White Paper

The Essentials to Delivering Persistent Storage for Container-based Applications

Red Hat Delivers Enterprise-grade Persistent Storage that Maximizes the Potential of Container-based Environments

By Scott Sinclair, ESG Senior Analyst
and Monya Keane, Senior Research Analyst

March 2019

This ESG White Paper was commissioned by Red Hat and is distributed under license from ESG.

Contents

Introduction

For organizations trying to increase their competitiveness in a digital age, containers can be a lifesaver.

Container technology is a way of packaging software that abstracts applications from the hardware environment. Containers give organizations assurance their apps will remain operational even when, for example, they are moved from dev/test into production or are redeployed from a physical host onsite to a VM in the cloud.

While often compared to virtual machines, containers are quite different. While virtual machines virtualize the underlying hardware, containers virtualize the operating system. As a result, containers are smaller in size and require less overhead. This lighter weight design improves portability and necessitates a more dynamic underlying infrastructure as containers often scale faster and to greater numbers than their virtual machine counterparts.

Companies using container platforms do so to adapt to evolving market demands and changing customer preferences. Those adaptations put pressure on IT teams—they need to be more agile in accelerating development and deployment of new, modern workloads. In a data-driven economy, organizations succeed largely by accelerating application development. Application developers and IT operations personnel are at the forefront of that important effort.

To put it another way, IT is no longer just a cost center providing basic "utility" services. Developers in particular are specialists in the nuances of translating data into a competitive asset. They now generate revenue and enable cost savings.

That is likely why new ESG research on application development spending is showing that 31% of the organizations surveyed intend to devote significant funding to implementing a formal DevOps practice in the next 12-18 months.[1] But transitioning to modern application development means IT must transform and move away from legacy infrastructures and support processes.

Containers Are Forming the Foundation of a Larger Trend

Implementing formal DevOps practices appears to be part of a larger trend—digital transformation. And containers will play an integral role.

Last year, 86% of IT decision makers surveyed by ESG said they agreed with the statement, "If we do not embrace digital transformation, we will be a less competitive and/or effective organization."[2] This year, ESG investigated that sentiment in more detail, finding that reasons for pursuing digital transformation vary, but all of them have clear potential to bring positive business impacts.

Some of the surveyed organizations believe digital transformation will help them be more operationally efficient (55%). Some want to provide a better, more differentiated customer experience (43%). Some hope to develop new data-centric or innovative products and services (a combined 73%). And more than a quarter of organizations surveyed said they are proceeding with digital transformation initiatives because they intend to develop an entirely new business model (28%).[3]

Digital transformation priorities invariably include app-dev-related improvements, and containers support those efforts. Developers now require a great degree of control over tools, processes, and platforms as they write and test their apps. Their counterparts in IT operations, however, are still charged with managing those same tools, processes, platforms, and environments consistently across the entire enterprise. Containers align and support both groups because they allow for application abstraction and portability.

Containers support *everyone* through application portability.

[1] Source: ESG Research, *2019 Technology Spending Intentions Survey*.
[2] Source: ESG Master Survey Results, *2018 IT Spending Intentions Survey*, December 2017.
[3] Source: ESG Research, *2019 Technology Spending Intentions Survey*.

A major function of a formal DevOps practice is to refine how developers and IT ops admins collaborate with one another. Here, containers play a vital role. They assist IT in streamlining everything: development, testing, and ultimately, the dependable, predictable, and frequent release of new and differentiated applications.

