

Mod-07 Lec-43 Cloud Service Providers

Balwinder Sodhi
Department of Computer Science and Engineering
IIT Ropar

Hello and welcome back. In this lecture, I will introduce you to some popular cloud service providers. I will talk about Google App Engine, which is a leading provider of cloud services for application development and I will also talk about Amazon Web Services, which is one of the earlier players in cloud computing landscape.

(Refer Slide Time 00:23)



Cloud Computing

SERVICE PROVIDERS



(Refer Slide Time 00:25)

Commercial providers

- Amazon Web Services (AWS)
 - IaaS, PaaS
- Microsoft Azure
 - IaaS, PaaS
- Google
 - IaaS, PaaS, SaaS
- Several other providers: RackSpace, IBM, Yahoo etc.



So let's start by looking at the commercial providers of various cloud services. We have Amazon Web Services, Microsoft Azure and then Google provides several of cloud services. Most of these providers offer all the three basic variants of services that is Infrastructure as a Service, Platform as a Service and Software as a Service and of course there are several other players such as RackSpace, IBM, Yahoo etc., which provide different services, different cloud based services.

(Refer Slide Time 01:00)

Open source projects

- OpenNebula
 - Nimbus
 - Eucalyptus
 - OpenStack
 - CloudStack
 - AppScale
- Mostly used for setting up IaaS clouds



On the open source space there are several projects which allow you to set up an in-house cloud that is a private cloud where you may configure different services such as Infrastructure as a Service. So all of these projects, which I have highlighted here, like OpenNebula, Nimbus, Eucalyptus, OpenStack etc., they're typically used for setting up IaaS clouds.

AppScale is an open source project, which provides necessary software infrastructure for you to set up a Google App Engine like PaaS platforms in-house.

(Refer Slide Time 01:40)

IaaS provider: AWS



- One of the earliest cloud services provider
 - Compute, storage, networking, applications etc.
- Has datacenters in several geographies
 - North America, Europe, Asia
- Offers excellent tooling support
 - For admins as well as developers

As I said Amazon Web Services is a Infrastructure as a Service provider and it is one of the earliest players in the cloud computing space and it has data centers across different geographies such as North America, Europe and Asia and it offers excellent APIs and tooling support for the cloud users so that they can make use of Amazon Web Services.

(Refer Slide Time 02:07)

PaaS/SaaS provider: Google



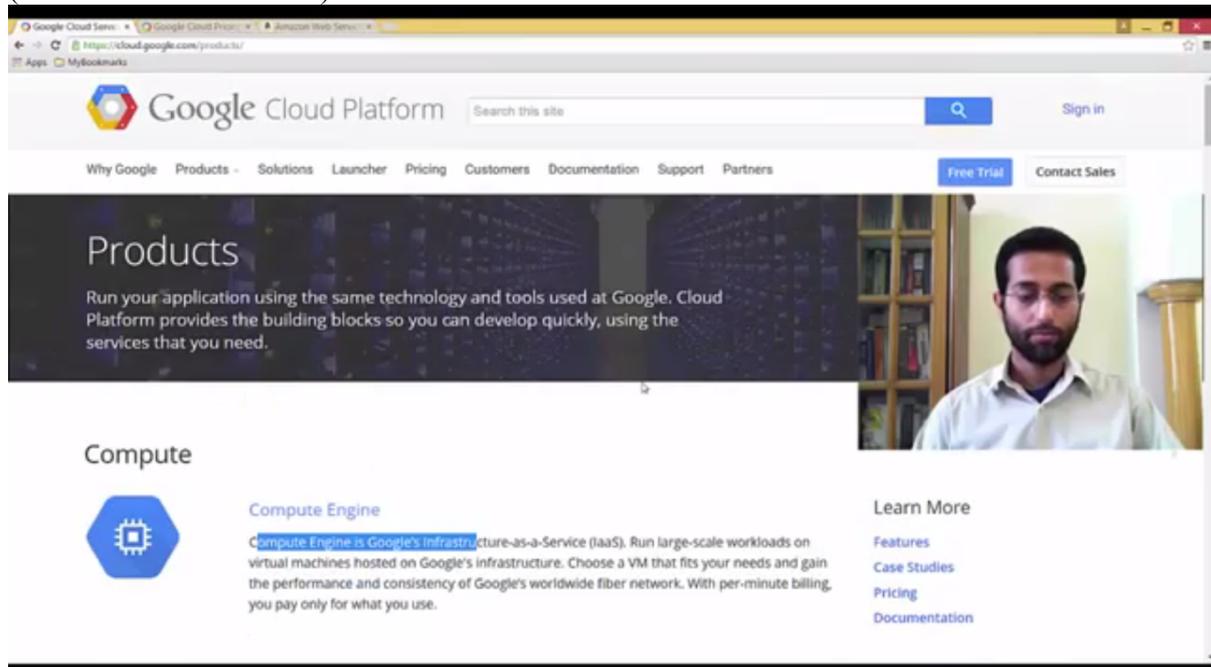
- Started with PaaS/SaaS offerings
 - Google App Engine (PaaS platform)
 - Develop web accessible applications
 - Deploy on Google managed platform
 - Google Apps (SaaS)
 - Enterprise productivity applications suite
 - Email, collaboration applications etc.

