



Exploring Amazon RDS MySQL Second Tier Read Replica

AWS recently introduced Second Tier Replica for RDS MySQL. This feature is used to shift the load from primary master DB to the replica in first tier and prepare the application for handling extremely high level of read traffic.

Exploring Amazon RDS MySQL Second Tier Read Replica


Content covered in this white paper

- Steps to configure Multi-Tiered Amazon RDS MySQL Read replicas
- Second Tier Read Replica Deployment architectures

Amazon RDS MySQL version 5.6.12 on AWS WEST region was used.

Steps to configure Multi- Tiered Amazon RDS MySQL Read Replicas

Configuration steps for the following architecture is given in this post :



Introduction to Amazon RDS

...

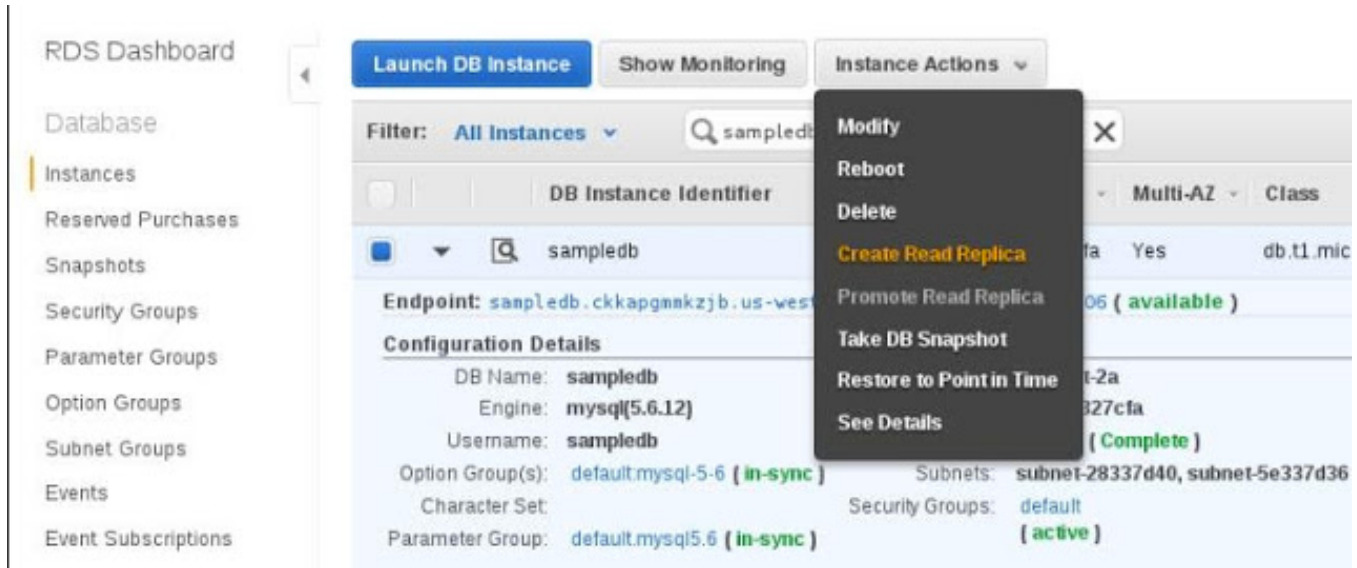
Amazon Relational Database Service (Amazon RDS) makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database administration tasks, freeing us up to focus on our applications and business. Amazon RDS provides us with six familiar database engines to choose from, including Amazon Aurora, Oracle, Microsoft SQL Server, PostgreSQL, MySQL and MariaDB.



