Guided Labs: Creating and restoring RDS backups using Snapshots

Description

Amazon RDS (Relational Database Service) allows users to create and restore backups using snapshots, providing a robust and reliable solution for data protection. Creating RDS backups involves taking a point-in-time snapshot of your database instance, capturing its entire state, including data, configuration, and transaction logs. These snapshots serve as a restore point, enabling you to recover your database to a specific moment if data is lost or the system encounters issues.

Restoring RDS backups from snapshots is a straightforward process. Users can initiate restoration by specifying the desired snapshot and creating a new RDS instance. This effectively recreates the database at the chosen time, ensuring data consistency and integrity. The ability to create and restore backups using snapshots is a critical feature for database administrators, offering peace of mind by providing a reliable method for data recovery and minimizing downtime in the event of unforeseen circumstances or data loss.

In this lab, you will gain hands-on experience creating and restoring RDS backups using snapshots. The lab typically guides participants through the step-by-step procedures of initiating a snapshot, understanding snapshot retention policies, and restoring a database to a specific point in time.

Objectives

In this lab, you will:

- Understand the importance of RDS backups for data integrity and disaster recovery
- Learn how to create snapshots of your RDS database instances
- Restore an RDS instance from a snapshot backup.
- Understand the differences between automated backups and manual snapshots.

Lab Steps

Creating a DB snapshot

1. Navigate to the search bar, type “RDS”, and click to open the RDS Dashboard.

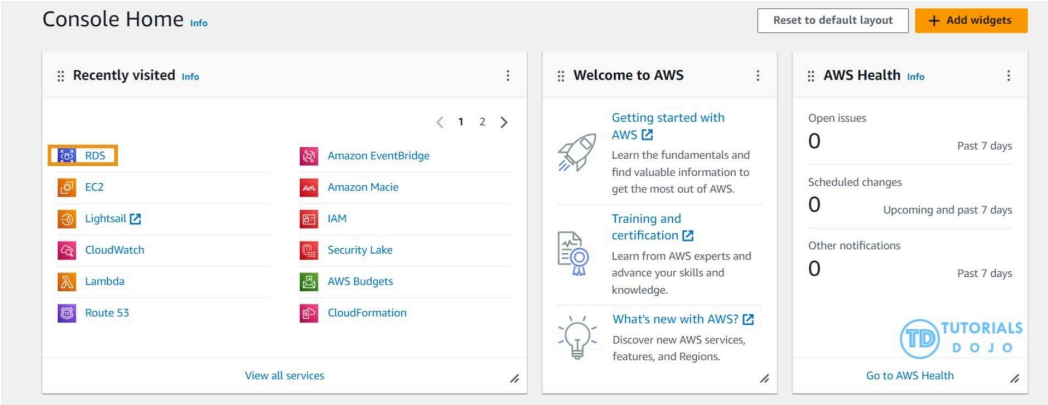

2. In the left-hand navigation panel, click on “Databases.”

3. From the list of available DB instances, select the one for which you wish to create a snapshot.

4. Next, locate the “Actions” menu and choose “Take snapshot.”

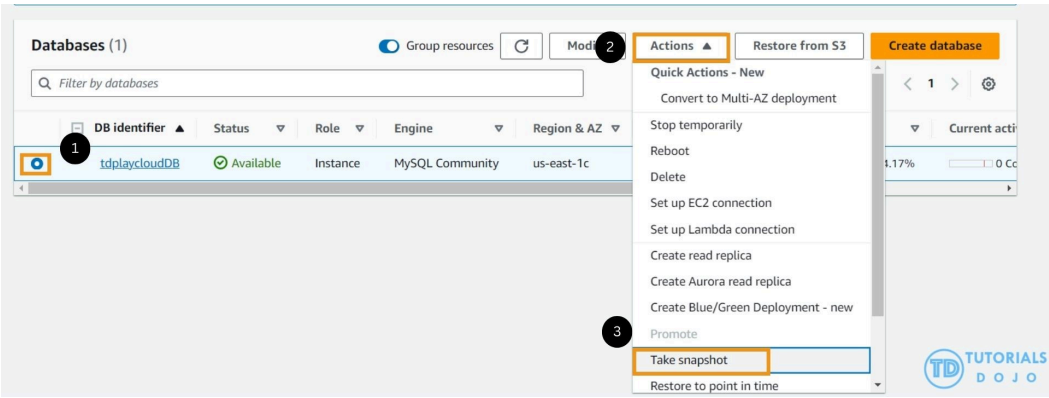

5. A new “Take DB snapshot” window will appear on your screen.A new “Take DB snapshot” window will appear on your screen.