Google is another provider, which started with mainly PaaS and SaaS offerings. Google App Engine was their initial offering. Google App Engine is a PaaS platform, which allows the

developers to build web accessible applications and Google Apps basically is a suite of productivity applications, which different enterprises may subscribe to such as for email, collaboration etc.

Just to give you an idea about what different services are available from both Google and Amazon Web Services, I will be showing their respective portals.

(Refer Slide Time 02:47)

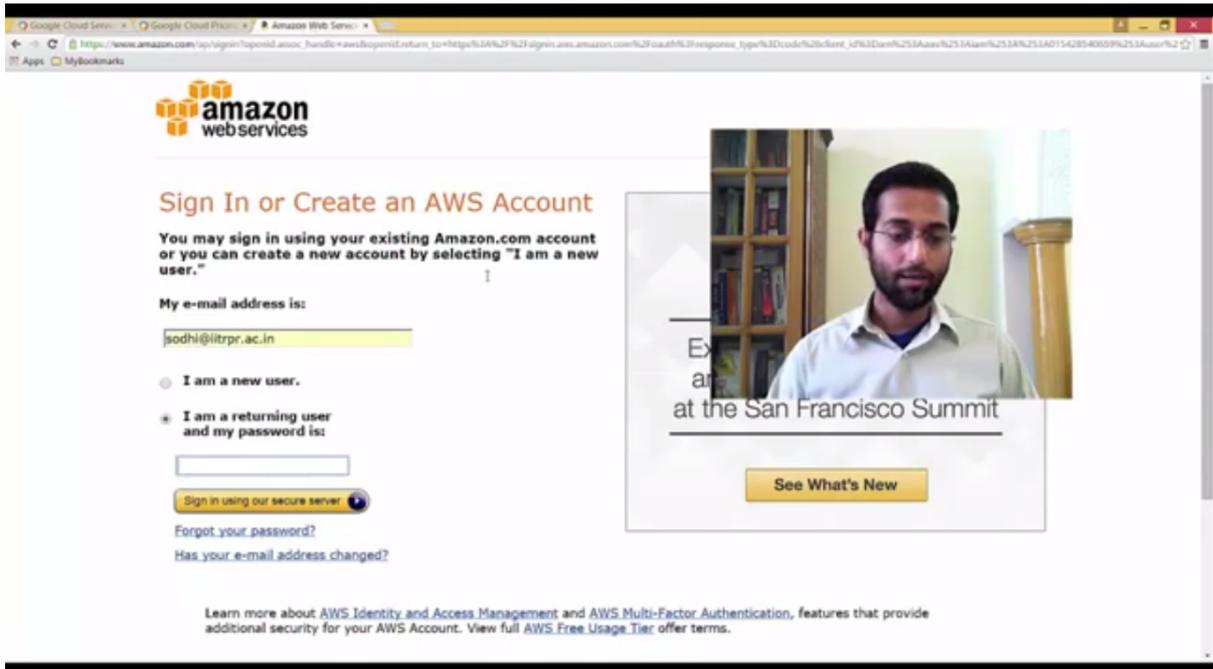


So Google provides a variety of services, cloud services. They provide the core infrastructure services such as Compute Engine. Then you have application development using App Engine. It's a PaaS platform.

