

## SOLUTION BRIEF

# A Data Platform for Developers

Joint solution from Cisco, Pure Storage and Red Hat speeds and simplifies the transition to Kubernetes-based software development

Organizations are modernizing applications by adopting containers and Kubernetes to take advantage of full-stack automated operations, a consistent experience across all environments, and self-service provisioning for developers that lets teams work together to move ideas from development to production.

Cisco and Pure Storage® have joined with Red Hat to deliver a powerful bare metal solution for deploying Kubernetes infrastructure. This platform solution combines Red Hat OpenShift with Cisco compute and networking, and Portworx® by Pure Storage as a storage management layer.

This solution is designed to address challenges in Kubernetes deployments, including:

- **Slow application development**  
Platform and DevOps teams struggle to onboard apps and often take the CSI “connector” approach to access the infrastructure, leading to high developer wait time and a poor quality of experience.
- **Operational complexities**  
Specific hardware lock-in hampers the ability to move applications across on-prem and cloud environments. Site Reliability Engineers and DevOps teams also require specialized skills to operate the infrastructure.
- **Lack of data availability**  
Mission-critical Kubernetes apps require data availability and resiliency in the event of failure of the underlying infrastructure.

To address these challenges and others, Portworx functions as a data services platform on top of Red Hat OpenShift and provides multiple value-added services that simplifies storage management for IT, makes developers more productive and helps de-risk the movement to a container-based development environment.



## Developer Self Service

- Free developers from IT dependencies
- Increase productivity of dev resources
- Automate Day 2 workflows



## Simplified Data Services

- Integration with popular data services
- Easy, automated deployment
- Supports Elasticsearch, Redis, Cassandra and many more



## Cost Efficient

- FlashArray data reduction reduces storage consumption
- Pay as you grow pricing available
- Reduces management overhead

## A Full Developer Platform

Developers are best served by a complete platform experience that eliminates the pain of having to use different tools to build the solution they want. Application teams must fully own the operations for their apps while staying compliant with their SLAs, cost estimates, and regulatory requirements.

But what is a true developer platform experience? And what does a cloud-native data platform look like? Let's examine the capabilities that a cloud-native data platform delivers to development teams.

### A Self-Service Developer Experience

The core value of this solution is that developers can be made more productive and effective via self-service. When you're serving a large number of end users and applications, you need a self-service experience where the users can consume the capabilities of the platform without intervention from the platform administrators. And when a large community uses the product, it's critical to have enterprise-class role-based access control.

For example, with the Portworx platform, a developer can fully manage their application deployment storage characteristics, set policies, take backups, and move their apps and data to and from anywhere without ever having to file a ticket or wait for someone else to do it for them. This is true Kubernetes-based DevOps at scale.

### Cost Savings

The sheer scale of operating and managing applications in production can cause cost overruns, overwhelm an organization, and burden the platform teams—and that affects the application team's productivity as well.

A cloud-native data platform helps rein in the expense by delivering infrastructure cost savings through a pay-as-you-grow model as well as operational cost savings by automating 90% of the day 2 workflows, thus reducing the operating expenses burden of maintaining a large platform team to run and manage apps in production.

### Manage Scale via Policy

Users often push the limits on cluster sizes. Development environments can include thousands of nodes in a cluster and run hundreds of clusters with hundreds of nodes. Kubernetes runs at staggering scales, and managing data on Kubernetes at that scale requires more automated management of the data infrastructure, capacity allocation, and performance management with little to no user intervention.

The developer solution delivers this through the Portworx Autopilot functionality, which can monitor and manage large-scale clusters, dynamically manage capacity, and balance the applications in a Kubernetes cluster to deliver the best utilization.

### Simplified Data Services

Development of any application requires creating multiple iterations, trying out different methods and schemas, and testing a lot of sample data sets that mimic production. When running these applications in production, these data services also need to be administered properly, and often the application teams heavily rely on data services teams to keep these data services up and running.

The Portworx platform includes Portworx Data Services, which delivers single-click deployment with fully automated day 2 management of 12 different data services as of September 2023, including SQL Server, MySQL, PostgreSQL, MongoDB, Redis, Elasticsearch, Cassandra, Couchbase, Kafka, Consul, RabbitMQ and ZooKeeper. It can handle thousands of database instances on any cloud with little to zero user intervention, essentially putting database management on autopilot with a single pane of glass across all clouds.