6. In the “Snapshot name” field, input a name for your snapshot. This name should be descriptive and relevant to the content of the snapshot.

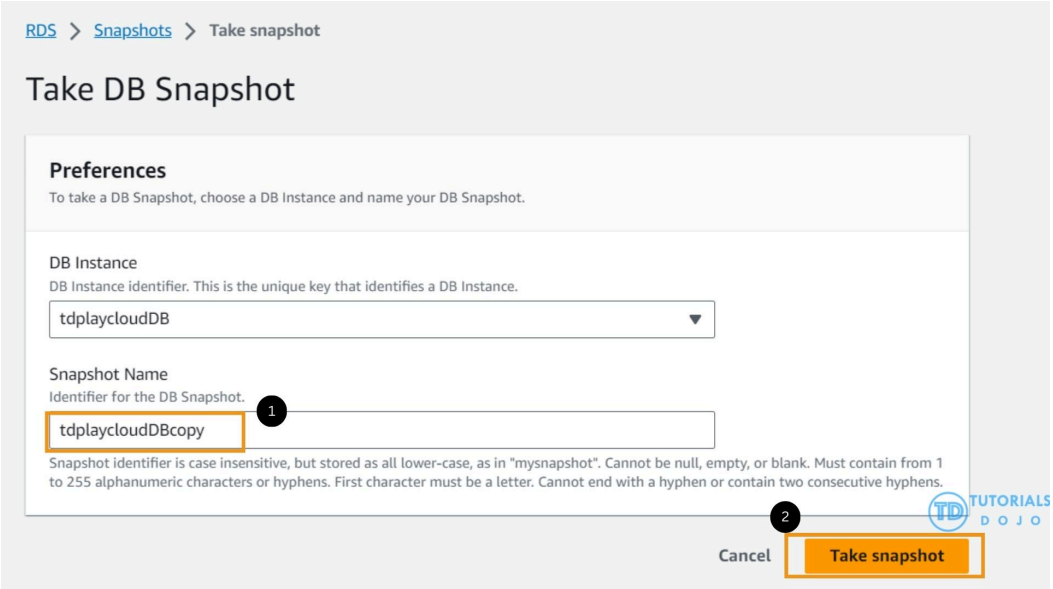

Restoring a DB instance from a DB Snapshot

1. In the left-hand navigation pane, click on “Snapshots.”

2. Select the specific DB snapshot you intend to restore.

3. Click the “Actions” button, then choose “Restore snapshot.” This will take you to the “Restore snapshot” page.

3. With the instance still selected, choose “Actions” > “Instance Settings” > “Change Instance Type.” Note that this option will be grayed out if the instance state is not “stopped.”

