

The Performance Computing of Oracle Database on AWS Cloud Storage

Arpitha M.S

Department of Computer Science
Yuvaraja's collage, Mysore
Karnataka - India

ABSTRACT

The Amazon Web Services (AWS) is a widely used cloud computing platform and it also more flexible, cost-effective and easy-to-use technique. In AWS, RDMS is used to store huge amount of data and widely deployed within the Amazon cloud. RDBMS is also helping us to understand how Oracle Database is deploying in an AWS.

The main goal is to explain how to run Oracle Database on both Amazon RDS and Amazon EC2, and to give an understanding of the advantages of each approach. We review in detail how to provision and monitor Oracle database, and how to manage scalability, performance, backup and recovery, high availability and security in both Amazon RDS and Amazon EC2. Here also describe how can set up a Disaster Recovery solution between an on-premise Oracle environment and AWS, and how you can perform a migration of your existing Oracle database to AWS.

Keywords:- Cloud Storage, oracle Database, AWS, Amazon EC2, Amazon RDS.

I. INTRODUCTION

Now a day's cloud computing is the most important platform is to make unique and indispensable changes in both technology and business. It has created an ongoing IT infrastructure challenge for many traditional technology executives. In the past few years, the typical business application architecture has evolved from a desktop-centric installation, then to client/server solutions, and now to mainly focusing on coupled web services and service-oriented architectures (SOA). Each evolutionary step has built on the previous one while adding new challenges, dimensions, and opportunities for IT departments and their business partners. Recently, virtualization has become a widely accepted way to reduce operating costs and increase the reliability of enterprise IT. Along with these technology changes, the speed of innovation and unprecedented acceleration in the introduction of new products has fundamentally changed the way markets work. Along with the wide acceptance of software as a service (SaaS) offerings, these changes have paved the way for the latest IT infrastructure challenge: cloud computing.

II. METHODOLOGY

There are two ways to run an Oracle database on AWS. One way is to use Amazon Relational Database Service (Amazon RDS). Another way is to deploy the database on Amazon Elastic Compute Cloud (EC2).

A. Oracle Database on Amazon RDS:

Amazon RDS is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. Amazon RDS offers automated Multi-AZ (Availability Zone) synchronous replication, allowing you to set up a highly available environment fully managed by AWS. Amazon RDS also automates installation, disk provisioning and management, patching, minor version upgrades, failed instance replacement, as well as backup and recovery of Oracle database. If we want Amazon to handle the day-to-day management of Oracle database, Amazon RDS is the preferred way. Amazon RDS enables us to focus on higher-level tasks, such as schema optimization, query tuning, and application development.

B. Oracle Database on Amazon EC2:

Amazon EC2 is a web service that provides resizable computing capacity in the cloud. It helps to deploy Oracle database on Amazon EC2, it has full control over the operating system, database installation, and configuration. It also have full control over system database administration, including backups and recovery, patching of the operating system and the database, tuning of the operating system and database parameters, security management, and configuration of high availability or replication. In terms of database administration, using Amazon EC2 is similar to running an Oracle database in-house. With Amazon EC2, it can quickly provision and configure instances and storage and also scale environment vertically or horizontally, and it can provision

databases in AWS regions across the world to provide low latency to end-users worldwide. Running the own relational database on Amazon EC2 is the ideal scenario if you require a maximum level of control and configurability.

C. Oracle Database instance on Amazon RDS:

The Amazon RDS allow us to starts an oracle Database instance. This can be done by using management console or just choose the Oracle engine version want to run, and we take care of choosing and configuring the underlying software. Using the AWS Management Console, select the RDS tab, click Launch a DB instance and then select the Oracle edition that you wish to use. Figure 1 shows an example of the Launch DB Instance Wizard.

Then configure the instance by specifying a few parameters, such as the size of the license model, the DB engine version, DB instance class, multi-AZ deployment, allocated storage, storage type, DB instance identifier, master user name, master user password and the confirm password. It can also accomplish all of this by using the command-line interface (CLI) or by invoking the application programming interface (API) from several programming languages including Java, Node.js, PHP, Python, Ruby and .NET.