Some companies are putting a lot of budget toward that effort—21% of organizations surveyed by ESG identified increasing infrastructure capacity supporting application development as one of their most significant app-dev investment areas, and 17% identified increasing the use of containers as a significant app-dev investment area.[4]

The Essential Role of Container Technology in Enhancing DevOps Success

As Figure 1 shows, a combined 56% of IT organizations surveyed by ESG now use containers in production to support applications, and 13% of them have already deployed an extensive number of containerized production applications.[5]

Figure 1. Containers in Production

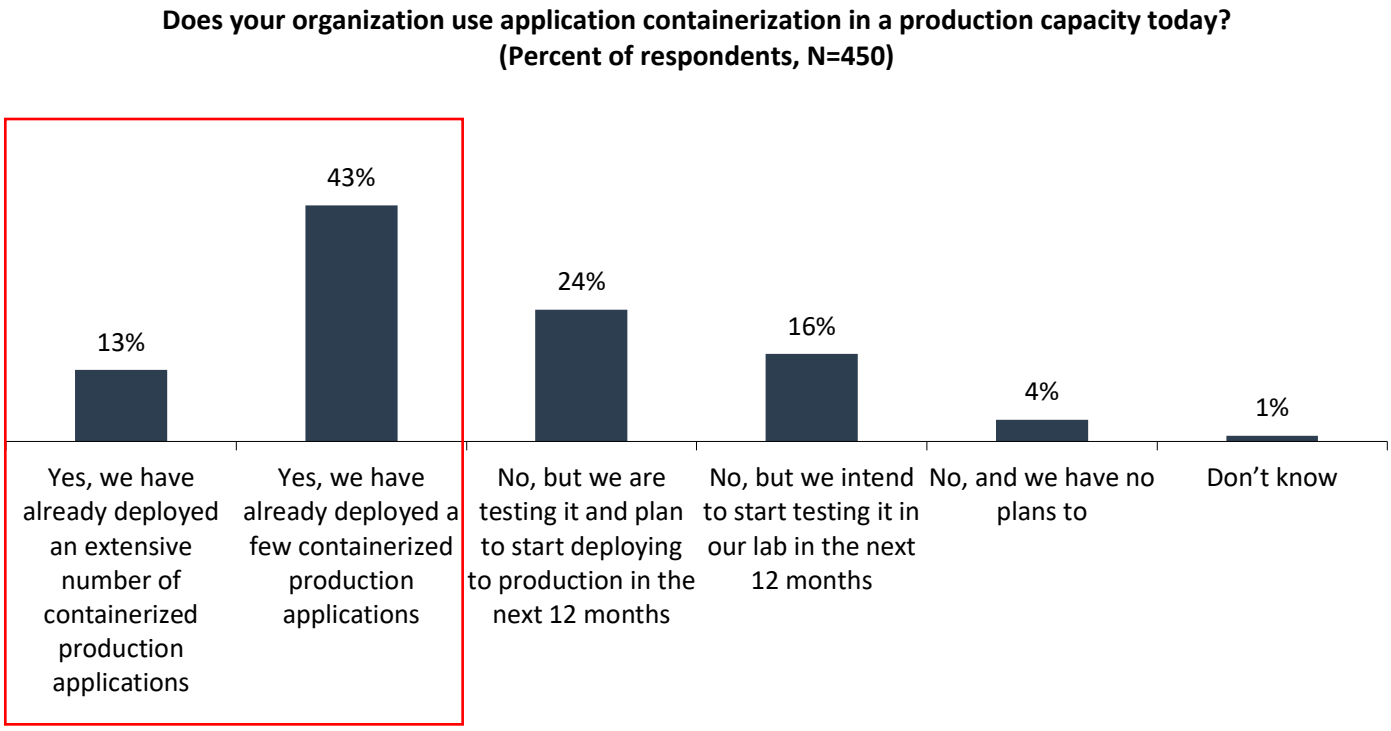

Source: Enterprise Strategy Group

That growth in implementation is unsurprising. IT organizations know that traditional IT systems and services impede the goal of being fast and frequent in developing/deploying production applications. The process to provision traditional infrastructure is too slow, usually involving multiple teams and excessive levels of bureaucracy. Adjusting to subsequent changes is slow and cumbersome as well. But container environments overcome this problem:

- Containers offer autonomy for applications and for the developers who create them. They do so by packaging applications with the libraries and bins those applications depend on to run.
- Containers combine application isolation with an image-based deployment scheme. Containers don't need an operating system kernel, which helps them stay lighter *and* more reliable than a VM environment.