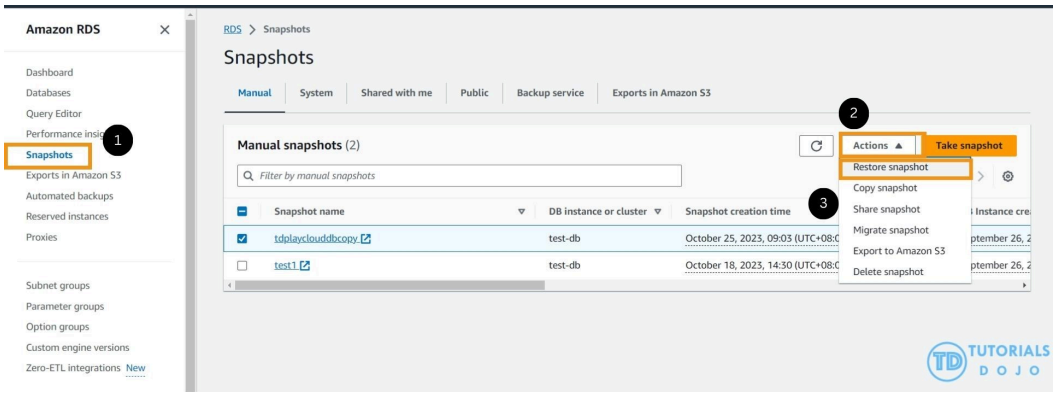

4. In the “DB instance settings” section, you can leave the default settings for the DB engine and license model as is. These settings are usually set according to the original snapshot.

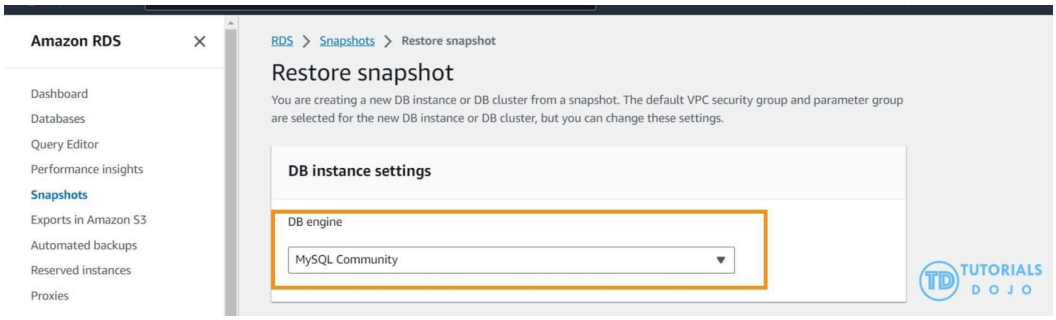

5. In the “Availability & durability” section, decide whether to create a standby instance in a different Availability Zone. For this lab, we will create a standby instance.

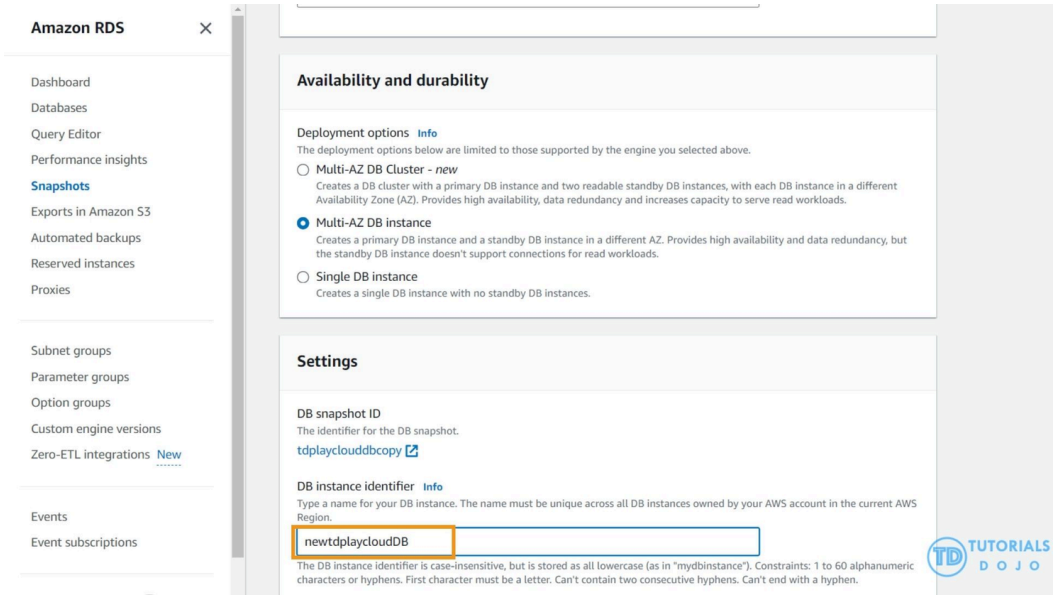

6. Under “Settings,” provide a unique name for the restored DB instance in the “DB instance identifier” field. For instance, you can use something like “newtdplaycloudDB.” If you’re restoring from a DB instance you’ve previously deleted, you can reuse its name.

