The topics deliberated during the FDP were:

- Steps to create cloud service on amazon
- Concepts of big data, hadoop
- Auto scaling
- Adding of volumes to server
- Hands-on practices – GCC lab

Steps to create cloud service on Amazon

Amazon Elastic Compute Cloud (EC2) is the Amazon Web Service used to create and run virtual machines in the cloud (we call these virtual machines 'instances'). The step-by-step guide to successfully launch a virtual machine on Amazon EC2 are.

- step 1: Launch an Amazon ec2 instance
- step 2: Configure your instance
- step 3: Connect to your instance
- step 4: Add storage
- step 5: Add tag
- Step 6: Configure security group
- Step 7: Review and launch

Concepts of big data, hadoop

Hadoop is an open-source software framework for storing data and running applications on clusters of commodity hardware. It provides massive storage for any kind of data, enormous processing power and the ability to handle virtually limitless concurrent tasks or jobs.

Why is Hadoop important?

Ability to store and process huge amounts of any kind of data, quickly. With data volumes and varieties constantly increasing, especially from social media and the Internet of Things (IoT), that's a key consideration.

Computing power. Hadoop's distributed computing model processes big data fast. The more computing nodes you use, the more processing power you have.

Fault tolerance. Data and application processing are protected against hardware failure. If a node goes down, jobs are automatically redirected to other nodes to make sure the distributed computing does not fail. Multiple copies of all data are stored automatically.

Flexibility. Unlike traditional relational databases, you don't have to preprocess data before storing it. You can store as much data as you want and decide how to use it later. That includes unstructured data like text, images and videos.

Low cost. The open-source framework is free and uses commodity hardware to store large quantities of data.

Scalability. You can easily grow your system to handle more data simply by adding nodes. Little administration is required.

Auto Scaling

Within the overall umbrella of Amazon Web Services are dozens of individual technologies that are used together to provision, launch, monitor and manage scalable web applications. Setting up intelligent auto scaling (AS) on AWS requires several of them, including:

- **Amazon Machine Image (AMIs)** – snapshot templates defining a launchable EC2 server instance
- **Elastic Load Balancer (ELB)** – a virtual load balancer platform with configurable events
- **CloudWatch (CW)** – tools to monitor and check your EC2 instances
- **Command line tools** – simple Java programs that call out to the AWS API using your credentials. Unfortunately, Amazon hasn't added all of the autoscaling configurations to the online AWS Console yet, so until further notice, you'll have to use a few command-line scripts to finish out the autoscaling configuration. Download the Auto Scaling command line tool from the AWS developer portal and run them on the command line to configure your autoscaling setup.

Adding of volumes to server

Amazon Elastic Block Store (Amazon EBS) provides block level storage volumes for use with EC2 instances. EBS volumes are highly available and reliable storage volumes that can be attached to any running instance that is in the same Availability Zone. EBS volumes that are attached to an EC2 instance are exposed as storage volumes that persist independently from the life of the instance.

Amazon EBS is recommended when data must be quickly accessible and requires long-term persistence. EBS volumes are particularly well-suited for use as the primary storage for file systems, databases, or for any applications that require fine granular updates and access to raw, unformatted, block-level storage. Amazon EBS is well suited to both database-style applications that rely on random reads and writes, and to throughput-intensive applications that perform long, continuous reads and writes.

For simplified data encryption, we launch EBS volumes as encrypted volumes. Amazon EBS encryption offers you a simple encryption solution for your EBS volumes without the need for you to build, manage, and secure your own key management infrastructure. When we create an encrypted EBS volume and attach it to a supported instance type, data stored at rest on the volume, disk I/O, and snapshots created from the volume are all encrypted. The encryption occurs on the servers that hosts EC2 instances, providing encryption of data-in-transit from EC2 instances to EBS storage.