Further down you have various types of storage services available ranging from SQL-based fully MySQL database to NoSQL instances using their cloud storage. Then you have various networking services such as Load Balancing, Interconnect through VPNs etc., and then DNS services.

Then further along they have Big Data related services available on their cloud infrastructure. Further they offer specialized services through the APIs such as Prediction API, Translate API etc. And then finally to allow managing your cloud subscriptions, cloud service subscriptions, they offer various other services for the management purposes. And then associated development tools are also provided. So they provide quite a comprehensive solution if you want to use different cloud services.

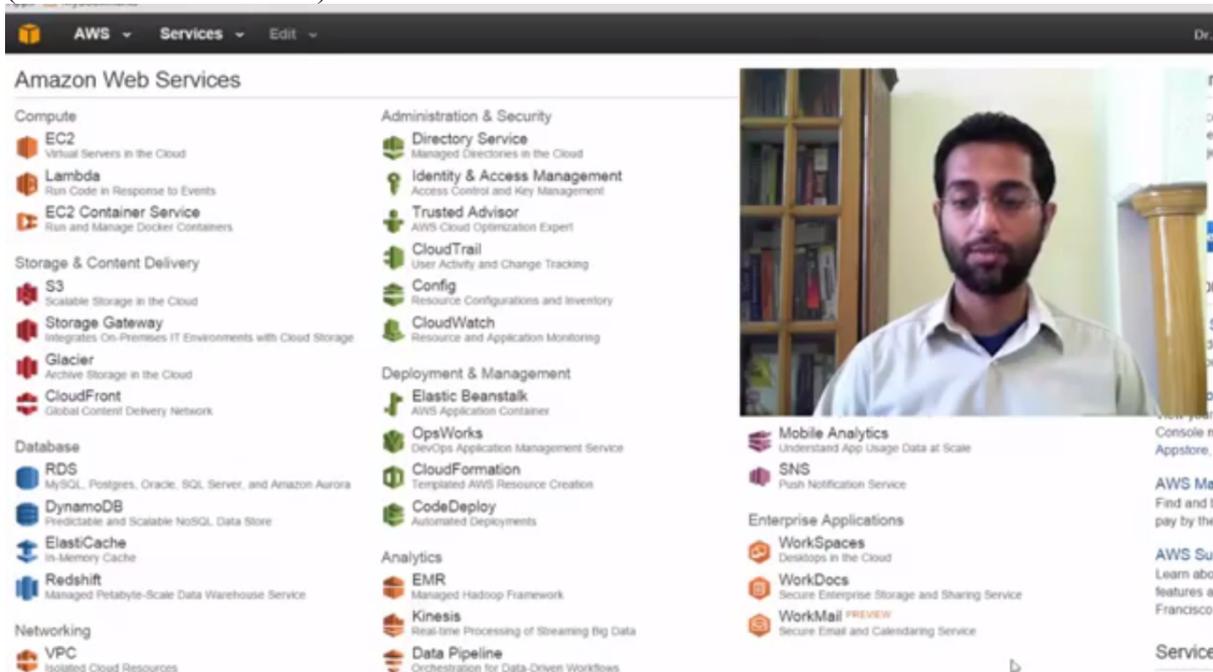
(Refer Slide Time 04:08)



Then we have Amazon Web Services. Amazon also provides a variety of different cloud services.

So I'm going to show you how a typical console, Amazon Web Services console looks like when you have subscribed to their cloud services. So I have got an account with them. So I will just try to log on to that account. I have enabled multi-factor authentication. That is what is prompting me to enter another code. So I have entered this code and it should log on.

