



MySQL Scalability on Amazon RDS: Scale Out to Multiple RDS Instances

A ScaleBase Whitepaper

info@scalebase.com

www.scalebase.com

©2014 ScaleBase. All rights reserved

ScaleBase

Table of Contents

- Optimizing Scalability of Amazon RDS for MySQL2
- Options for MySQL Scalability3
 - Larger RDS Instances3
 - Provisioned IOPs3
 - Leverage Read Replicas3
 - Multiple MySQL Instances3
- Unlimited MySQL Scalability on RDS4
 - A Distributed Amazon RDS Database: What you keep, everything!5
- Summary6
- About ScaleBase7

Optimizing Scalability of Amazon RDS for MySQL

MySQL is a very popular a database choice on Amazon’s Relational Database Service (RDS). Thousands of applications leverage MySQL databases that run inside the Amazon Cloud. This frees IT organizations from hardware capital expenditures and saves IT from spending time on database maintenance and administration chores.

What can you do when a MySQL database needs to scale write-intensive workloads beyond the capabilities of the largest available Amazon RDS MySQL instance?

Let’s take a look. In a typical EC2/RDS set-up, users connect to app servers from their mobile devices and tablets, computers, browsers, etc. Then app servers (or web/cloud services) connect to an RDS instance and in some cases they might leverage some read-only replicas.



Typical AWS EC2/RDS Architecture

Options for MySQL Scalability

As an application becomes more popular it can expect an increasing number of users, more transactions, and more accumulated data. As an application adds new capabilities and experiences, user interactions can require and generate more data. The result of all this positive activity is that your MySQL database will inevitably begin to experience scalability pressures.

Broadly speaking, there are four options available to improve RDS scalability.

Larger RDS Instances

If you're not already using the maximum available RDS instance, you can always scale up – to larger hardware to leverage bigger CPUs with more compute power and more memory. But the largest available RDS instance is still limited, as revealed by the specifications of Amazon's largest RDS instance:

"Memory optimized db.r3.8xlarge Instance"

- 244 GB of memory
- 32 vCPUs
- 64-bit platform
- High I/O Capacity

Provisioned IOPs

You can get provisioned IOPs and higher throughput on the I/O level. However, there is a hard limit with a maximum instance size and maximum number of provisioned IOPs that you can buy from Amazon and you simply cannot scale beyond these hardware specifications.

Leverage Read Replicas

If your application permits, you can leverage read replicas to offload some read workloads from the master databases. But there are a limited number of replicas that you can utilize and this approach generally requires some modifications to your existing application. Unfortunately however, read-replicas don't help with write intensive applications.

Multiple MySQL Instances

Amazon offers a fourth option:

["You can implement partitioning, thereby spreading your data across multiple database Instances"](#)

However, Amazon does not offer any guidance or facilities to help you with this. "Multiple database instances" is not an RDS feature, and Amazon doesn't explain how to implement this idea.

In fact, when asked, this is the response on an Amazon forum:

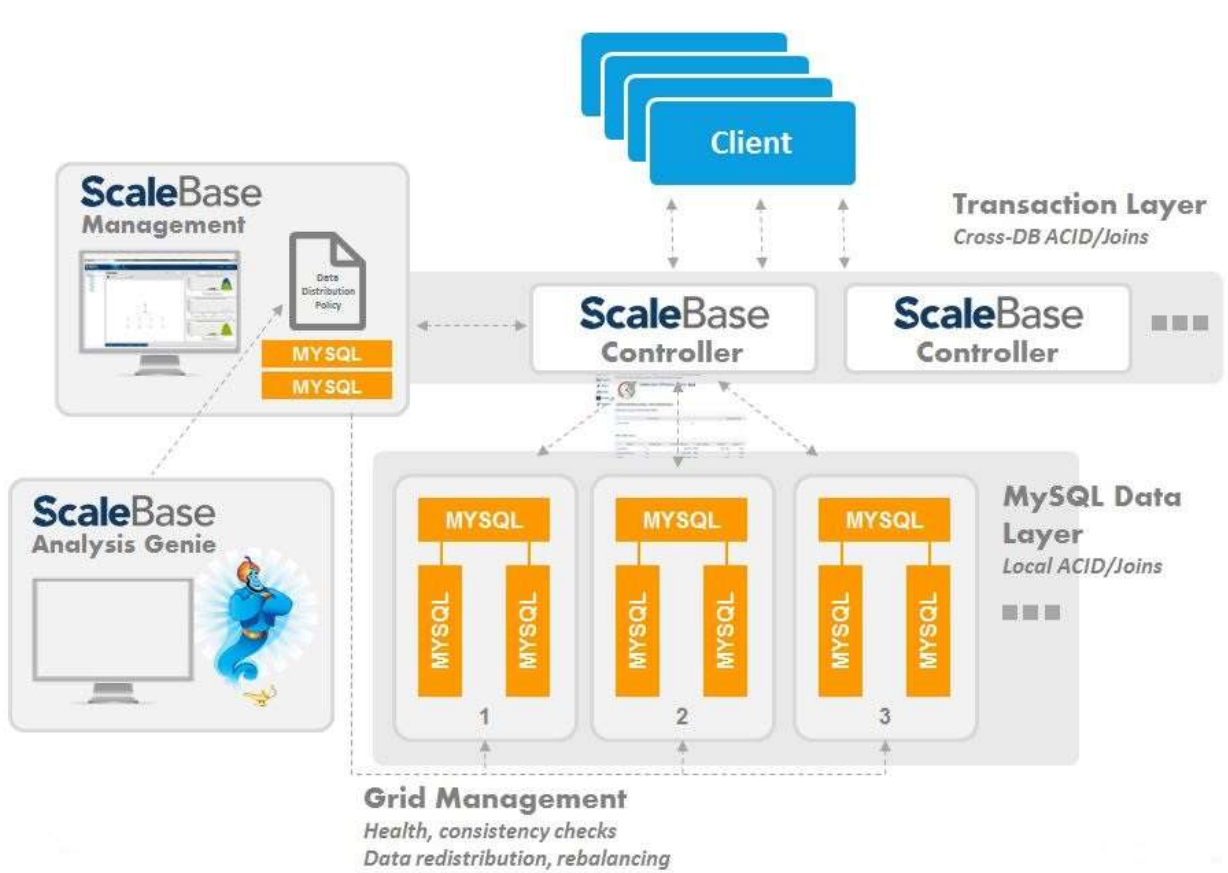
Q: *Is there any documents that describe the partition DB across multiple RDS? I need to use DB with more 1TB but exist a limitation during the create process, but I read in the any FAQ that you need to partition database, but I don't find any documents that describe it.*

A: *"DB partitioning/sharding is not an official feature of Amazon RDS or MySQL, but a technique to scale out database by using multiple database instances. The appropriate way to split data depends on the characteristics of the application or data set. Therefore, there is no concrete and specific guidance."*

Unlimited MySQL Scalability on RDS

ScaleBase is specifically designed to scale out a single MySQL RDS instance into multiple MySQL RDS instances. ScaleBase gives organizations the relational data integrity of MySQL combined with the scalability and flexibility of a modern distributed, multi-site database to support an unlimited numbers of users, larger data volumes and continuously growing transaction volume.

The ability to migrate an application from a hosted environment with a single growing database to a virtualized environment with smaller, more manageable data nodes gives companies agility, flexibility and competitiveness. ScaleBase was purpose built for cloud deployment – seamlessly distributing a single database into smaller ones to easily facilitate the move to the cloud.



ScaleBase Architecture

Critically, this is accomplished with no changes to your application code. Your application continues to “see” one database. ScaleBase does all the work of managing and enforcing an optimized data distribution policy that distributes data and workloads among multiple MySQL instances.