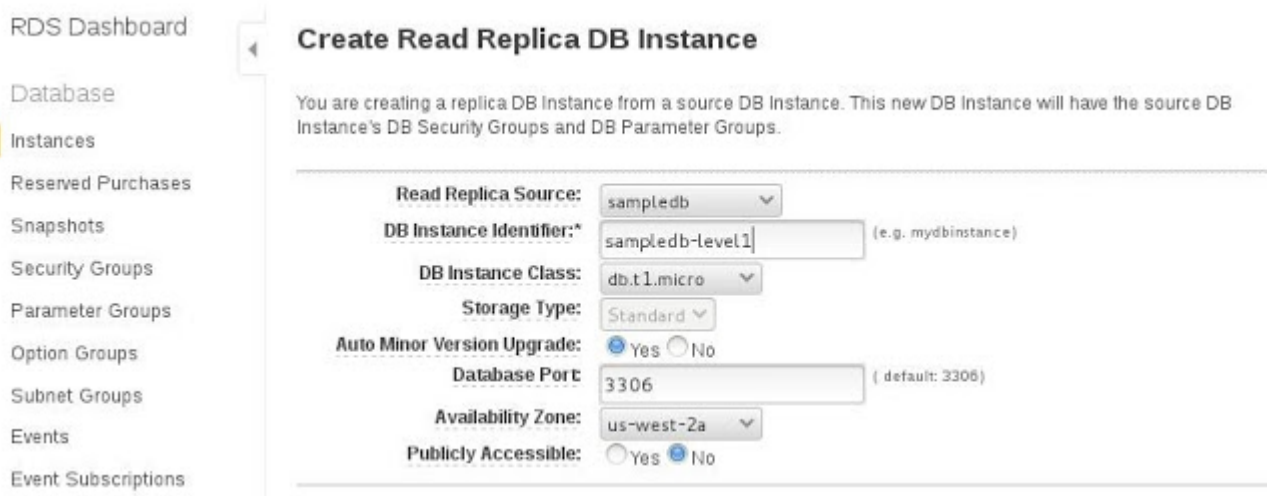


Step 1: Creating Read Replica from Master and Place at Tier 1:

To create RDS MySQL Read replica navigate to the dashboard of Amazon RDS, select the Amazon RDS MySQL Master(named"sampledb") and use the option of "Create Read Replica".



Name the newly created Amazon RDS Read replica as "sampledb-level1" and place it in Tier 1. The Tier 1 Amazon RDS MySQL read Replica can be created in same AZ of Master or in a different AZ for High Availability. When the Tier replica is placed in Different AZ , we should factor few extra milliseconds of latency during replication.



Explore the status from Master DB: Post successful creation of the Tier 1 Read Replica, We can see the Read Replica Id's from exploring the details of the Master DB. Illustrated in Below Screen shot :

