

Data resiliency for PostgreSQL

Crunchy Data PostgreSQL on Red Hat OpenShift

Highlights

Deploy reliable, scalable, and highly available persistent storage for critical PostgreSQL applications.

Extend failover protections for important enterprise PostgreSQL data across AWS Availability Zones.

Choose cloud provider-supplied, general-purpose, or direct-attached storage with Red Hat OpenShift Data Foundation replication.

Support on-premise, public cloud, or hybrid cloud deployments with a single, software-defined storage solution.

Resiliency with performance for PostgreSQL

Organizations are increasingly deploying database and data cache workloads in containers.¹ Red Hat® OpenShift® Platform Plus includes [Red Hat OpenShift Data Foundation](#), a cluster data management solution that offers scalable and available persistent storage for cloud-native applications based on PostgreSQL. Container-native, software-defined storage lets application and development teams dynamically provision persistent volumes (PVs), quickly scaling or deprovisioning storage on demand. Moreover, OpenShift Data Foundation can provide business continuity with resilience across multiple cloud provider availability zones, maintaining performance comparable to cloud providers' storage offerings that only run within a single availability zone.

Crunchy Data PostgreSQL on Red Hat OpenShift

With many PostgreSQL database deployments taking place in cloud environments, software-defined storage increasingly plays a critical role in both database performance and resilience for enterprise applications. As a full-featured, tier-1 relational database management system (RDBMS)², PostgreSQL is growing rapidly³. DB-Engines ranking tool for popular database management systems, regularly shows PostgreSQL as one of the most popular open source RDBMS over the last several years⁴.

Enterprises deploying PostgreSQL for critical applications in cloud environments need more than a full feature set. They need both robust performance and resilience for critical data. Ultimately, software-defined storage is essential for achieving both. Like most databases, organizations have multiple methods to provide resiliency, including:

- ▶ **Application-layer resiliency.** With PostgreSQL replication, the database itself manages database resiliency. While this approach offers application awareness, it introduces greater complexity, requiring more in-depth PostgreSQL knowledge (or third-party software) to manage data replication. Moreover, any configured resilience applies only to PostgreSQL. Other applications and databases would need their own resilience methods, adding complexity and duplication.
- ▶ **Storage-layer resiliency.** In contrast, storage-layer resiliency relies on underlying storage services to manage data replication. This approach is usually more straightforward than implementing application-layer resiliency and provides potentially greater flexibility. Storage-layer resiliency protects not only PostgreSQL databases but also other types of databases and applications. It can also offer more control over replication fine-tuning.

¹ Pulse, sponsored by Red Hat. "State of workloads adoption on containers and Kubernetes", Nov. 2021.

² PostgreSQL features ACID properties (atomicity, consistency, isolation, and durability) and primary and unique indexes, updatable views, triggers, foreign keys (FKs) and even stored procedures (SPs).

³ Maxwell, John. "The Growth of PostgreSQL as a Tier-1 RDBMS," Database Trends and Applications, 10 June 2020.

⁴ "DB-Engines Ranking," DB-Engines, accessed 6 Dec. 2022.

Red Hat testing of the latest releases demonstrated storage performance comparable to cloud-native storage, even while OpenShift Data Foundation provided resilience across three Amazon Web Services (AWS) Availability Zones.

Crunchy Data PostgreSQL

PostgreSQL is a popular open source, object-relational database system with more than 20 years of continuous development. For testing, Red Hat engineers chose Crunchy Data PostgreSQL. Crunchy Data provides commercial support for PostgreSQL on a subscription basis, ensuring that enterprises of all sizes have access to certified software packages, updates, bug fixes, security patches, and 24x7x365 technical support from PostgreSQL experts. Crunchy Certified PostgreSQL is a trusted, commercially supported, and Common Criteria EAL 2+ certified distribution of open source PostgreSQL. Crunchy PostgreSQL for Kubernetes is a containerized PostgreSQL deployment that uses the operator pattern for Kubernetes and has achieved the autopilot capability level as part of Red Hat OpenShift Operator Certification⁵.

