

The Data Foundation for Modern Infrastructure

Erik Zandboer
Cloud Native Architect
Portworx | EMEA-North

Priorities Driving CIO Agendas



Enterprises Don't Want Separate Platforms



They want ONE platform

Market trends are shaping Kubernetes adoption.
The platform is ready.
The data layer is not.

Plan to modernize and/or migrate VMs to Kubernetes

74%

Run databases on Kubernetes, the #1 workload four years running

66%

Now run AI/ML workloads on Kubernetes, up sharply

44%

Cite storage as their #1 AI cost concern, ahead of GPUs

50%




Now view Data on Kubernetes as foundational or a competitive advantage

86%

Red Hat as likely K8s platform for VMs

81%

Customers think they're *only* making a compute decision



In practice, they *also* need to make a data architecture decision



Modern applications & AI
are **orchestrated** flows of data.

The durable design decision is not the compute layer alone. **It is the data architecture beneath it.**

The Everpure Enterprise Data Cloud Portworx Extends into Kubernetes



Agents



Analytics



Apps

Universal Data Intelligence

Delivered
as a Service

Intelligent Control Plane

Unified Data Plane

Evergreen Architecture



Edge



On-prem



Public Cloud



Fusion Automation

Portworx and FlashArray Integration

Unified Storage Infrastructure

Integrated Kubernetes-native
Data Management



Delivering the full enterprise stack on Red Hat OpenShift



Modern Virtualization

Virtual machine alternative with Red Hat OpenShift Virtualization

Everpure and Portworx bring critical storage capabilities for enterprise virtualization solutions with Red Hat.



Application Modernization

Self-service access to storage for apps running on Red Hat OpenShift

Portworx and OpenShift deliver a solution to automate, protect, and unify Kubernetes applications and data.



Artificial Intelligence (AI)

AI Data Management and Performance aligned with Red Hat OpenShift and Red Hat AI

Everpure and Portworx pave the way to enterprise AI on OpenShift

CSI plug-ins start the journey, but don't finish it...



CSI drivers haven't kept up with the innovations in the Kubernetes space and they're definitely **pinning the customers to the hardware** implementations.

— Blue Cross Blue Shield of Alabama

Create Data Silos
and Fragmentation



Require Manual
Ops with Inconsistent
Capabilities



Lead to Complex
Security & Compliance



Cause Upgrade
& Migration Disruption



Extending advanced capabilities beyond CSI spec

Enterprise-grade kubernetes data management capabilities and support

Portworx Enterprise

App Aware, Granular
Backup & Restore

Namespace
scoped Security

App Aware
Disaster Recovery

Live Workload &
Storage Mobility



Hybrid and
Multi Cloud



Heterogeneous
Storage



Edge Devices
Local Disk

Everpure CSI (PX-CSI)

One CSI for all Everpure storage

FA/FB Direct Access

Unified block + file

Array-native snapshots and clones


Everpure Storage Platform

FlashArray


FlashBlade

Everpure Cloud


Most enterprises are not single vendor environments



You need a single data control plane across your entire application estate



You need a single data control plane across your entire application estate





Modern Applications



Customer Challenge

Fiserv powers real-time payments for thousands of financial institutions. Legacy infrastructure bottlenecked and couldn't scale under strict regulation, pushing Fiserv to modernize.



Portworx & Red Hat Solution

Modern apps on OpenShift with Portworx by Everpure as the Kubernetes-native data layer: automation, high availability, disaster recovery, and self-service across on-prem and cloud.



Customer Outcomes

Developers self-serve their own data, shipping multiple deployments a week. Fiserv runs a resilient payments platform scaling from 1,000 to 2,000 nodes across hybrid and multi-cloud.



The impacts of the business are clearly reliable software deployments, where we're deploying more features...this allows developers to get towards multiple deployments perhaps per month, per week."

John Durko
Director of IT –
Container Services
Fiserv

Red Hat OpenShift

Portworx

Everpure

ARO

AKS



Modern Virtualization



Customer Challenge

Tier 0 rail, where failure is not an option. CSX moved VMs to OpenShift to avoid lock-in, but needed resilient, high-performance data services across hybrid environments.



Portworx & Red Hat Solution

VMs and containers on OpenShift, across datacenter and cloud. Portworx adds the resilient Tier 0 data layer beneath.



Customer Outcomes

A two-person team runs the data behind 400+ apps on one platform.