(Refer Slide Time 04:49)

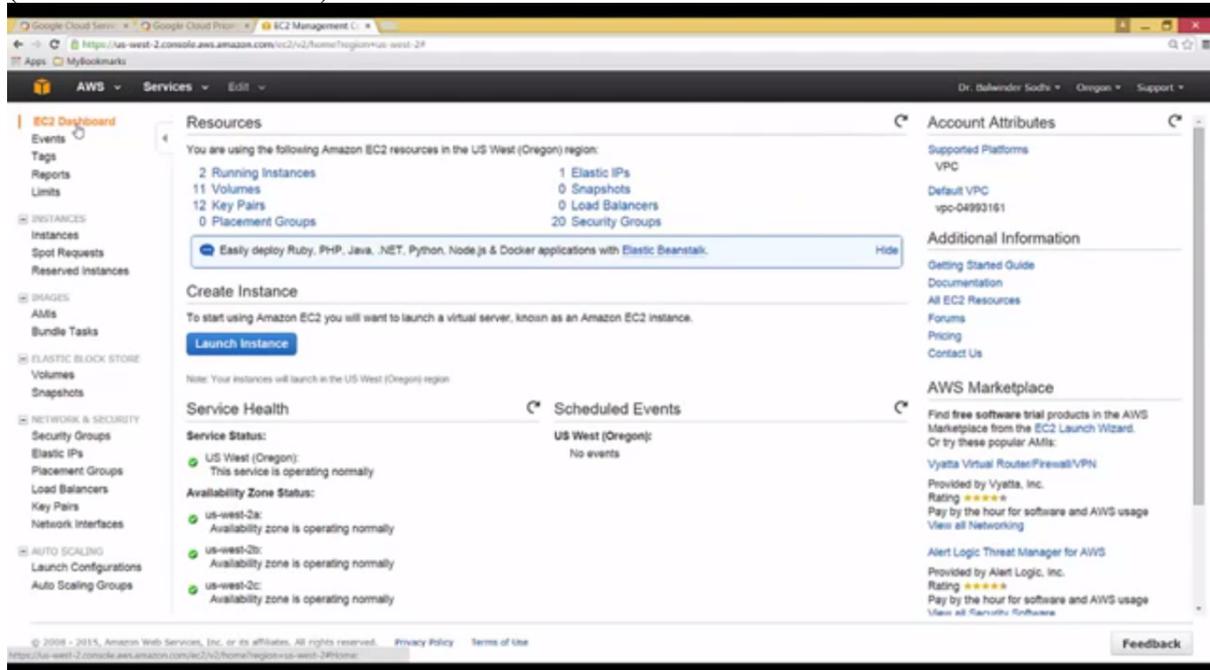


Okay. So after I successfully log on, authenticate with the Amazon Web Services console, I see all the services that is available to my account. So you see here Amazon provides Elastic

Compute. Then you have various storage services here and database services, networking and so on and so forth.

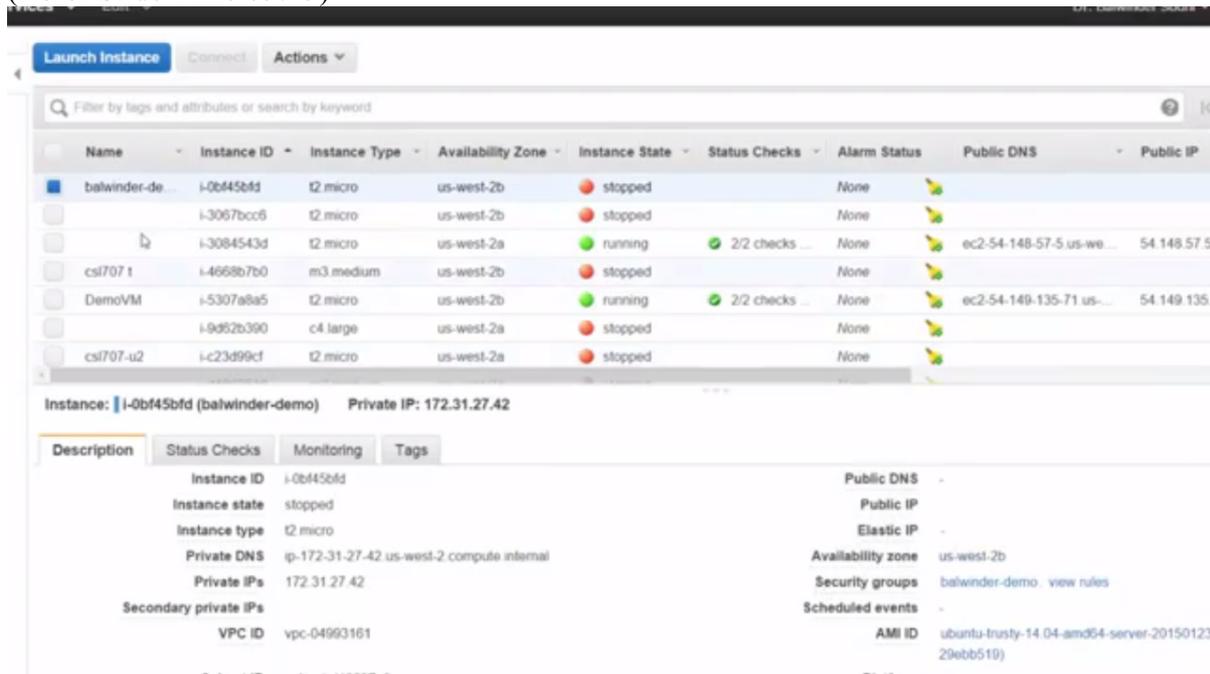
I will demonstrate launching of a virtual machine. A virtual machine can be launched through EC2 that is Elastic Compute cloud.