[4] Source: ESG Brief, *2018 Application Development Spending Priorities*, February 2018.
[5] Source: ESG Brief, *The Growth in the Use of Application Containers,* May 2018.

Such benefits have a positive impact on application developers and IT operations administrators alike (see Figure 2).

Figure 2. Containers Benefit Both Application Developers and IT Operations Administrators

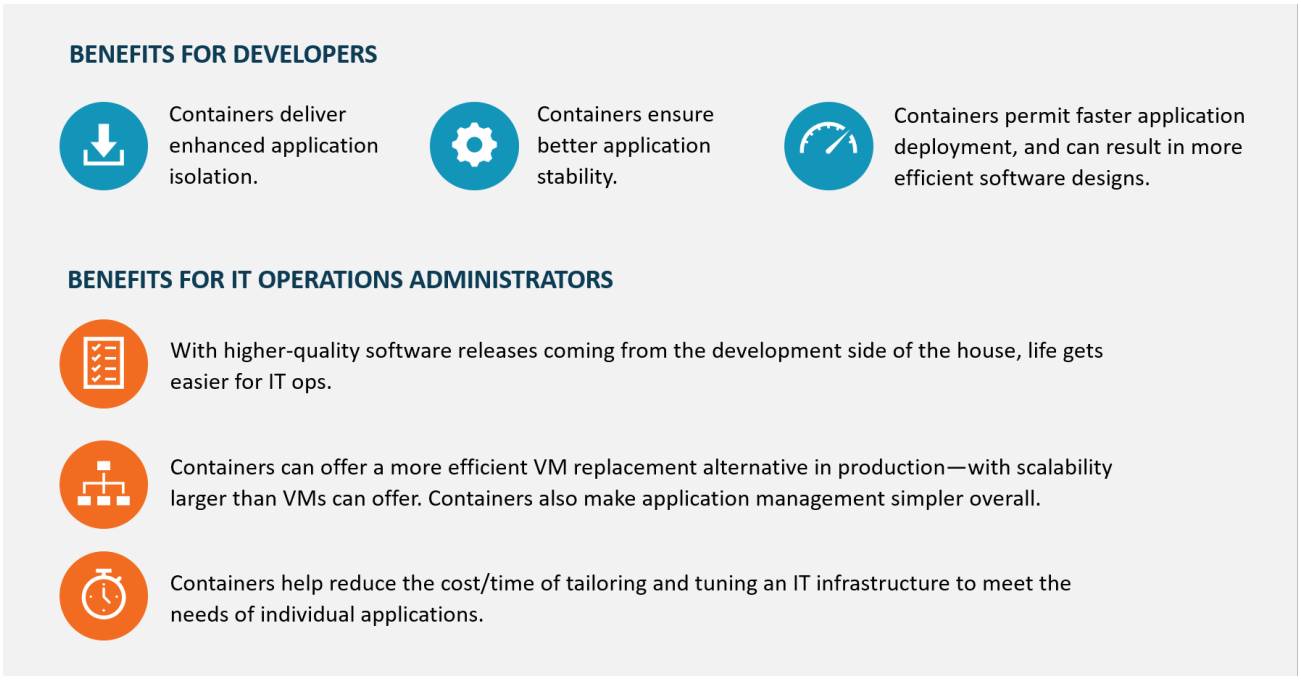

Source: Enterprise Strategy Group

By providing these benefits, containers can go a long way to resolving the management-oversight-related conflicts that developers and operations personnel otherwise will experience. Those conflicts deter and delay development work, which in turn hampers business success.

It's another reason organizations are becoming increasingly bullish on containers. Consider that the IT managers ESG surveyed report they expect container use to soon increase enough to cause measurable reductions in both bare metal hardware deployments *and* VM deployments (see Figure 3).[6]

[6] Source: ESG Master Survey Results, *Trends in Hybrid Cloud Security*, March 2018.