We did a proof of concept to prove we could quickly build a VM from scratch and migrate a live VM from host to host... We chose Portworx because it was the fastest.”

Eric Grabill
Lead Senior Product
Manager of IT, CSX

Red Hat OpenShift

Portworx

Everpure

ARO

AKS



ROSEN

AI Transformation



Customer Challenge

ROSEN inspects the world's pipelines:
10+ PB of data, growing 1 PB a year.
Manual review of 1,000+ variables
per inspection couldn't scale.
They needed an AI-ready platform.



Portworx & Red Hat Solution

An AI and ML platform on OpenShift
with Kubeflow. Portworx delivers
the automated data layer beneath,
feeding models petabyte-scale data
across clouds.



Customer Outcomes

99% runs as automated code,
cutting analysis time 70%. ML narrows
1,000 variables to 100, letting customers
safely run pipelines up to 5 years longer.



With our AI platform,
powered by Pure
Storage, we've
achieved a 70%
reduction in the time
it takes to deliver
analyses to
customers.
This helps our
customers make
quicker data-driven
decisions, while
allowing us to
complete more
pipeline inspections
each year."

Thorsten de Vos
IT Service Architect
ROSEN Group

Red Hat OpenShift

Portworx

Everpure


ARO

AKS

Hybrid Cloud Architecture with Red Hat OpenShift and Portworx Enterprise

Portworx and Red Hat OpenShift:

Built on the Same Design Principles



Platform and Developer Value

Stop Building Infrastructure. Start Shipping Applications Faster.

Platform Team

01 Platform teams stop managing storage exceptions and start delivering capabilities

02 Storage, protection, security, and DR delivered as a service: one policy model, one runbook, every cluster.

03 The same data platform across cloud, core, and edge, no re-architecting for new infrastructure

04 New app teams onboard with an existing policy, not a new storage architecture.

05 Platform teams shift from bottleneck to service provider.

Development Team

01 Developers self-serve storage, databases, snapshots, and protection via Kubernetes APIs, no tickets.

02 Golden paths and templates deploy any app or database in a few clicks, any storage, any cloud.


03 Data and apps stay portable across on-prem, cloud, and edge. Build once, run anywhere.

04 Developers backup and restore their own data in one click, no ticket required.

05 Development teams ship apps in days, not months


Platform and Developer Value

Stop Building Infrastructure. Start Shipping Applications Faster.



Seamless Workload & Storage Mobility

Across Hybrid and Multi-Cloud



Business Continuity

Hybrid Cloud Region Replication



Azure-EMEA-West-1



AWS-EMEA-East-2



ARO - Cluster

Red Hat OpenShift
(ARO)

Portworx



Worker Nodes + Cloud Drives

K8s
Resource
Replication

Async DR


PX
Volume
Replication



ROSA - DR Cluster

Red Hat OpenShift
(ROSA)

Portworx




Worker Nodes + Cloud Drives



Business Continuity

Hybrid Cloud Region Replication



What Portworx and OpenShift Deliver

One hybrid cloud platform that lowers cost, simplifies operations, and speeds delivery, on every cluster and every cloud.

Reduce Overall Cost



Pool capacity and optimize performance to pay for what you use—not what you reserve.

Streamline Platform Operations



One policy model and runbook across every cluster. Stop reinventing storage for every workload and cloud.

Increase Developer Velocity



Give developers storage, protection, and security without tickets or expertise—so teams ship faster.

Portworx Dynamic OpenShift Plug-in with ACM

Portworx Dynamic OpenShift Plug-in with ACM

VM Storage Dashboard

Integrated storage Observability for KubeVirt VM disks in OpenShift Console, providing capacity usage and performance stats.

Portworx Cluster Dashboard

A single OpenShift-native Portworx dashboard giving at-a-glance health, capacity, and performance info for a Portworx cluster.

Portworx DR with Red Hat ACM

Fleet-wide Portworx cluster visibility and Async DR orchestration natively managed through ACM.