7. Select the desired DB instance class. In this tutorial, you can choose “Burstable classes,” including “t” classes, and then pick “db.t3.small.”

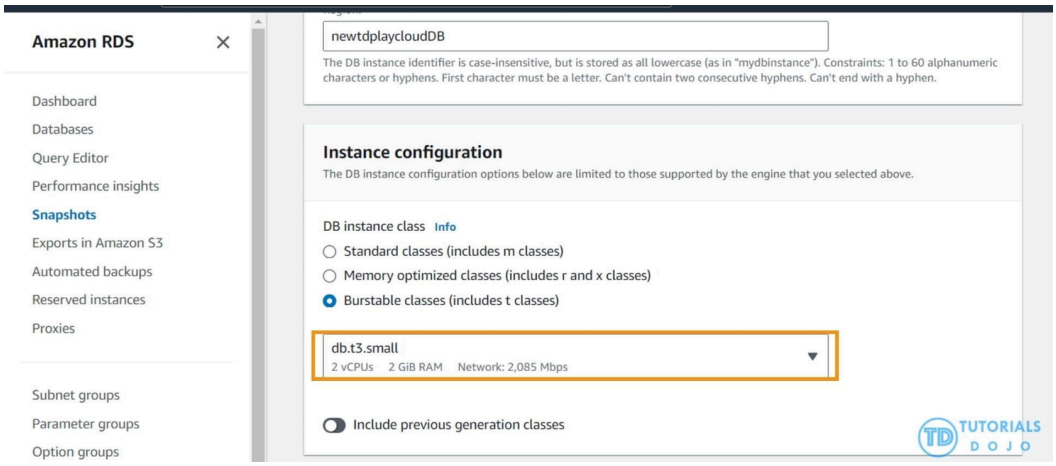

8. In the “Connectivity” section, you can keep the default settings for the following:

- Virtual private cloud (VPC)
- DB subnet group
- Public access
- VPC security group (firewall)

9. For “Encryption,” use the default settings. If the source DB instance from the snapshot was encrypted, the restored DB instance will also be encrypted, and you can’t change this.

10. Expand the “Additional configuration” section at the bottom of the page.

11. Under “Database options,” do the following:

- **DB parameter group:** For this lab, stick with the default parameter group.
- **Option group:** In this lab, the default option group is recommended. In some cases, you may need to select an option group that matches the options used by the original DB instance. **Note:** If you’re restoring from a DB snapshot of a DB instance with specific options, ensure the chosen option group includes those same options.
- Enable deletion protection by checking the “Enable deletion protection” box.

12. Once configuring all the settings, click “Restore DB instance.” You will see the restored DB instance on the Databases page with the status of “Creating.”

