cloud solutions, is that these deployments are treated as an extension to the core cloud solution to be managed through the respective UIs. For example, once AWS Outposts are connected to the parent AWS Region, they appear in the same AWS Management Console. Customers can launch Amazon EC2 instances and use EBS storage as well as other AWS services on Outposts using the same APIs in the AWS Regions.

As previously noted, AWS focuses on AWS only, or “all-in on AWS”. This means that the AWS management strategy works for hybrid implementations. Similarly, through GKE Connect Google customers can use a single control plane to manage container workloads running across its public (and in the future AWS and Azure) and private cloud offerings.

Unified Multi-Cloud Management

While Microsoft Azure and Azure Stack use the same Azure Resource Manager technology, customers have noted that since each Azure Stack scale unit is a region by itself, users must access multiple portals. Customers are asking for a unified Azure Stack portal that will allow administration of multiple scale units and regions. The announcement of Microsoft Azure Arc (still in Preview at the time of writing) might help address this with

the management of container-based applications across multi-cloud Kubernetes clusters, but it is not expected to address infrastructure provisioning which still requires separate access. More specifically, the issue is that resources will need to be manually pre-provisioned through the native means of the third-party provider and added to Azure Arc (using a locally installed agent) before they can be managed.

IBM addresses an enterprise's need to support and manage multi-cloud deployments today through its IBM Cloud Pak® for Multicloud Management. Not only does IBM Multicloud Manager support administration of IBM Cloud and Red Hat OpenShift applications through a single user experience, but it also provides visibility into workloads across multiple Kubernetes clusters and clouds; even those deployed on third party cloud platforms such as Microsoft Azure and Amazon Web Services. Furthermore, IBM Cloud Automation Manager supports automated provisioning and service initiation for multiple cloud vendor IaaS solutions through a single UI with optional workflow orchestration.

4 Support for Enterprise Application Modernization to Maximize the Benefits of Containerization

Encapsulation

To help modernize existing applications and allow enterprises to capitalize on the benefits enabled through a cloud model, Google and Microsoft have essentially come to market with an encapsulation solution.

With the technology Google acquired from Alooma in 2019, Google Migrate for Anthos provides an automated migration experience to Google Cloud where it will eventually help extract the application and data portions in most Virtual Machines (initially only VMWare and GCE will be supported), converting them into containers running specifically on the Google Kubernetes Engine (i.e., Anthos GKE deployed on-prem is not supported). While this helps containerize compatible workloads, they are essentially still monolithic applications encapsulated in a container.

With Microsoft, facilities are provided to migrate .NET applications by scanning an app with an endpoint scan, downloading the migration assistant and then migrating the application to Azure. In short, Azure Stack can give customers a way to move .NET workloads to a cloud model, replacing database VMs with PaaS offerings, IIS VMs with App Services, and linking them together with Azure Functions. Thanks to its IaaS support, Azure Stack can also provision virtual machines as part of an application deployment. This allows customers to create mixed applications without forcing them to create containerized implementations.

Migration

Red Hat Container-Native Virtualization (currently in Tech Preview) will help customers move and manage traditional VM-based applications using OpenShift. Running inside a pod, OpenShift will manage the virtualized workload alongside the containerized workloads as part of a single application allowing customers to move application components to containers over time.

While AWS relies on partners for migration services, Google announced in February 2020 that it acquired Cornerstone, a Dutch company that specializes in helping enterprises migrate their legacy workloads from mainframes to public clouds through extensive hands-on migration assistance. While this acquisition will form the basis of Google Cloud's mainframe-to-GCP (Google Cloud Platform) solutions, it does not address the on-premises perspective discussed throughout this analysis.

Of these vendors, IBM is the only one that provides a significant focus on helping organizations modernize their existing enterprise applications to capitalize on a cloud model while maintaining on-premises requirements. For example, Red Hat OpenShift helps

address situations where regulatory and critical sensitive data needs to stay behind the firewall (also addressed by IBM Cloud Hyper Protect Services, a set of enterprise-grade Cloud services for securing business critical data). Supported on IBM Z and LinuxONE enterprise platforms, Red Hat OpenShift makes cloud-native workload modernization on these enterprise platforms a much more streamlined process (e.g., Linux on Z applications can be moved to Red Hat OpenShift with little to no application changes).

Accelerators

Ultimately it is cost prohibitive to re-write applications from the ground up, and perpetual evidence shows that the mainframe continues to play a key role in business transactions as workloads continue to increase. While numbers vary between studies, it is commonly recognized that 70% to 80% or more of business transactions touch the mainframe at some point or another. It is commonplace to run a variety of diverse middleware products and functions on the IBM Z platform, often with multiple instances of some.

IBM has been effective in providing tools to facilitate mainframe application modernization to help migrate existing enterprise applications and capitalize on the benefits provided through a cloud management model. For example, IBM Transformation Advisor helps identify Java applications that can be deployed to a cloud model for faster ROI. Furthermore, IBM Cloud Paks allow commonly used IBM components (e.g., IBM WebSphere®, IBM MQ, etc.) to be moved to a container model with minimal application changes, allowing these components to capitalize on the benefits of a cloud implementation.