20250813T153538-pxc01-iden.x

Red Hat OpenShift

Pods - Red Hat OpenShift

Overview - Cluster - Red Hat

Everpure - Calendar - 4 da

Notes - Red Hat

Not Secure console-openshift-console.apps.pxe1.vlan142.local/k8s/ns/bbq-bookkeeper/core-v1=Pod?page=1&perPage=50

TIME Personal Pure PX Spreadsheets Meeting Notes Documents Slides blog WJP HOL CNSEA CNPDX Scratch Google Gemini

Red Hat OpenShift

Core platform

Home

Overview

Projects

Search

API Explorer

Events

Favorites

Ecosystem

Workloads

Topology

Pods

Deployments

DeploymentConfigs

StatefulSets

Secrets

ConfigMaps

CronJobs

Jobs

DaemonSets

ReplicaSets

ReplicationControllers

HorizontalPodAutoscalers

PodDisruptionBudgets

Networking

Services

Routes

You are logged in as a temporary administrative user. Update the cluster OAuth configuration to allow others to log in.

Project: **bbq-bookkeeper**

Pods

Name

Filter by name

1-4 of 4

1 of 1

Name	Status	Ready	Restarts	Owner	Memory	CPU	Created
bbqbookkeeper-web-7cb7c8c74c-j87gt	Running	1/1	2	RS bbqbookkeeper-web-7cb7c8c74c	14.6 MiB	0.000 cores	Aug 12, 2026, 6:07 PM
bbqbookkeeper-web-7cb7c8c74c-twix4	Running	1/1	2	RS bbqbookkeeper-web-7cb7c8c74c	7.3 MiB	0.000 cores	Aug 12, 2026, 6:07 PM
bbqbookkeeper-web-7cb7c8c74c-w9rkp	Running	1/1	1	RS bbqbookkeeper-web-7cb7c8c74c	7.3 MiB	0.000 cores	Aug 12, 2026, 6:07 PM
mssql-0	Running	1/1	0	SS mssql	1192.5 MiB	0.008 cores	Aug 12, 2026, 6:07 PM

1-4 of 4

1 of 1

61.0°C

3.2%

57.3%

16:45 (UTC+00:00) 22:15 (UTC+05:30) 9:45 (PDT)

8/13/26 8/13/26 8/13/26

Portworx Async DR with ACM

The screenshot shows the OpenShift web console interface. On the left is a dark sidebar with navigation links: Home, Search, Infrastructure (expanded), Clusters (selected), Automation, Host inventory, Applications, Governance, Credentials, Portworx, Clusters, Disaster Recovery Pairs, Protection Groups, and Disaster Recovery Actions. The main area has a top blue banner stating "You are logged in as a temporary administrative user. Update the cluster OAuth configuration to allow others to log in." Below this is the "Clusters" section header with a help icon. A tab bar contains "Cluster list" (active), "Cluster sets", "Cluster pools", "Discovered clusters", and "Placements". Under "Cluster list", there's a search bar, filter dropdowns for "Filter" and "Labels", and buttons for "Create cluster", "Import cluster", and "Actions". The main content is a table of clusters:

	Name	Namespace	Status	Distribution version	Infrastructure	Control plane type	Cluster set	Labels
☐	acm	acm	Ready	OpenShift 4.21.6 Update available	Other	Hub	default	local-cluster=true openshiftVersion=major-minor-4 openshiftVersion-major-minor-4 velero.io/exclude-from-backups 15 more
☐	pxe1	pxe1	Ready	OpenShift 4.21.6 Update available	Bare metal	Standalone	default	openshiftVersion=major-minor-4 openshiftVersion-major-minor-4 tsslab=true 14 more
☐	pse2	pse2	Ready	OpenShift 4.21.6 Update available	Bare metal	Standalone	default	openshiftVersion=major-minor-4 openshiftVersion-major-minor-4 tsslab=true 14 more

1 - 3 of 3



Portworx Async DR with ACM

The screenshot displays the Red Hat OpenShift console interface for managing Portworx Disaster Recovery (DR) pairs. The left sidebar shows the navigation menu with 'Disaster Recovery Pairs' selected under the 'Portworx' section. The main content area shows the configuration for a specific DR pair named 'dr'.

Disaster Recovery Pairs > dr DRP Ready

Overview | **YAML**

Disaster Recovery Pair Summary

Disaster Recovery Pair Name	Disaster Recovery Pair status
dr	Ready

Pair Type	Primary Cluster (Source)
Bidirectional	pxa1 Scheduler: Ready Storage: Ready

Disaster Recovery Type	Disaster Recovery Cluster (Destination)
Async	pxa2 Scheduler: Ready Storage: Ready