In a hybrid cloud environment, this significantly reduces the operational complexity and frees up platform engineers from having to be responsible for different data services. It also reduces DBA fatigue from having to administer so many data services, provides application owners total control and transparency over the data sovereignty, and gives line of business owners complete control over the costs. Plus, it enables development teams to spin up as many instances of their test data sets and test out different versions of the software instantaneously without expensive setup time.

### Enhanced Application Mobility

In many cases, what holds rapid application prototyping and development back is the infrastructure. Developers try to build modern applications on outdated infrastructure. They're held back from delivering great customer experiences because it's exponentially harder to deliver application experience closer to where their customers are because they're limited by what their infrastructure can support. Platform teams often struggle with keeping their infrastructure up to date as a result.

A true cloud-native data platform gives users the ability to create a global data fabric, which enables applications to be deployed and to run everywhere. This enables easy mobility of applications from one infrastructure to another, from data center to cloud and vice versa, across different clouds, and between different versions of hardware and software infrastructure. When mobility increases, productivity increases and value creation accelerates.

## Solution Components

The following are the major components to the solution. This is not a complete inventory.

### Cisco Unified Computing System (Cisco UCS)

Cisco Unified Computing System (Cisco UCS) is a next-generation datacenter platform that integrates computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership and increase business agility.

### Easy Data Protection SLAs

A cloud-native data platform delivers application-aware and application-consistent business continuity and disaster recovery so each application owner can easily define how and where their application instances are deployed and what kind of SLAs—commonly called RPO (Recovery Point Objective) and RTO (Recovery Time Objective) — the application requires and get it at the granularity of each app. This provides tremendous control back to the application owner and frees up the platform team from having to understand and manage different application DR needs.

### Observability

Platform teams can struggle to understand how their applications are performing and where to look when there's an issue.

A cloud-native data platform helps users measure performance from the application level and provides the ability to drill down to the individual infrastructure components, including CPU, memory, and I/O devices. It integrates natively with the customers' dashboards and other observability tools by continuously exporting all the relevant metrics. In short, it gives the application owner and the platform engineer total visibility into the efficiency of the application infrastructure and helps them run the application in the most efficient manner.

### Cisco Nexus Switching Fabric

The Cisco Nexus 9000 Series Switches offer both modular and fixed 1/10/25/40/100 Gigabit Ethernet switch configurations with scalability up to 60 Tbps of nonblocking performance with less than five-microsecond latency, wire speed VXLAN gateway, bridging, and routing support.



## RedHat OpenShift

The Red Hat OpenShift Container Platform is a consistent hybrid cloud foundation for building and scaling containerized applications. Once installed, Red Hat OpenShift uses Kubernetes operators for push-button, automatic platform updates for the container host, Kubernetes cluster, and application services running on the cluster.

## Portworx

Portworx is a data management solution that serves applications and deployments in Kubernetes clusters. Portworx is deployed natively within Kubernetes and extends automation capabilities down into the infrastructure to eliminate all the complexities of managing data. Portworx provides simple and easy-to-consume storage classes that stateful applications can use in a Kubernetes cluster.

## Pure Storage FlashArray//XL

FlashArray//XL™ offers powerful enterprise storage that combines cloud-like scale with lightning-fast performance, supporting today's largest applications and databases. Pure's FlashArray//XL delivers high-density enterprise storage with lightning-fast performance and cloud-like scale to support the most-demanding applications and to maximize workload consolidation.

Portworx combined with Pure Storage FlashArray™ can be used as a cloud storage provider. This allows administrators to store your data on-premises with FlashArray while benefiting from Portworx cloud drive features, automatically provisioning block volumes, expanding a cluster by adding new drives or expanding existing ones. Pure Storage FlashArray with Portworx on Kubernetes can attach FlashArray as a Direct Access volume. Used in this way, Portworx directly provisions FlashArray volumes, maps them to a user PVC, and mounts them to pods.

---

## Additional Resources

- Read more about the high-performance [FlashArray//XL](#)
  - Get the details about the [Portworx](#) cloud-native data platform
  - Read about joint Cisco and Pure Storage [FlashStack® solutions](#)
- 

[purestorage.com](https://purestorage.com)

800.379.PURE