Enterprise Development Tools

While each of these vendors provide tools for creating containerized applications with deployment through Kubernetes and Helm/Operators, IBM also addresses the needs of larger enterprise customers that rely on the mainframe to meet the most demanding enterprise workloads. IBM z/OS® open cloud native development tools expedite the creation of containerized enterprise applications.

Today, new enterprise applications can be built to natively support cloud models while still benefiting from the reliability and security provided by the IBM Z and LinuxONE platforms. In short, developers can build new mainframe-based applications using a cloud containerized model to achieve the best of both worlds (reliability, security, portability, scalability, etc.).

FINAL THOUGHT

The value of containerized applications is more than simply a new way of packaging and delivering applications. The real value lies in the longer-term ability to move applications and application components into services layers for simplified management, scalability, and the ability to derive new applications without a requirement to rebuild the entire infrastructure layer.

Although the examined solutions are all based on the Kubernetes platform, IBM provides more focus on helping large enterprise customers that arguably have more value stored in legacy mainframe enterprise applications compared to their distributed VM housed counterparts.

Selecting the right on-premises Kubernetes solution will be very individual given diverse and unique business requirements. As each of these vendors pursue different philosophies and strategies regarding an on-premises, hybrid, and/or multi-cloud Kubernetes platform, their different implementation perspectives play a significant role when comparing the individual strengths and weaknesses of the respective solutions. It is not simply a matter of looking at an overall rating average (see below), but rather, it requires prospective customers to consider these strategies fully when assessing which solution best fits their individual development and delivery culture.

FEATURES COMPARISON

– AWS/Google/IBM/Microsoft –

Core Evaluation Factors				
Feature	**AWS**	**Google**	**IBM**	**Microsoft**
Platform Support				
Hardware Support (incl. custom, HCI, and architecture support)				
OS Support (considers Installation and VM Guest OS support)				
Client OS Support (incl. web-browser and mobile UIs)				
Database Support (incl. Container / On-Premises Options, and via Public Cloud)				
Directory Server Support (e.g., LDAP, Microsoft Active Directory, OIDC, etc.)				
Scalability and Availability (focus is on-prem implementations)				
Services Support				
Vendor Public Cloud Services (generally available and Cloud PaaS)				
Private Cloud (considers on-prem, hybrid, and multi-cloud implementations)				
Infrastructure as a Service (looks at on-prem and cloud availability)				
Platform as a Service (primarily on-premises availability)				
VM Support (considers on-premises options and dedicated solutions)				
Storage Services (considers on-prem, public cloud, and partner support)				
Messaging / Middleware Support (on-prem and generally available options)				
Machine Learning / AI (incl. platforms, engines, and tools)				
Marketplace / Service Catalog (with public access and on-prem availability)				
Extension of Enterprise Assets and Industry Standards Support				
Integration Support (with existing enterprise assets)				
SDLC / DevOps Integration (incl. native and 3rd party, CI/CD Tools, etc.)				
Enterprise Application Modernization (encapsulation and migration options)				

Core Evaluation Factors				
Feature	**AWS**	**Google**	**IBM**	**Microsoft**
Cloud Native Development / Developer Tools (such as IDE, Built/Test Options, SDKs, etc.)				
Application Deployment (both deployment and management)				
Security / Encryption (incl. services, key mgmt., data at rest and in motion)				
Administration and Management				
Infrastructure Support (physical and virtual using custom or HCI)				
User Management (incl. Role Based Access Control and Identity Management)				
Hybrid / Multi Cloud Management (disparate vs. unified centralized UIs)				
Policy Based Management / Automation ("as Code" and image/container automation)				
Service Management (such as Service Mesh support and other extended cloud services)				
Monitoring, Alerts and Management (centralized and 3rd party options)				
Auditing (of on-prem and hybrid / multi-cloud deployments)				
Maintenance (concerning upgrades and required downtime)				
Backup and Disaster Recovery (considers innate facilities, third party support, etc.)				
Vulnerability Scanning (examining container / image scanning and DAST)				

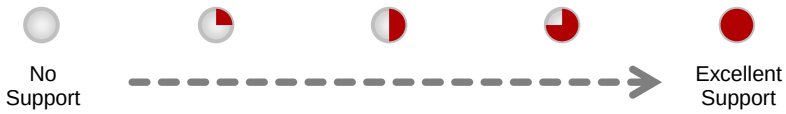

Note: This assessment was commissioned by IBM Corp. Zibis Group does not endorse any vendor solution. This report is provided as a guide to help evaluate products based on several critical areas of consideration when choosing an on-premises, hybrid, and/or multi-cloud solution for the enterprise.

ABOUT ZIBIS GROUP

Zibis Group Inc. ("Zibis Group") is a leading source for market research and advisory services on the Information and Communications Technology (ICT) industry and related sectors. We assist vendors, government institutions, academia, non-profits and other key industry stakeholders by providing them with the tools they need to make informed decisions.

Zibis Group is a niche play in the marketplace for customized ICT sector research. We understand that all companies have distinct interests and unique business challenges. We thrive on rising to the occasion and in seeing our clients succeed. Technology is changing and we're here to help.

For more information regarding Zibis Group, please visit **www.zibisgroup.com**.

ZibisGroup
555 Legget Drive, Suite 304 Tower A
Ottawa, Ontario
Canada K2K 2X3
Tel: +1 613-518-8006
www.zibisgroup.com