

Never Lose MySQL Data

# Percona XtraDB Cluster On Amazon EC2

This solution brief outlines replication and multi-primary capabilities using a Percona XtraDB Cluster EC2 architecture built with Percona Server for MySQL and the Percona enhanced Codership Galera library.

### Best fit for:

- Amazon EC2 environments geared to maintaining uptime and high availability
- Industries with high-read environments (financial or healthcare)
- Companies with in-house or external dedicated database resources
- Applications with read-heavy workloads

## Summary

An Amazon EC2 environment running [Percona XtraDB Cluster](#) provides availability and data consistency. This architecture provides a strong foundation for more advanced deployments capable of surviving disaster scenarios.

It does add the expense of complexity compared to deploying and running a standard MySQL® environment.

Percona XtraDB Cluster's automated failover and recovery allow the database to continually service the application transparently with ProxySQL directing traffic to available nodes. Percona Monitoring and Management (PMM) provides advanced visibility into the environment.

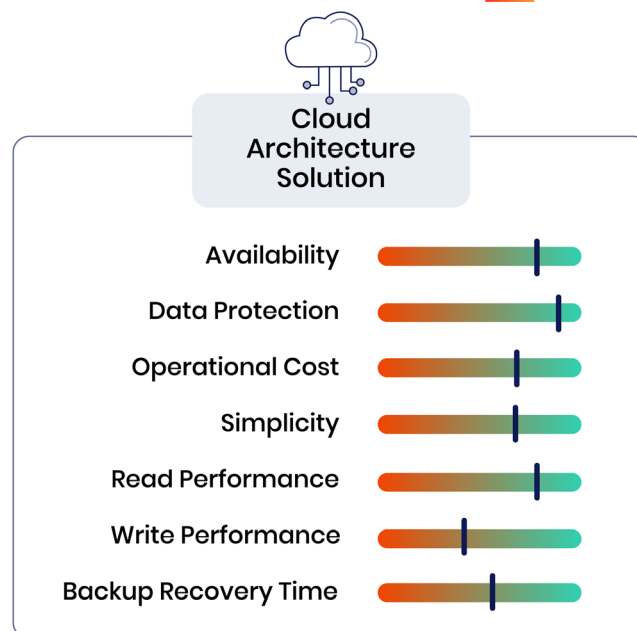
This document describes a proven standard Percona XtraDB Cluster EC2 architecture that is built on Percona Server for MySQL and the Percona enhanced Codership Galera library for replication and multi-primary capabilities.

## Use case

This solution deploys Percona XtraDB Cluster on Amazon Web Services (AWS) Elastic Compute Cloud (EC2) environment. This Percona XtraDB Cluster environment is appropriate for applications that need a high level of data consistency and availability. This architecture provides some read scaling capabilities but limited write scalability.

This environment has a high level of data consistency and high availability. Because Amazon EC2 SLAs provide uptime SLAs as high as 99.99% and Percona XtraDB Cluster's built-in election protocol and multi-primary capability provide auto recovery of nodes, this environment is meant for applications where there needs to be a guarantee of uptime and data integrity.

This solution is tailored to companies with in-house or external MySQL resources or expertise dedicated to the database environment.