Exploring Amazon RDS MySQL Second Tier Read Replica

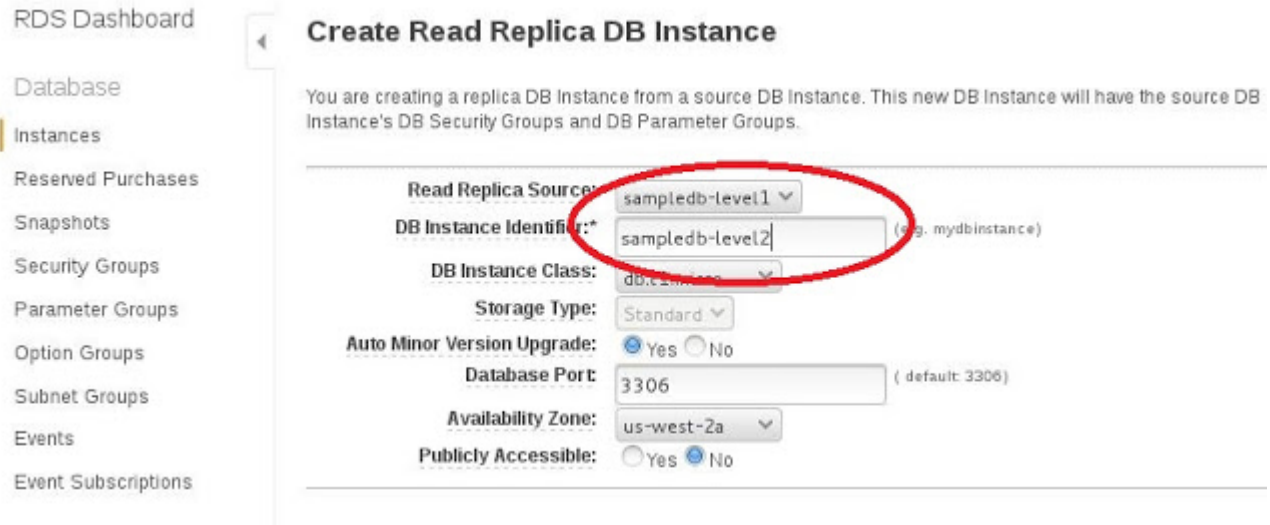
• • •

Configuration Details			Security and Network			Instance and IOPS		
DB Name:	sampledb		Availability Zone:	us-west-2a		Storage:	10 GB	
Engine:	mysql(5.6.12)		VPC ID:	vpc-92327cfa		Instance Class:	db.t1.micro	
Username:	sampledb		Subnet Group:	sample (Complete)		IOPS:	disabled	
Option Group(s):	default:mysql-5-6 (in-sync)		Subnets:	subnet-28337d40, subnet-5e337d36				
Character Set:			Security Groups:	default (active)				
Parameter Group:	default:mysql5.6 (in-sync)							
Availability and Durability: Read Master			Maintenance Details					
Replication State:	-		Minor Version Upgrade:	Yes				
Replication Error:	-		Maintenance Window:	sat06:12-sat06:42				
Multi AZ:	Yes		Backup Window:	08:13-08:43				
Secondary Zone:	us-west-2b							
Automated Backups:	Enabled (1 Day)							
Latest Restore Time:	2013 September 3 10:55:00 UTC+5:30							
Read Replica Ids:	sampledb-level1							

Explore the status from Tier 1 Read Replica : When we explore the Tier 1 Read Replica details, we will find it is pointing to the to Master DB. Illustrated in Below Screen Shot :

Configuration Details			Security and Network			Instance and IOPS		
DB Name:	sampledb		Availability Zone:	us-west-2a		Storage:	10 GB	
Engine:	mysql(5.6.12)		VPC ID:	vpc-92327cfa		Instance Class:	db.t1.micro	
Username:	sampledb		Subnet Group:	sample (Complete)		IOPS:	disabled	
Option Group(s):	default:mysql-5-6 (in-sync)		Subnets:	subnet-28337d40, subnet-5e337d36				
Character Set:			Security Groups:	default (active)				
Parameter Group:	default:mysql5.6 (in-sync)							
Availability and Durability: Read Replica			Maintenance Details					
Read Replica Source:	sampledb		Minor Version Upgrade:	Yes				
Replication State:	replicating		Maintenance Window:	sat06:12-sat06:42				
Replication Error:	-		Backup Window:	08:13-08:43				
Multi AZ:	No							
Secondary Zone:								
Automated Backups:	Enabled (1 Day)							
Latest Restore Time:	2013 September 3 10:20:00 UTC+5:30							

Step 2: Creating Second Tier Read Replica from Tier 1 Read replica : To create Second Tier Read replica navigate to the dashboard of Amazon RDS, select the Amazon RDS Tier 1 Read Replica (**sampledb-level1**) as the **source** and use the option of “Create Read Replica”.



Name the newly created Amazon RDS Read replica as "sampleddb-level2" and place it in Second Tier 2. The Second Tier -2 Amazon RDS MySQL read Replica can be created in same AZ of Tier 1/Master or in a different AZ for High Availability.

Explore the status from Tier 2 Read Replica : When we explore the Tier 2 Read Replica details, we will find it is pointing to the to Tier 1 - sampleddb-level1 as replication source. Illustrated in Below Screen Shot :




To Load Balance AWS RDS Read Replica's Refer this article

: <http://harish11g.blogspot.com/2013/08/Load-balancing-Amaon-RDS-MySQL-read-replica-slaves-using-HAProxy.html>


Second Tier Read Replica Deployment Architectures

One of the main complexity behind Multi-Level replication is that, if Tier 1 Read Replica if not properly architected/placed, it can be a single point of failure. Imagine a case where we need 4 Read Replicas for wer Master DB. We can take following approaches as illustrated below while designing wer infrastructure for this requirement.



In Architecture-1, Tier-1 Read replica is a Single Point of Failure. Instead if we split Tier-1 itself into two separate fleets it offers better availability than architecture-1. Since both Tier-1 RR put replication load on Master DB , we can by pass this using Hot Standby instance. Benefits of this approach is explained below in the best practice architecture.

Best Practice High Availability Architecture for Second Tier Read Replica:



Condition 1 - Master DB Failure : Hot Standby becomes new master and Tier-1 Read Replicas points automatically to new source

Condition 2- Tier-1 Read Replica Failure : Entire Tier-1 and Associated Tier-2 has to be recreated. The Alternate Active Fleet of Tier-1+Tier-2 will serve the requests for high availability.