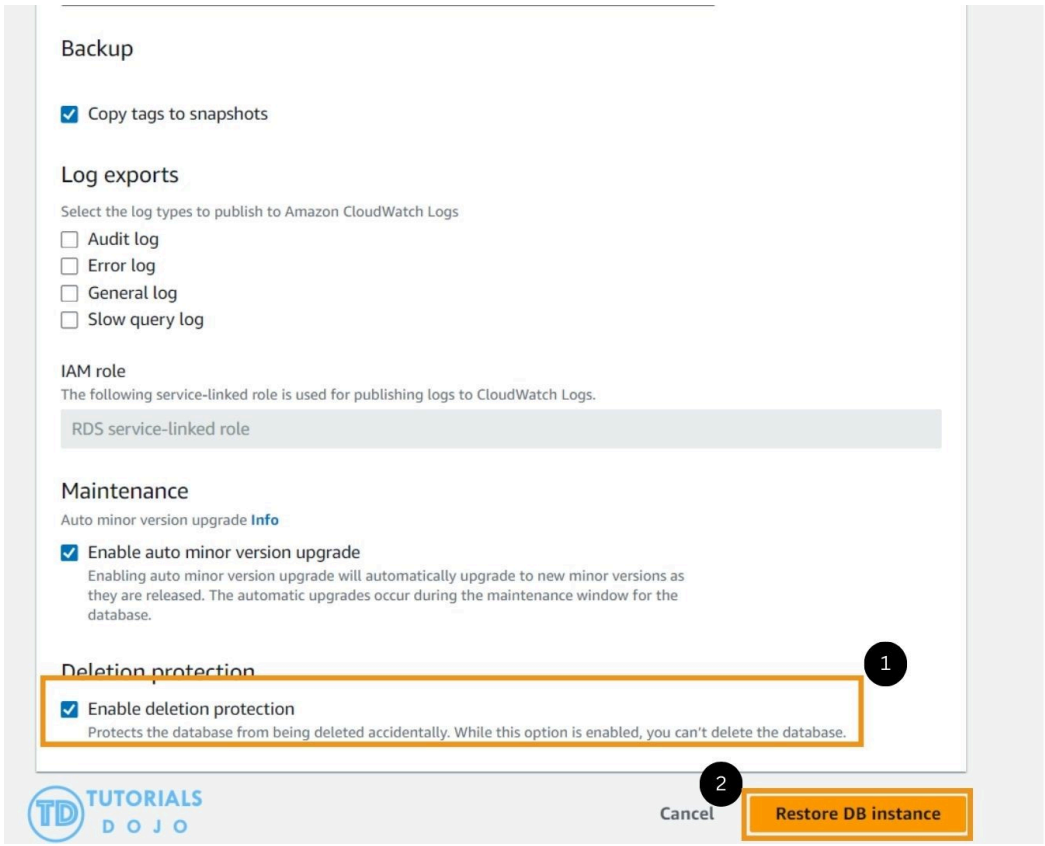

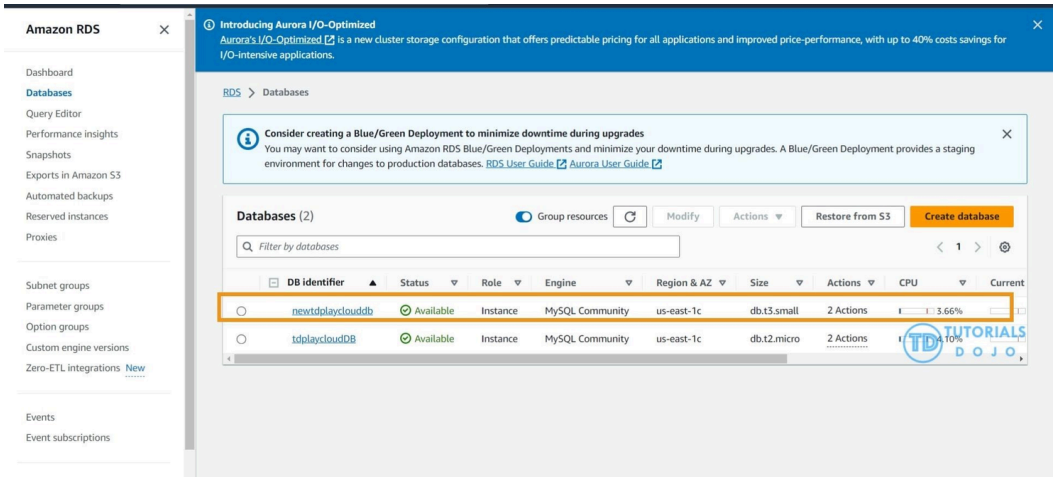

Creating and restoring RDS backups using snapshots is crucial for data protection and disaster recovery in AWS. Manual snapshots are user-initiated and customizable and provide control over retention, making them suitable for specific backup needs. In contrast, automated snapshots are system-generated, recurring backups with limited retention, offering a convenient daily routine for database protection. The choice between manual and automated snapshots depends on your specific backup and retention requirements, with many users often employing a combination of both for comprehensive data management and security.