Crunchy Data helps enterprises benefit from the power and efficiency of PostgreSQL for critical applications through its suite of open source products and services, offering:

- ▶ More secure and high-availability PostgreSQL deployments.
- ▶ Elastic, hybrid cloud PostgreSQL solutions on all infrastructures.
- ▶ Geospatial, big data, and artificial intelligence (AI) architectures backed by PostgreSQL.
- ▶ Certified PostgreSQL installations and automated compliance verification.

Red Hat OpenShift Data Foundation

OpenShift Data Foundation offers more reliable and scalable persistent storage for cloud-native applications like PostgreSQL running in the cloud. It provides agile, scalable, portable, and highly available storage that can be provisioned and deprovisioned on demand. Application teams can dynamically provision PVs for many workload categories. The platform offers:

- ▶ *Agility* to streamline application and development workflows across hybrid cloud environments.
- ▶ *Scalability* to support emerging data-intensive workloads.
- ▶ *Portability* to allow simple data placement and access across cloud environments.

OpenShift Data Foundation features a software-defined storage platform based on Ceph®, which supports the needs of modern stateful applications. The use of the Kubernetes orchestration framework and Kubernetes operators makes OpenShift Data Foundation less complex and easier to install. Operators are software extensions to Kubernetes that use custom resources to automate and manage applications and their components.

Storage-based resiliency options

Within a cloud-based platform, organizations have choices for configuring and deploying storage-based resiliency. These choices can have ramifications for both performance and cost. Additionally, public cloud customers can choose between general-purpose storage classes or


⁵ [“Crunchy PostgreSQL for Kubernetes 4.2 Receives Red Hat OpenShift Operator Certification.”](#)
Crunchy Data, 10 Feb. 2020.

higher-performance, direct-attached storage volumes for their PVs. These storage options typically limit recovery from data failures to within a single AWS Availability Zone, which may not satisfy application requirements.⁶

In contrast, adding OpenShift Data Foundation to AWS storage volumes can provide data failover protection across multiple AWS Availability Zones— independent of the cloud-provider storage class selected. Red Hat testing has shown that this additional resilience can be accomplished while providing consistent performance for small databases. Table 1 summarizes the advantages and disadvantages of different AWS instance and storage classes, both with and without OpenShift Data Foundation.

While Elastic Block Store (EBS) general-purpose (gp2) offers failover within a single AWS Availability Zone, OpenShift Data Foundation adds automatic failover for AWS instances with direct-attached storage—resulting in additional performance and resiliency. AWS provides no storage failover options across multiple Availability Zones. In contrast, OpenShift Data Foundation provides automatic failover across multiple Availability Zones, while ensuring performance for applications.

Table 1. Red Hat OpenShift Data Foundation has performance, failover, and cost implications for single and multiple AWS Availability Zones

		 Red Hat OpenShift Data Foundation			
AWS instance		EBS gp2	EC2 I3en	EBS gp2	EC2 I3en
Single Activity Zone					
Performance database failover cost		☆☆☆	☆☆	☆☆	☆☆☆
		auto/free	manual/\$	auto/free	auto/free
		\$\$	\$\$	\$\$\$	\$\$\$
Multiple Activity Zones					
Performance database failover cost		--	--	☆☆	☆☆☆
		manual/\$	manual/\$	auto/free	auto/free
		\$\$	\$\$	\$\$\$	\$\$\$

⁶ For example, Amazon Web Services (AWS) Elastic Block Storage (EBS) general-purpose (gp2) storage classes as well as AWS direct-attached storage do not offer failover across AWS Availability Zones.

Performance testing

To evaluate the performance of different software-defined storage options, Red Hat engineers used the [Sysbench benchmark suite](#) to load a Crunchy Data PostgreSQL cluster with both small (20GB) and large (120GB) databases. AWS instances and storage volumes used are shown in Table 2.