Condition 3 - Tier-2 Read Replica Failure : Only the non performing Tier-2 Read Replica instance has to be recreated

Condition 4- AZ NW problem: In event AZ-3 is failed, Requests are served by alternate fleet of Tier-1+Tier-2 in another AZ(AZ-4).

Exploring Amazon RDS MySQL Second Tier Read Replica



Multi-AZ Hot Standby Implementation is a recommended best practice when it comes to Multi-tiered Read replica implementation

- It is recommended to create Tier 1 Read Replicas from Multi-AZ DB instance to offload read queries from the source master DB instance for high traffic sites. If the source instance of a Multi-AZ deployment fails over to the secondary, any associated read replicas will be switched to use the secondary as their replication source automatically. This model guarantees high availability.
- Also when we initiate the creation of a Tier 1 read replica, Amazon RDS takes the DB Snapshot of wer Standby DB instance (instead of Source DB) and begins replication. This model saves I/O suspension on wer source DB during the snapshot process

Other Points to Note:

Note :

To Load Balance AWS RDS Read Replica's Refer this article : <http://harish11g.blogspot.com/2013/08/Load-balancing-AWS-RDS-MySQL-read-replica-slaves-using-HAProxy.html>

P1) Circular replication are not allowed in this tiered replica creation process.

P2) Third Tier cannot be created. Actually in production, very rarely we need third tier and it is not important feature

P3) When the "X" Tier replica is placed in Different AZ , we should factor few extra milliseconds of latency during replication.

P4) Before a DB instance (Master or Tier 1) can serve as a replication source, we must enable automatic backups on the source DB instance by setting the backup retention period to a value other than 0. This requirement does not apply to second tier replica as they are not source DB instance for another read replica. So create "Read Replica" option will appear only when the Backup Retention is set to minimum of 1 day for any level Read Replica creation. When we "create Read Replica" by default RR is created with Backup Retention set to 0. Use the modify option Illustrated in below screenshot and set the Backup Retention Period. :

The screenshot displays the configuration page for a new Amazon RDS MySQL instance. The following settings are visible:

- DB Instance Identifier:** sampledb-level1
- DB Engine Version:** MySQL 5.6.12 (default)
- DB Instance Class:** db.t1.micro
- Auto Minor Version Upgrade:** Yes (selected)
- Allocated Storage:** 10 GB (Minimum: 5 GB, Maximum: 3072 GB)
- Use Provisioned IOPS:** No
- Provisioned IOPS:** (Empty field)
- Parameter Group:** default.mysql5.6
- Security Group:** default (VPC) (selected)
- Option Group:** default.mysql-5-6
- New Master Password:** (Empty field)
- Backup Retention Period:** 1 days
- Backup Window:** Start Time: 08:13 UTC, Duration: 0.5 hours
- Maintenance Window:** Start Time: Saturday 06:12 UTC, Duration: 0.5 hours
- Apply Immediately:** Checked