Associated protection groups
0

Created At
Aug 12, 2026, 8:22 AM

Backup Location Configuration

Type	Backup Location
Amazon S3	Use IAM role for authentication: Disabled

Bucket name	Access Key ID
px-dr	*****

Region	Secret Access Key
us-west-2	*****

Endpoint
http://10.13.18.39:9000

Advance Settings

S3 Storage Class

Portworx Async DR with ACM

The screenshot displays the Red Hat OpenShift Portworx console interface. On the left is a navigation sidebar with a 'Fleet management' dropdown menu. The main content area is titled 'Create Protection Group' and includes a breadcrumb 'Protection Groups > Create Protection Group'. Below the title is a description: 'Enroll your namespaces to improve resilience by implementing disaster recovery protection.' The interface is divided into two panels. The left panel contains a numbered list of steps: 1. Protection Details (active), 2. Resource Configuration, 3. Advanced Settings, 4. Replication Settings, and 5. Review. The right panel, titled 'Protection Details', contains the following fields and options: a text input for 'Protection group name' with the value 'bookkeeper'; a dropdown for 'Disaster Recovery Pair' with the value 'dr'; a section for 'Namespace' with instructions to choose how to select namespaces, radio buttons for 'Select from namespace list' (selected) and 'Use label selector', and a search bar with the text '82 results found for source cluster : 1 selected' and a search input containing 'book'. Below the search bar is a table with two rows: one with a checked checkbox and the name 'Name', and another with a checked checkbox and the name 'bbq-bookkeeper'.

Red Hat OpenShift

Fleet management

Home >

Search

Infrastructure

Clusters

Automation

Host inventory

Applications

Governance

Credentials

Portworx

Clusters

Disaster Recovery Pairs

Protection Groups

Disaster Recovery Actions

Protection Groups > Create Protection Group

You are logged in as a temp

Create Protection Group

Enroll your namespaces to improve resilience by implementing disaster recovery protection.

- Protection Details
- Resource Configuration
- Advanced Settings
- Replication Settings
- Review

Protection Details

Define the protection group, select the DR pair and backup location, and choose the workloads to protect.

Protection group name *

bookkeeper

Disaster Recovery Pair *

dr

Namespace

Choose how you want to select the namespaces to protect in this group.

☒ Select from namespace list ☐ Use label selector

82 results found for source cluster : 1 selected

Name 🔍 book ✕

☒	Name
☒	bbq-bookkeeper

Portworx Async DR with ACM

The screenshot displays the Red Hat OpenShift Portworx Async DR interface. The left sidebar shows the navigation menu with 'Protection Groups' selected. The main content area shows the 'bookkeeper' Protection Group in a 'Protected' state. The 'Operational Status' section shows a progress bar with stages: Sync, Protected, Failover, and Failback. The 'Cluster direction' section shows a flow from 'Primary' (pxe1) to 'DR' (pxe2). The 'Status' section shows 'Protected state' as Healthy, 'Replication state' as Active, and 'Sync state' as Partial Success. The 'Data & last sync' section shows the last successful sync time as Aug 13, 2026, 9:50 AM. The 'Data transferred' section shows 1 Volume and 8 Resources. The 'Current Stage' is 'Replication completed successfully'. The 'Sites' section shows 'Primary site' as Active and 'DR site' as Standby. The 'Sync execution details' section shows 'Started' at Aug 13, 2026, 9:47 AM and 'Finished' at Aug 13, 2026, 9:50 AM, with a duration of 2m 52s.

Red Hat OpenShift

You are logged in as a temporary administrative user. Update the [cluster OAuth configuration](#) to allow others to log in.

Protection Groups → bookkeeper

PG bookkeeper Healthy

Overview YAML

Operational Status

DR lifecycle

Sync Protected Failover Failback

Cluster direction

Primary pxe1 Active → DR pxe2 Standby

Status

Protected state
Healthy

Replication state
Active

Sync state
Partial Success

Data & last sync

Last successful sync time (Failover will use this sync data)
Aug 13, 2026, 9:50 AM

Data transferred
1 Volume
8 Resources

Current Stage
Replication completed successfully

Sites

Primary site
State: Active
Placement: Primary
Namespaces: 1

DR site
State: Standby

Sync execution details

Started
Aug 13, 2026, 9:47 AM

Finished
Aug 13, 2026, 9:50 AM

Duration
2m 52s

Portworx Async DR with ACM

The screenshot displays the Red Hat OpenShift console interface for managing Portworx disaster recovery. The left sidebar shows the navigation menu with 'Protection Groups' selected. The main content area shows the 'bookkeeper' Protection Group in the 'Preparing Failback' state.