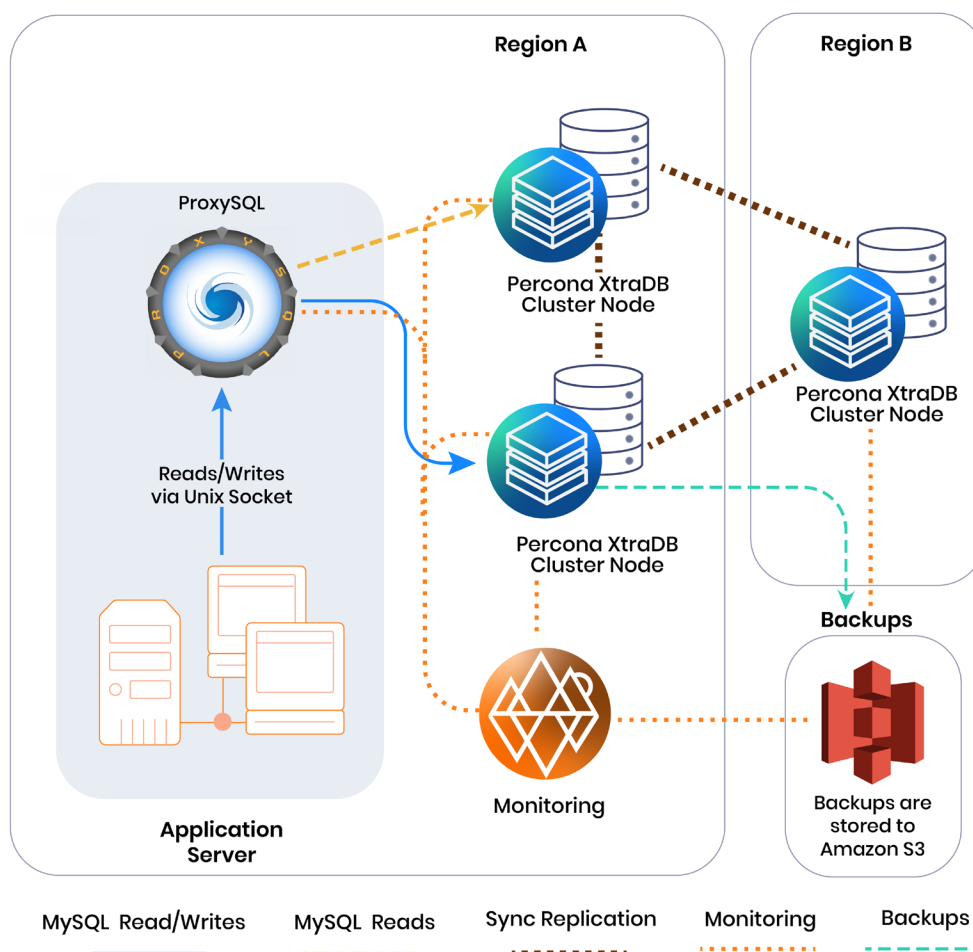
## Pros

- Quicker failover means higher application uptime with little to no application downtime, with consistent data across nodes.
- Failover is transparent to applications and doesn't affect application performance.
- Provides detailed data analytics from enhanced Percona software packages.
- Cloud solution lets you only pay for the infrastructure you need.
- Allows you to scale read and writes as needed.
- Allows you to distribute application workload across your entire system, minimizing updates and work on applications while maximizing flexibility.
- Allows you to keep your applications up during database maintenance.
- EC2 provides access and control of your operating system and infrastructure variables and configuration.

## Cons

- Solution is more complex, and DBAs need more understanding of multiple technologies.
- Write performance is traded-off for increased data protection with added write latency dependent on inter-datacenter latency.
- Not a good fit for applications with low write latency or high write throughput.

## ARCHITECTURE DIAGRAM



## Architecture

The design is straightforward: Your application servers connect to ProxySQL, which directs traffic transparently to the Percona XtraDB Cluster EC2 cluster. To mitigate losing an entire region, one node must be located in another region. Due to the synchronous replication mechanism, this impacts write latency based on the network latency between the two regions.

To allow for better workload analysis, deploying an additional EC2 instance running **Percona Monitoring and Management (PMM)** launched via the Amazon Marketplace is recommended.

## Components

- Three EC2 instances for Percona XtraDB Cluster nodes\*
  - Use i3 instances with local, attached NVMe storage
  - Two instances are located in a single region within different Availability Zones. The third instance is located in another region.
- One EC2 instance running PMM launched via the Amazon Marketplace for monitoring and analytics\*
- ProxySQL running on each application node
- Physical backups taken with Percona XtraBackup and real-time binary log stream with mysqlbinlog
- S3 bucket for backups

\* Enabling encryption is recommended for AWS storage..

## Failover

Failover is handled automatically by Percona XtraDB Cluster through a quorum vote. ProxySQL automatically removes failed nodes from the active read pool or shifts the primary. This requires no changes to the application. As consistency is the primary function of Percona XtraDB Cluster, no data loss occurs when a node fails and ProxySQL redirects traffic.

## Disaster

In the event of a disaster in the primary region, a simple manual process can be initiated to use the secondary region without data loss, keeping in mind that this requires the infrastructure to be capable of serving the entire application.

## Testing backups

Backups are not automatically tested. It is strongly recommended that backups are frequently tested to validate the backup and restore process.

## Backups

Backups should be taken from one of the Percona XtraDB Cluster nodes at regular intervals. Percona XtraDB Cluster ships with mysqldump, which can make logical backups of the database environment, but the recovery time of a logical backup grows exponentially with the size of the dataset. Hence, it is strongly recommended to use the open source tool Percona XtraBackup, which takes physical backups of the dataset that are much faster to restore. To be able to do point-in-time recovery, it is highly recommended to back up binary logs using mysqlbinlog.

Backups should be encrypted, compressed, and uploaded to S3.

## Monitoring

For query analytics and time-based database performance information, the use of the open source tool Percona Monitoring and Management is strongly recommended. This should be installed on a third host using the Amazon Marketplace. This monitors OS and MySQL metrics and provide advanced Query Analytics.

## Percona can help

Managing your organization's database operations, on-premises or in the cloud, requires in-depth knowledge of potential issues plus diligent, dedicated practice. Being aware of the issues above helps protect your organization's data-based applications when migrating your database to the cloud. It also significantly enhances both performance and scalability to deliver a better user experience.

Percona Support services are accessible 24x7x365 online or by phone to ensure that your database installation is running optimally. We can also provide on-site or remote Percona Consulting for current or planned projects, or in emergency situations. Every engagement is unique, and we will work with you to create the most effective solution for your business.

**Percona Managed Services** can support your existing database infrastructure, whether it is hosted on-premises or at a colocation facility, or if you purchase services from a cloud provider or Database-as-a-Service provider.

To learn more about how Percona can help, and for pricing information, please contact us at **+1-888-316-9775 (USA), +44 203 608 6727 (Europe)**, or email us at [sales@percona.com](mailto:sales@percona.com).