Figure 3. Containers in Production

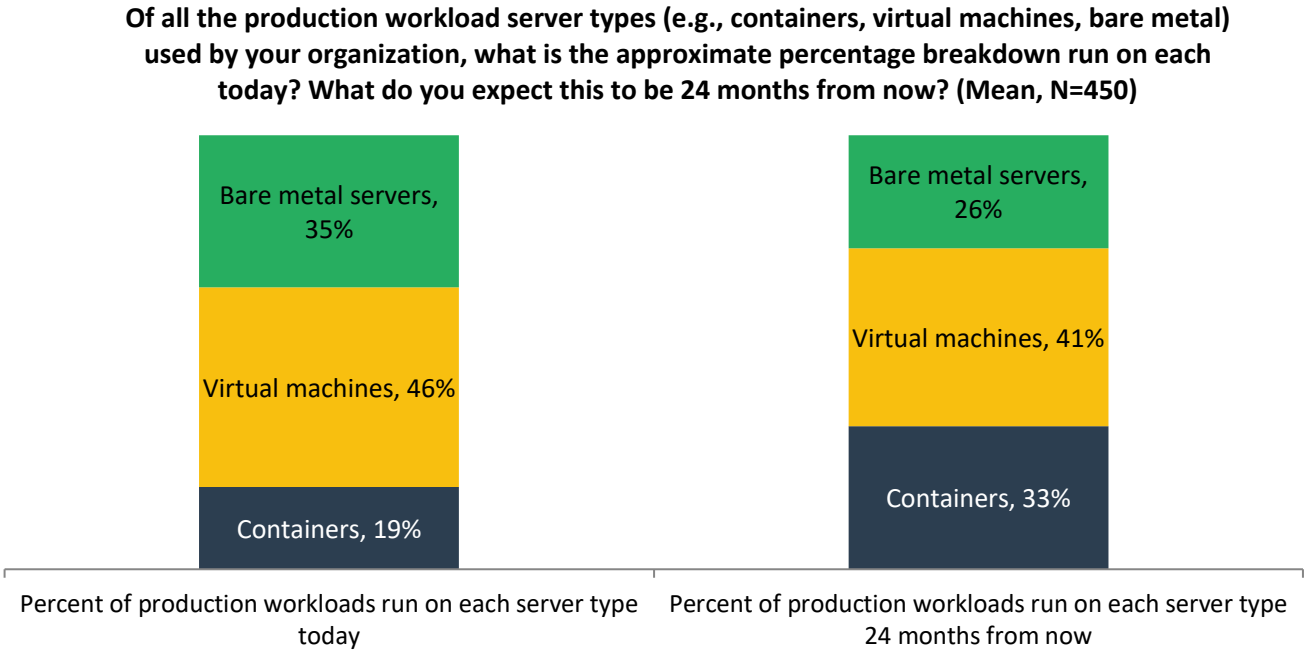

Source: Enterprise Strategy Group

Containers Also Support a Multi-cloud Infrastructure

Public cloud services—including multi-cloud environments—also come into play in relation to containerized application development.

ESG is seeing widespread use of multi-cloud environments. Recent findings show a combined 76% of surveyed public cloud infrastructure service users are now leveraging more than one provider (IaaS and/or PaaS), with only 23% using just one cloud provider at this point.[7]

Containers can ease those organizations' transition to using public cloud IT by making it much easier to transition an application between different cloud providers or move it between the onsite data center and a public cloud service. That capability also helps organizations avoid cloud vendor lock-in.

In general, ESG has found that 51% of organizations leveraging containerized applications still deploy them solely onsite. But the remainder are deploying them either partially *or exclusively* in public cloud services (see Figure 4).[8]

[7] Source, ESG Research, *2019 Technology Spending Intentions Survey*.
[8] Source: ESG Brief, *The Growth in the Use of Application Containers*, May 2018.