Operational Status

DR lifecycle

Sync → Protected → Failover → Failback

Cluster direction

Primary: pxe1 (Standby) → DR: pxe2 (Active)

Status

Protected state: Preparing Failback

Replication state: Suspended

Sync state: -

Data & last sync

Last successful sync time (Failback will use this sync data): -

Data transferred: -

Current Stage: Reverse sync is in progress.

Sites

Primary site: State: Standby

DR site: State: Active, Placement: DR, Namespaces: -

Sync execution details

Started: -

Finished: -

Duration: -

Better Together: Everpure + Red Hat Delivering Customer Results

350+

Joint customers running, VMs,
Containerized Apps, & AI Workloads

38%


Lower total cost of
ownership across the
virtualization estate*

90%

Faster provisioning:
From 5–10 days
to under 2 hours*

68%

Faster application delivery Container
& VM workloads on one platform*





Thank you.





EVERPURE BLOG

Everpure AI Factory

AI at the Edge

COMPUTE • PLATFORM • DATA

LEARN
MORE
AT OUR
BOOT
H



Built for resilient edge inference



Architecture shortfalls without a common data foundation

The Shortfall	The Challenge	Triggers	Warning Signs	The Mitigation Plan
Treating all CSI drivers as equal	Feature sets (block vs file, snapshot semantics, topology, cloning, plugin maturity) vary	Most enterprises are, or become mixed storage environments over time	Missing snapshot support, failed live-migration, unexplained volume attach errors	Create a short vendor/driver checklist and TEST
Ignoring performance tests	Real workloads (VMs, DBs) have different I/O patterns — synthetic defaults can wildly under/over-estimate results	Workloads change or new teams and apps get added to platforms	High tail latency under load, noisy-neighbor IO spikes, CPU stalls in VMs	Run benchmarks, measure throughput, IOPS, and 99th/99.9th percentile latency. Test at expected scale
No backup/failover plan	Replication isn't a backup; edge cases (corruption, accidental deletion) need point-in-time recovery. VMs need consistent backups/app-level quiesce	Consistent DR services needed across hybrid/multi-cloud infra	No tested restore procedure, unclear RTO/RPO, manual recovery steps, inability to restore VM to a prior point	Define RTO/RPO, implement tested snapshot + off-cluster backup workflow, automate restores, practice runbooks regularly
Running outdated drivers	Older drivers miss fixes for attach/mount races, security patches, compatibility with newer K8s versions, leading to instability and data risk	Common Day 2 ops like update and upgrade cycles	Driver crashes, CSI sidecar restarts, incompatibility warnings during K8s upgrades, open CVEs	Adopt a driver lifecycle policy: pin versions in manifests, run a canary cluster for upgrades, subscribe to vendor/security advisories, automate upgrades where feasible after smoke tests



Modern Applications



Customer Challenge

Fiserv powers real-time payments for thousands of financial institutions. Legacy infrastructure bottlenecked and couldn't scale under strict regulation, pushing Fiserv to modernize.



Portworx & Red Hat Solution

Modern apps on OpenShift with Portworx by Everpure as the Kubernetes-native data layer: automation, high availability, disaster recovery, and self-service across on-prem and cloud.



Customer Outcomes

Developers self-serve their own data, shipping multiple deployments a week. Fiserv runs a resilient payments platform scaling from 1,000 to 2,000 nodes across hybrid and multi-cloud.



The impacts of the business are clearly reliable software deployments, where we're deploying more features...this allows developers to get towards multiple deployments perhaps per month, per week."

John Durko
Director of IT -
Container Services
Fiserv

Red Hat OpenShift

Everpure

ARO

AKS





Modern Virtualization



Customer Challenge

Tier 0 rail, where failure is not an option. CSX moved VMs to OpenShift to avoid lock-in, but needed resilient, high-performance data services across hybrid environments.



Portworx & Red Hat Solution

VMs and containers on OpenShift, across data center and cloud. Portworx adds the data layer beneath, with the resilience and speed Tier 0 demands.



Customer Outcomes

A two-person team runs the data behind 400+ apps on one platform. CSX modernizes on its own timeline, VMs to containers, without adding tools or teams.



We did a proof of concept to prove we could quickly build a VM from scratch and migrate a live VM from host to host... We chose [Everpure] and Portworx because it was the fastest.”

Eric Grabill
Lead Senior Product
Manager of IT, CSX

Red Hat OpenShift

Everpure

ARO

AKS