The result: now you can cost-effectively leverage multiple MySQL RDS instances to scale out write-intensive workloads to an unlimited number of users, transactions, and data.

A Distributed Amazon RDS Database: What you keep, everything!

ScaleBase create a distributed MySQL RDS database whose complexity is transparent to your applications. This has three very important advantages to you.

- 1. Keep your application, unchanged** – There is no change your application development life-cycle at all. You still use your existing development tools, frameworks and libraries. Application quality assurance and testing cycles stay the same. And, critically, you stay with an ACID-compliant MySQL environment.
- 2. Keep your RDS value-added services** – The value-added services that you rely on are all still available. Amazon will continue to handle database maintenance and updates for you. You can still leverage High Availability via Multi A-Z. And, if it benefits your application throughput, you can still use read replicas.
- 3. Keep your RDS administration** – The RDS monitoring and provisioning tools you rely on still work as they did before.

With your one large MySQL instance now split into multiple instances, you can actually use smaller, less expensive hardware and continue to see better database throughput and performance.



AWS EC2/RDS Architecture with ScaleBase

Summary

Amazon RDS is a tremendous service, but it doesn't offer solutions to scale write-intensive workloads beyond a single MySQL instance. When MySQL maxes-out on the available hardware, you're stuck. Amazon recommends scaling out your single instance into multiple instances for transaction-intensive apps, but offers no services or guidance to help you. This is where ScaleBase provides a unique solution.

ScaleBase gives you a simple and effective way to transparently work with multiple MySQL RDS instances, while removing all the complexities typically caused by "DIY" partitioning and with changes to your applications.

With ScaleBase you achieve unlimited MySQL scalability and continue to leverage the AWS/RDS ecosystem: commodity hardware and value added services like read replicas, multi A-Z, maintenance/updates and administration with monitoring tools and provisioning.

About ScaleBase

[ScaleBase](#) is a modern, distributed MySQL database optimized for the cloud which deploys in minutes and enables you to scale out to an unlimited number of users, data and transactions. It is a horizontally scalable database cluster built on MySQL that dynamically optimizes workloads and availability by logically distributing data across public, private and geo-distributed clouds.

ScaleBase allows you to design new applications for web scale, while requiring no changes to existing applications. ScaleBase is an authorized Oracle Partner, Rackspace Marketplace Partner, IBM Cloud marketplace partner and Technology Partner in the AWS Partner Network. Its customers span a wide range of industries including online gaming companies, digital media providers and social networks. The company is venture-backed and based in Newton, MA

ScaleBase is the only distributed database that provides a logical data distribution policy tuned and optimized for individual applications to ensure maximum database efficiency. ScaleBase delivers the following capabilities:

- **Simple, Elastic Scale Out** - ScaleBase deploys in minutes. With its policy-based data distribution, it enables scale out to an unlimited number of users, data and transactions. ScaleBase provides dynamic scale out with online database rebalancing.
- **Perfect for New Apps** - Reduce your time-to-market and design your application for web scale from the start with a modern cloud-based database, leveraging the power of MySQL as a scalable, elastic database that can grow with your business from day one.
- **No Changes Required to Existing Apps** - The existing application code or database requires no changes; allowing you to leverage your existing investments in MySQL skills, tools and infrastructure.
- **Database Availability and Resiliency** - Provides automatic failover during database failure, protecting users against downtime and delays, ensuring users & applications remain connected to their documents, data files and business systems at all time.
- **Hybrid Private/Public Cloud** - Distribute the database so sub-sets of the database can be split across multiple sites, private and public cloud infrastructure.

Additional Resources:

- ScaleBase software is available for free download [here](#).
- ScaleBase software is available on the AWS Marketplace [here](#).
- You can register for access to the ScaleBase Portal [here](#).
- You can learn more about ScaleBase [here](#).