Figure 4. Container-based Application Deployment

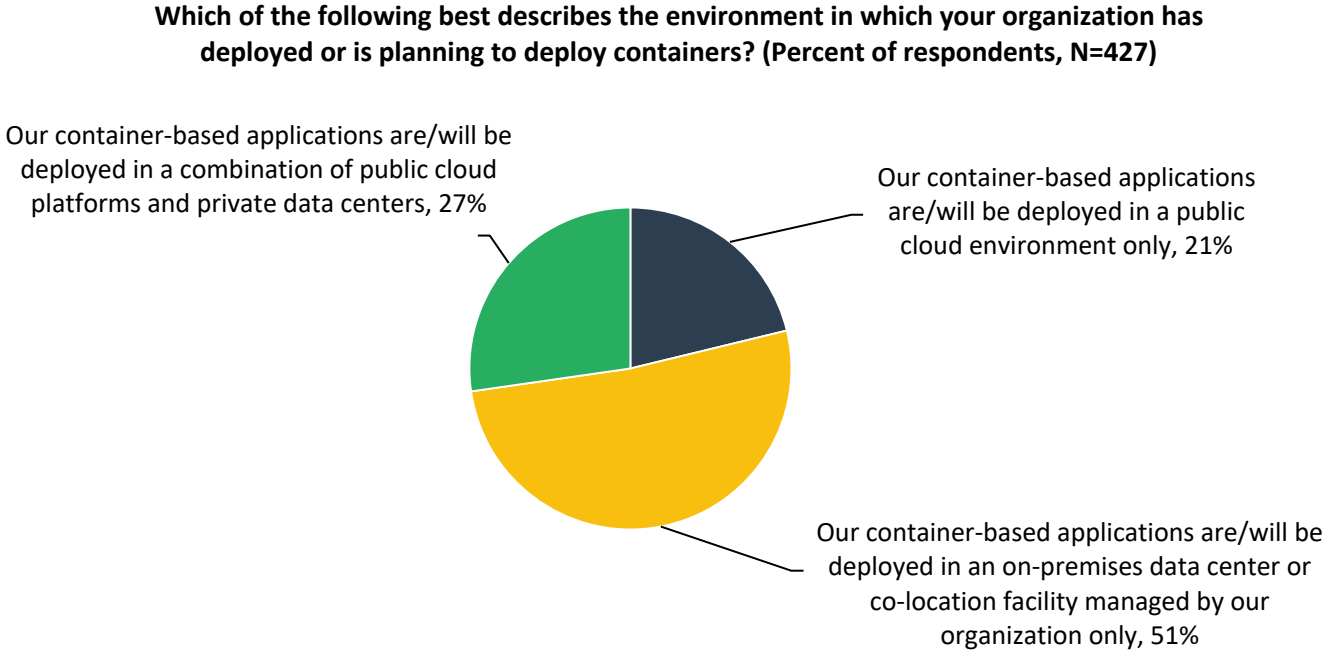

Source: Enterprise Strategy Group

The Essentials for Providing Persistent Storage for Containers

As the benefits of containers fuel adoption, demands arise for access to persistent data storage. Containers are ephemeral by nature, able to start up and shut down automatically. Without being stored on persistent storage, the data will die with the container. When addressing the data storage demands of containers, traditional storage arrays or even those fitted with a software "wrapper" between the array and the container-based environment are not entirely sufficient. They fill a checkbox rather than fully supporting the containerized application.

Yes, containers need data persistence, but standardization, abstraction, scalability, performance, resiliency, and automation are essential as well.

Standardization & Abstraction

Standardization delivers a foundation for portability. Without being able to standardize, porting containerized applications across a wide variety of environments, including public clouds, private clouds, and traditional data center infrastructures would be too cumbersome to be viable. The challenge to achieving standardization, though, is that everyone's needs and use cases are different.