Table 2. Sysbench test configuration


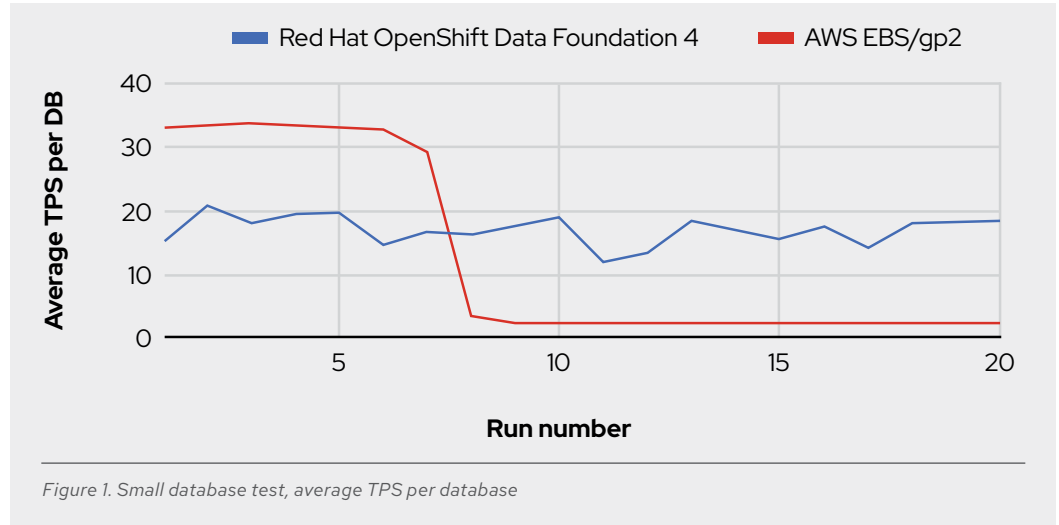
	Small database (20GB)	Large database (120GB)
Master nodes	3x M5.xlarge instances	
Compute nodes	3x M5.xlarge instances (Crunchy Data PostgreSQL)	
Storage nodes	3x M5.4xlarge (Red Hat OpenShift Container Storage 4.2)	3x i3en.2xlarge (Red Hat OpenShift Container Storage 4.3)
Storage devices	3x 2TB EBS gp2 volumes per node	2x 2.3 TB direct-attached NVMe solid-state drives (SSDs) per node

Small database tests

For the small 20GB database tests, each PostgreSQL pod requirement specification included one vCPU and 3GB of memory. Each OpenShift Container Platform compute node held 12 PostgreSQL pods and another set of 12 pods running Sysbench. Test runs compared systems using AWS EBS gp2 volumes against the same systems with OpenShift Data Foundation running on the EBS gp2 volumes.

Testing showed that when PostgreSQL is backed directly by EBS gp2 PVs, latency grows, and performance drops dramatically due to the [gp2 credit burst calculation](#) for small volumes. In contrast, the performance in terms of transactions per second (TPS) was consistent when using OpenShift Data Foundation running on the same EBS gp2 volumes (Figures 1 and 2).⁷

⁷ Each Sysbench run was 10 minutes, with a 75% read and a 25% write workload ratio.