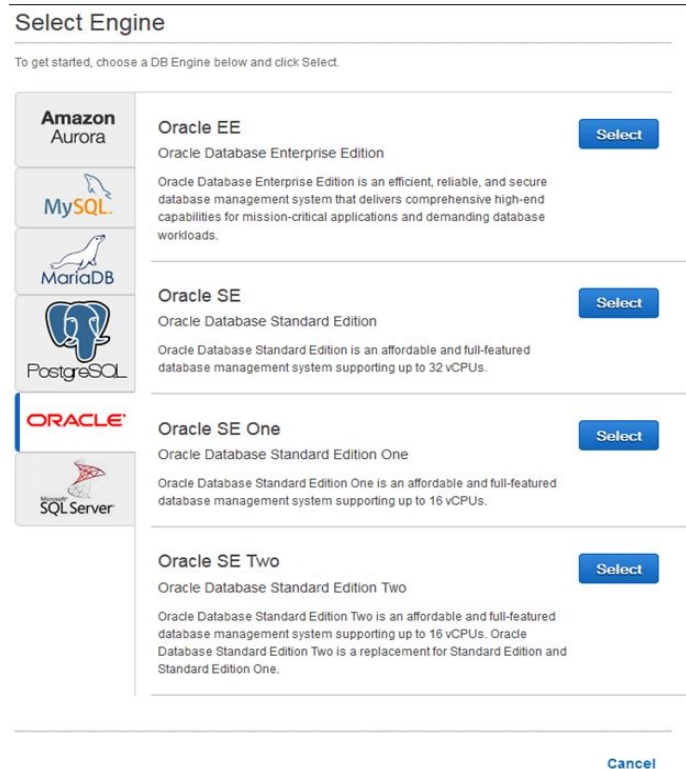


Fig. 1: AWS management console launch DB instance wizard.

D. Oracle Database instance on Amazon EC2:

To start an Oracle Database instance in Amazon EC2, first have to do is to choose an Amazon Machine Image (AMI). An

AMI is a special type of pre-configured operating system and virtual application software, which is used to create a virtual machine in Amazon EC2. The main thing is we can start an Amazon EC2 instance using the AWS Management Console, the CLI or the API.

III. PERFORMANCE MANAGEMENT

The performance of a relational database instance on AWS depends on many factors, the configuration of the database software, the application workload, including the Amazon RDS or Amazon EC2 instance type, the configuration of the database software, the application workload, and for Oracle databases running on Amazon EC2 instances, the storage configuration. The following sections describe various options that are available to show the performance of the AWS infrastructure on which Oracle database is running.

A. Instance sizing.

The performance of a database increasing to requires an understanding of which of the server’s resources is the performance constraint. If the database performance is limited by CPU, memory or network throughput, it can scale up the memory, compute, and network performance by choosing a larger instance type.

For many customers, increasing the performance of a single DB instance is the easiest way to increase the performance of their application overall. One way to achieve this is with vertical scalability. The scaling up of instance size is to address the hardware performance requirements of the database. In the Amazon RDS and Amazon EC2 environments, vertical scalability is very easy.

1) Instance Sizing in Amazon RDS: The choosing of Amazon RDS can that fits in specified workload. Amazon RDS supports several database instance classes. It has range from the very small Micro to the High-Memory Quadruple Extra Large, which features eight virtual cores, 68GB of memory and a high I/O capacity. The Amazon RDS instance classes are essentially identical, in terms of CPU, memory and I/O capacity, to the Micro, Standard and High-Memory Amazon EC2 instance types.

If you want to change the size of instance, here can modify the Amazon RDS instance, using the AWS Management Console, the CLI or the API and can also modify the Database Instance Class. This will cause a restart of database instance. Instance Sizing in Amazon EC2:

Amazon EC2 instances are grouped into eight families: Standard (first and second generation), Micro, High-Memory, High-CPU, Cluster Compute, Cluster GPU, High I/O, High Storage, and High Memory Cluster.

B. Disk I/O Management in Amazon RDS

Amazon RDS uses Amazon EBS for database and log storage. Depending on the size of storage requested, Amazon RDS automatically stripes across multiple EBS volumes to enhance IOPS performance. An IOPS means I/O operation per second which helps to scale the database storage performance.

1) Increasing Amazon RDS Storage Size: Amazon RDS makes it easy to scale up the storage by using the AWS Management Console, the CLI or the ModifyDBInstance API. The main reason for increasing the Amazon RDS storage size is to accommodate the database growth, but this can also be done to increase I/O.