Fortunately, the Kubernetes community is working toward ways to better enable hybrid cloud management for containers. Kubernetes is an open source orchestration system that automates the deployment, scaling, and management of containerized applications. Operating at the container, rather than the hardware, level, Kubernetes has become a technology of choice for increasing the speed and efficiency of containerized application development and deployment.

While most implementations of Kubernetes adhere to the Cloud Native Computing Foundation (CNCF) software conformance program, additional production elements are often required, and can vary from deployment to deployment, or cloud to cloud. Developers need a standardized abstraction layer that exists between the various iterations of

Kubernetes (or other container orchestration frameworks such as Mesos or Cloud Foundry) so that they can easily and seamlessly "talk" to each other.

This is what the Container Storage Interface (CSI) open source project will bring to the table. The CSI is a standardized mechanism for storage across the different container orchestration systems, including Kubernetes, Mesos, Docker, and others. With the CSI, storage should work uniformly across different container orchestrator systems, regardless of the storage provider.

Scalability, Performance, & Resiliency

Microservices and cloud-native applications require scalability beyond traditional applications. Refactoring many of today's applications could result in hundreds, if not thousands, of microservices running simultaneously.

This need for extreme scale introduces a couple of vital questions that must be answered:

- How to efficiently and quickly attach persistent storage to those microservices?
- How to ensure data integrity and availability as microservices scale?

It's an intractable problem for both developers and storage administrators. The time and effort required to manually provision storage is not practical. Development teams need a scalable and expedited approach, especially when moving that new application/microservice into production.

Persistent storage, when served as a native service of the Kubernetes platform, can be enabled via dynamic provisioning. This can allow developers to easily deploy storage without having to submit a request to a storage administrator. This may help accelerate development since developers do not have to wait hours, days, or weeks to have their requests fulfilled.

Once provisioned, the infrastructure must be performant and resilient. Containers applications run by the hundreds, even the thousands, and access to persistent storage must be predictable and immediate. Without the right supporting infrastructure, the data cannot keep pace with the container demand.

Automation with Kubernetes Operators

Kubernetes Operators, announced in 2018, enable the use of automation, thus providing a simple, expedited way to package, deploy, and manage a Kubernetes application. As a result, Kubernetes applications can function like cloud services, complete with self-provisioning and self-healing features.

This delivers automated management and maintenance at production scale, including automated installation and upgrading, which becomes incredibly valuable when applied to storage provisioning, the migration of data, and other potentially time-consuming tasks. It allows organizations to access the necessary storage without jumping through a lot of hoops.

Fortunately, there are several options for automating storage-related tasks, including the open source Rook project and API-based storage. Rook is an emerging persistent storage orchestrator that is designed to run as a native Kubernetes service—key for making storage automation work. And API-based storage is fundamental for automation. As demand for new microservice applications increases, reliance upon APIs is expected to increase as well.

Deploying container-native storage (software-defined storage running in the container) allows for setting up automated storage management. Kubernetes itself doesn't have a storage component, but storage can be automated through the Kubernetes orchestration engine just like many other operational tasks. Enterprise Kubernetes platforms should provide a

self-service way to request whatever storage (i.e., a persistent volume) is needed for the application. Kubernetes can also be used to mount and add storage to containers and have the engine automatically assign storage resources appropriately.

As Operators, designed via the Kubernetes Operator Framework, increase in development, adoption, and maturity, container-based storage benefits of automated installations and upgrades will be extended even further. Thus, Kubernetes can take care of the storage housekeeping, allowing development teams to focus on what they do best - development.

It is hard to overstate how valuable Kubernetes Operators are in addressing challenges experienced by both developers and IT operations staff: Kubernetes Operators give greater control to developers, simultaneously reducing the strain on operations administrators.

Red Hat Leading the Way for Container Storage with Kubernetes

Red Hat, a leader in open source development since the early 1990s, with a wealth of experience supporting development communities, offers software-defined storage built for container environments with its Red Hat OpenShift Container Storage (OCS) product.