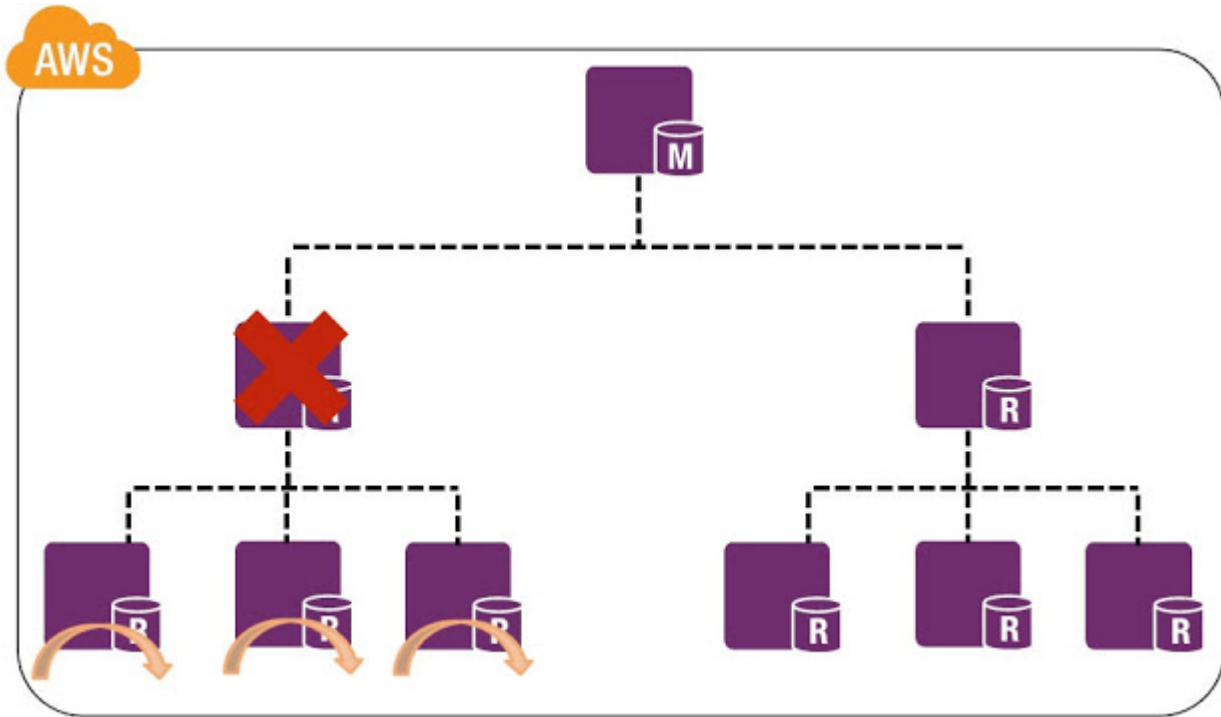
A "Continue" button is located at the bottom right of the configuration area.

P5) We have used t1.micro as RDS DB for article explanatory purpose. For production use cases please use proper

Exploring Amazon RDS MySQL Second Tier Read Replica

• • •

instance types after proper capacity planning. P6) AWS has released Parallel replica creation process, where we can create multiple Tier 1 and Tier 2 replicas in parallel. Since we no longer need to wait for one replica creation before starting the next one, it becomes easy to create multiple RR quickly. Without this feature, it would take hours to create a large Multi-tiered Replica setup. P7) If a replica lags too far behind for wer environment, the normal practice is to scale up or consider deleting and recreating the read replica. Imagine we have architecture where there are two Tier 1 and three Second Tier RR as illustrated in below architecture:



Now since Tier -1 RR is lagging , we are planning to delete and recreate the same. When we delete the Tier-1 RR, All the Second Tier Read Replica's will now become Standalone , Single AZ DB's. Once we have recreated the Tier -1 RR again and managed to retain the original end point as well, we cannot re-point these Second Tier RR to Tier -1 again. We need to recreate all 3 second tier from the Tier -1 again. For a application with heavy DB dependency and read traffic, it means 1/2 of the fleet is now down. This can lead to uncomfortable performance situation. This is not a ideal condition and AWS RDS team can take this in their Road Map. For such cases, it is better to create a new Tier -1 and Second tier fleet first , update wer LB/App configs and then delete the old lagging fleet

P8)Why it is better to have Multi-tiered Read Replica as the last resort in wer architecture ?

8.a)Currently 30 Read Replica can be created overall for a master in two tiers. 30 Read replica is more than sufficient and usually turns out to be costly architecture approach. Use this approach only for cases, which demand heavy read and when application code cannot accommodate changes and are highly DB dependent.

8.b)It is recommended in DB world to stay away from Multi-level replication as much possible. Wer architecture will be much simpler with one master and "X" replica slaves, rather than having tiered replica's. As we observed in above deployment architectures, the second tier replica slave will be a trouble to manage in event of replication delay, crashes and network problems affecting the Tier-1 RR or the Master DB.

8.c)In case wer application code can be redesigned, it is recommended to take following approaches before resorting to Second Tier replicas architecture

- Functional partition the RDS MySQL with Hot Stand By and Read Replica's

Exploring Amazon RDS MySQL Second Tier Read Replica



- Re-balance the DB load by using alternate data stores provided by AWS like DynamoDB, ElastiCache, CloudSearch etc.

If the above methods does not work for we, take the Second tier Read replica approach.