2) Increasing Amazon RDS Instance Size: Increasing the size of the Amazon RDS instance may provide better performance because larger instances offer the following:

- More memory which allows for more data to be cached.
- More network bandwidth which can help provide better EBS performance.

3) Amazon RDS Provisioned IOPS: An Amazon RDS Provisioned I/O operation per second is a storage option that gives control over the database storage performance by allowing specifying IOPS rate at database creation time. This feature is designed to deliver fast, predictable, and consistent I/O performance. When it creates new Oracle database instances using the AWS Management Console or the Amazon RDS APIs, It can provision from 1,000 IOPS to 30,000 IOPS with corresponding storage ranging from 100GB to 3TB.

C. Disk I/O Management in Amazon EC2

The Amazon EC2 has Disk I/O management system which offers disk storage options that will choose for Oracle database.

1) Instance Storage: An Amazon EC2 instance comes with a certain amount of “local” storage. The data has been saved on an instance will not be available after that instance is terminated by the user, or if the specified hardware fails, which would cause an instance restart to happen on a different server.

2)

3) EBS Volumes: AWS offers a storage service called Amazon Elastic Block Store (Amazon EBS), which provides persistent block-level storage volumes. Amazon EBS volumes are off-instance storage that persists independently from the life of an instance. Amazon EBS volume data is uses the multiple servers in an Availability Zone to prevent the

loss of data from the failure of any single component. It is easy to back them up to Amazon Simple Storage Service (Amazon S3) using snapshots. These attributes make EBS volumes suitable for log files, data files and for the flash recovery area. The maximum size of an EBS volume is 1TB, but it can address larger database sizes by removing the data across multiple volumes. There are two types of EBS volumes, standard EBS volumes and Provisioned IOPS volumes, as described in the following sections.

➤ Standard EBS Volumes

Standard EBS volumes provide about 100 IOPS on average, with the ability to burst to hundreds of IOPS on a best-effort basis. Standard EBS volumes are great for applications with moderate or burst I/O requirements, as well as for boot volumes. Standard EBS volumes tend to perform better for sequential reads and writes.

➤ Provisioned IOPS Volumes

Provisioned IOPS volumes are designed to deliver predictable and consistent high performance for I/O intensive workloads such as databases. With Provisioned IOPS volumes, it can specify an IOPS rate when creating a volume, and then Amazon EBS provisions that rate for the lifetime of the volume.

D. Caching

Whether using Oracle on Amazon EC2 or Amazon RDS, Oracle users confronted with heavy workloads should look into reducing this load by caching data so that the web and application servers do not have to repeatedly access the database.

E. Database Replicas

A technique to provide higher performance is to spread the database query load across multiple instances. This technique is often referred to as scaling out or horizontal scalability. It also provides virtual storage space to store large amount of data.

1) Creating Read Replicas in Amazon RDS: Amazon RDS currently does not support read-replicas for Oracle. To increase database throughput, it needs to scale vertically by using larger instance types. Data could be shared based on real-world criteria or it could be spread across multiple shards using a hashing algorithm.

2)

3) Creating Read Replicas in Amazon EC2: Oracle Active Data Guard is an Oracle Database add-on, which allows setting up standby databases that can be open for read-

only requests, while continuing to archive transactions from the primary database. The standby databases can be used as read replicas of primary database. The replication between the primary and the standby databases can be configured to be synchronous. This allows to scale the database layer horizontally by adding read replicas and to offload read-only queries from the primary database.

III. CONCLUSION

AWS provides two deployment platforms to deploy Oracle databases: Amazon EC2 and Amazon RDS. It will be more helps to reduce operating cost and uploading cost by using fast network. It also provides overview of how to manage performance, high availability, monitoring and security in both environments.

REFERENCES

- [1] <http://aws.amazon.com/ec2/instance-types/>.
- [2] <http://aws.amazon.com/rds/oracle/>.
- [3] Jean-pierre Lw Goaller, Carlos Conde, and shakila langa, "RDBMS in the cloud: Oracle Database on AWS", Amazon Web Service ,October 2013
- [4] Jinesh varia, Sajee Mathew."Overview of Amazon web services".Amazon web services,January 2014