



Portworx Enterprise in the Cloud

Improve application reliability with container-granular storage, data availability, data security, backup, and disaster recovery, while slashing your cloud container storage costs in half

Block storage in the cloud is convenient and easy to provision. But it was designed to provide storage to VMs, not Containers. In highly dynamic Kubernetes environments, this leads to problems such as limited density of volumes per host, slow and unreliable failover, environmental lock-in, and ultimately increased costs.

Portworx Enterprise is a software-defined container storage platform that uses your cloud block storage to provide containerized applications with scale-out storage, data availability, data security, backup, and disaster recovery running in a single cloud or across multiple clouds. Portworx has helped dozens of Global 2000 companies such as Carrefour, Comcast, GE Digital, Lufthansa, T-Mobile, and SAIC run containerized data services in production.



We're running 200-300 stateful pods per host. Because of these densities enabled by Kubernetes and Portworx, we're easily saving 60-90% on our compute costs."



– Sergey Pronin, SaaS Ops Software Engineering Manager at Aurea Software
portworx.com/customers/aurea


Why Global 2000 customers run Portworx in the cloud

- Cut cloud storage costs in half
- Increase container density beyond cloud provider limits
- Multi-availability zone and region availability and data protection
- Consistent storage experience across clouds
- Integration with cloud auto-scaling groups (ASG)
- Data availability, data security, backup, and disaster recovery

Slash cloud storage costs in half with PX-Autopilot

While Kubernetes automates the deployment of applications, enterprises must also be able to automate their underlying cloud infrastructure to ensure they have sufficient compute and storage available to scale applications. Despite the cloud's promise of *pay-for-use*, the reality is that enterprises must overprovision block storage, often by 2-3x, and manage the additional complexity to achieve the required application performance.

PX-Autopilot enables you to cut your cloud storage costs in half by intelligently provisioning storage only when it is needed, delivering on the cloud promise of *pay-for-use* rather than *pay-when-provisioned*. By automatically scaling on demand at an individual volume level, or even an entire cluster, you can save money and avoid application outages. Save further with policy driven storage tiering to reduce the time wasted by manually configuring your applications. All with native integration of the major cloud providers' block storage—on Amazon EBS, Google Persistent Disk, and Azure Managed Disks—as well as VMware vSphere.



[Learn more at portworx.com/cloud](http://portworx.com/cloud)

Unlock full performance

In addition to driving cost savings through auto-provisioning, Portworx Enterprise also allows enterprises to achieve higher performance at a lower cost.

IOPS in the cloud is correlated with the size of a cloud block device. For performance sensitive apps, the result is that enterprises often deploy bigger drives than are needed in order to achieve their performance targets. By using Portworx storage virtualization, a user can achieve a 70% EBS savings for the same level of performance for multiple applications.

For example, compare the cost difference of four (4) Kubernetes applications that each require 1TB of storage capacity and a max of 10,000 IOPS each.

	Kubernetes Volume Capacity	Total EBS Capacity	Max IOPS per Kubernetes Volume	Annual Cost	Cost Saving
AWS EBS only	3.3 TB	13.2 TB	10,137	\$16,220	
Portworx on top of AWS EBS	1 TB	4 TB	12,288	\$4,915	70%

Escape density jail

With Portworx, these raw storage cost savings combine with increased density to dramatically lower overall infrastructure costs. When running stateful containers on the public cloud, each cloud provider limits the number of block devices that can be attached to a given VM. For example, on Amazon, this number is 40 EBS volumes. That means even if a VM has enough compute capacity to serve more than 40 stateful containers, you will have to provision an additional VM beyond this limit if you need an additional volume. This overprovisioning drives up compute costs, not to mention the operational overhead of running more VMs than is strictly necessary. Because Portworx virtualizes cloud block volumes, you can easily run 1,000 or more stateful pods per host, eliminating VM-overprovisioning and further reducing cloud infrastructure costs.

Application protection and business continuity

Important applications running on Kubernetes require a host of data protection and disaster recovery options. Portworx provides a range of data protection options on top of your cloud block storage from static backups to zero-RPO and < 1 minute RTO for applications running in the cloud. In all instances, rather than focusing only on data protection, Portworx ensures that application data, Kubernetes objects, and application configuration are backed up, moved, and restored together - dramatically reducing the time it takes to recover from failures. Relying only on cloud block storage alone for data protection is expensive, limited to specific availability zones or regions, and only solves the data backup portion of recovery, leaving you to manage application recovery as a separate process.

Improved failover speed in the cloud

Kubernetes failover using cloud block devices relies on detaching block devices from one host and reattaching them to another. This process is fragile and slow, resulting in application downtime. Failover with Portworx up to 3x faster than using cloud block storage directly.¹ In order to achieve the same failover speed in the cloud without Portworx requires three times as much compute.

Configuration	Failover (sec)	Failover Improvement (%)	Number of records written	% Improvements
MongoDB + GKE Persistent Disk	180		15,343	
MongoDB + Portworx	45	300%	49,014	319%

¹<https://portworx.com/kubernetes-failover-mongodb>