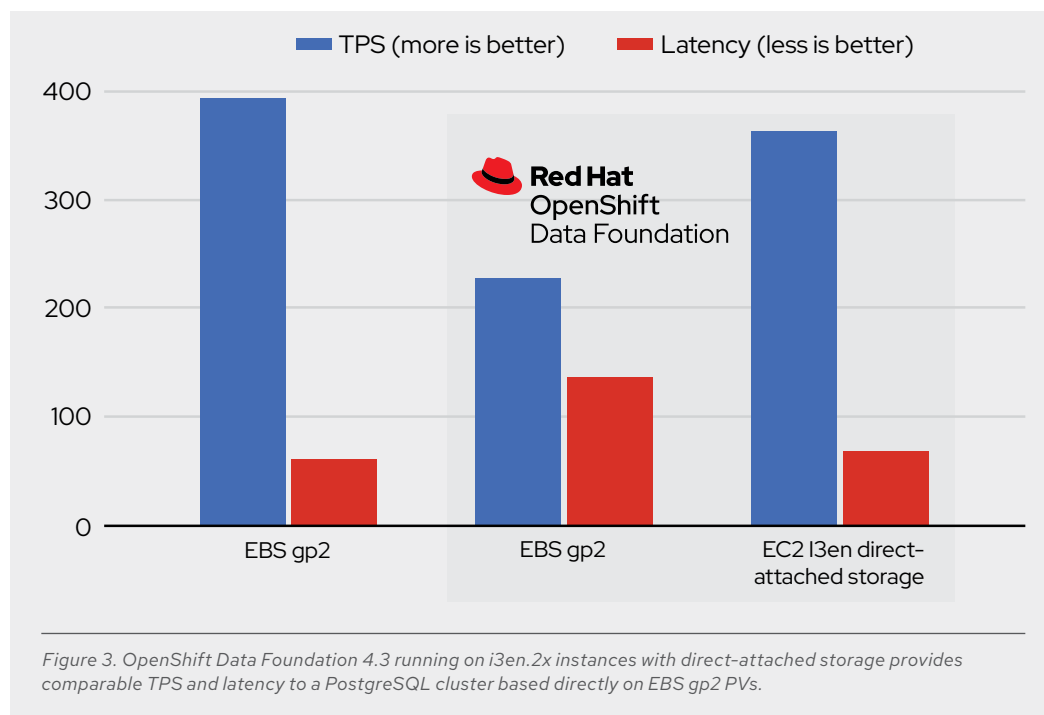


Large database tests

For the large 120GB database tests, each PostgreSQL pod requirement specification included eight vCPUs and 32GB of memory. Each OpenShift Container Platform compute node held a single PostgreSQL pod and a single Sysbench pod.

In OpenShift Data Foundation 4.2, EB2 gp2 volumes form the basis of the cluster. As such, workload performance is necessarily lower than when using the EBS gp2 volumes directly (Figure 3). However, it is important to note that the performance shown for OpenShift Data Foundation 4.2 includes replication across three AWS Availability Zones while the EBS gp2 solution is only measuring performance from within a single Availability Zone.

OpenShift Data Foundation 4.3 includes support for direct-attached storage. In the third set of columns, OpenShift Data Foundation used direct-attached storage instances (i3en.2x instance store volumes) instead of EBS gp2 volumes. With direct-attached storage, OpenShift Data Foundation has comparable performance to that of a cluster based on EBS gp2 alone—while still providing resilience across three AWS Availability Zones.



Conclusion

Red Hat OpenShift Data Foundation provides a flexible storage platform for PostgreSQL databases. OpenShift Data Foundation demonstrated consistent performance when tested using Crunchy Data PostgreSQL pods based on EBS gp2 volumes. Moreover, this approach added the ability to support high-performance, direct-attached storage and provide replication across three AWS Availability Zones. This functionality gives those deploying PostgreSQL the flexibility they need to supply resiliency and performance that matches their most demanding database applications.



About Red Hat

Red Hat helps customers standardize across environments, develop cloud-native applications, and integrate, automate, secure, and manage complex environments with [award-winning](#) support, training, and consulting services.

[f facebook.com/redhatinc](https://facebook.com/redhatinc)
[@RedHat](https://twitter.com/RedHat)
[in linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)

North America
 1888 REDHAT1
www.redhat.com

**Europe, Middle East,
and Africa**
 00800 7334 2835
europa@redhat.com

Asia Pacific
 +65 6490 4200
apac@redhat.com

Latin America
 +54 11 4329 7300
info-latam@redhat.com

redhat.com
 #000000_0021

Copyright © 2022 Red Hat, Inc. Red Hat, the Red Hat logo, OpenShift, and Ceph are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.