OCS integrates into the larger Kubernetes-based Red Hat OpenShift Container Platform. OCS scales across container, bare metal, VM, and cloud deployments and gives developers' application data a permanent place to live, even when containers are spinning up or down.

Red Hat OpenShift Container Storage

Red Hat OpenShift Container Storage delivers the essential architectural benefits of persistence, standardization, abstraction, performance, and scalability vital to containerized environments, along with the other essential storage requirements of resiliency and automation.

- Red Hat co-developed it with the Red Hat OpenShift Container Platform. That means OCS is optimized to serve as an ideal storage back-end for a Red Hat OpenShift environment—it is not only optimized for the container application data, but can also serve as the underlying storage for the OpenShift Container Platform deployment (storage for the registry, logging, and metrics infrastructure elements).
- Its software-defined storage architecture provides infrastructure efficiency and flexibility and enables expedited app deployments.
- It allows for data, and the application that consumes the data, to coexist on the same infrastructure. This translates into a denser, more integrated deployment.
- It is manageable through a single control plane, the OpenShift UI, which covers the entire container-based development environment—both the compute and storage resources.
- It can run anywhere OpenShift is deployed (bare metal, virtualized, on-premises, or in the cloud) and provides a consistent experience across hybrid and multi-cloud deployments.

Business Benefits of Partnering with Red Hat for Containers

Red Hat OpenShift Container Storage delivers enterprise-level persistent, abstracted storage to maximize the benefits of a container environment. It is not only optimized for container environments, but also managed with Red Hat OpenShift under a single interface. And thus, one vendor, Red Hat, handles all necessary support for the entire container-based environment in OpenShift. That level of consolidation translates into business benefits including reduced costs, simplicity, and easier acquisition/support. Those benefits are going to be especially helpful to any company that is still in the early stages of exploring container technology.

Another business benefit of OCS is the support of multiple protocols in one architecture—block, file, and object—which obviously translates into lower footprint CapEx for a business.

And when it comes to leveraging the cloud, as mentioned, Red Hat's multi-cloud support will alleviate the business risks of cloud vendor lock-in. (It runs on-premises, in virtualized environments, and on three public cloud services.) OCS offers a unified experience across platforms, with features such as dynamic provisioning adding to the seamlessness.

Red Hat also leverages the open source project Prometheus to allow deep, advanced container storage visibility and monitoring. This capability gives a business a quick view of *what* is happening with its OpenShift storage without having to worry about *how* it's happening.

Lastly, transformational business value arises when development teams control their own infrastructure. OCS lets developers self-manage the storage they consume. They no longer need a storage administrator's involvement. Red tape is reduced, and services are expedited. This is a case of a leading-edge IT solution delivering a very classic business benefit—*efficiency*.

The Core Features of OpenShift Container Storage

- Highly scalable, multi-petabyte architecture
- Supports up to 1,000 persistent volumes per OpenShift cluster.
- Supports file, block, and object storage.
- Addresses both container application storage needs and OpenShift registry, logging, and metrics storage.
- Supported in a single contract with Red Hat OpenShift.

The Bigger Truth

Organizations often find themselves needing to resolve technology concerns that have business-level ramifications as well. For example, maybe they are running on old infrastructure, and they have reached the point at which they are just limping along.

Recently, many such organizations have decided they want to put together a foundation that allows for container use to begin or continue. They will most likely be running stateful applications in those containers—SQL databases, messaging platforms, web applications, and similar applications characterized by a need for persistent data and persistent registry, log files, and metrics files. Even if they are running stateless applications exclusively, they will still need persistence.

IT decision makers who don't want to go down the wrong path should know that Red Hat is a company that believes it can provide them with the best solution. Red Hat has long been an expert in open source and now has containers and OpenShift Container Platform as a key strategic direction—making sure that containers are part of its DNA as a company. Compare that focus to what some other storage vendors offer: They tend to see support for containerized applications as little more than a box to be checked off.

Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides actionable insight and intelligence to the global IT community.