(Refer Slide Time 05:22)



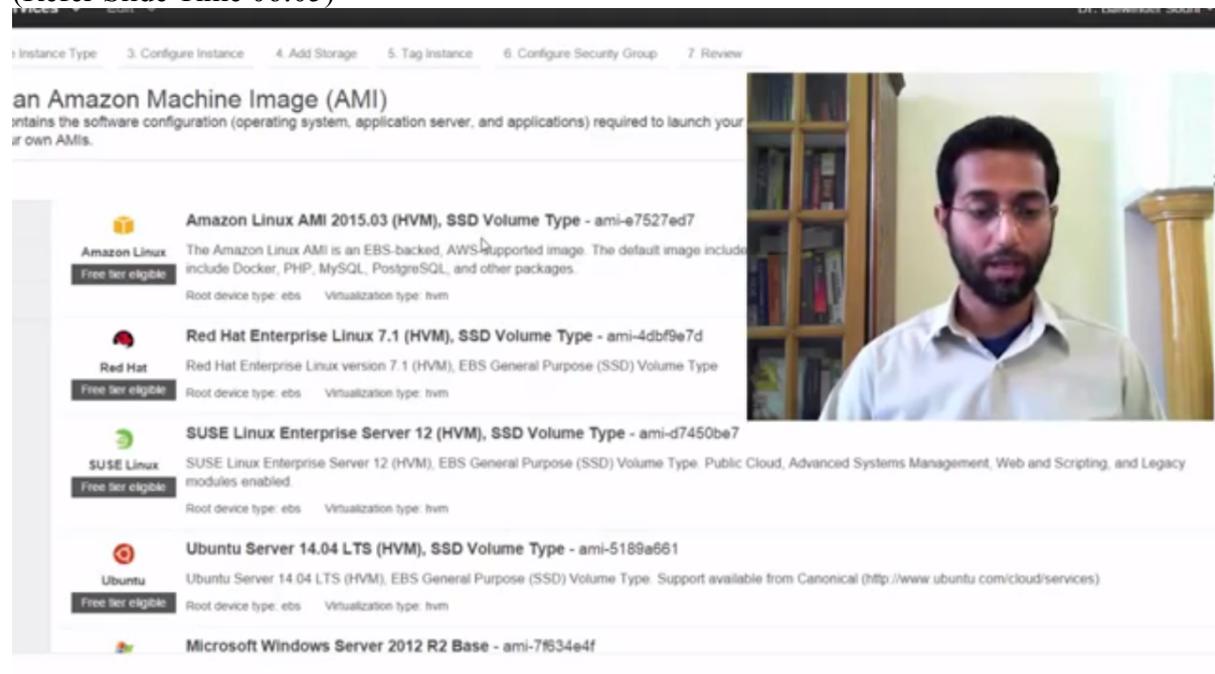
So this is after I go to EC2 dashboard, it shows the summary of all my resources. So I have two instances running. There are 11 volumes, disk volumes and 12 key pairs created and 1 elastic IP. So let's look at the existing instances. That is existing virtual machines that are running on my account.

(Refer Slide Time 05:45)



So this is how typically it looks like. So once you click on the name of an individual instance, you can see the details of that instance underneath here.

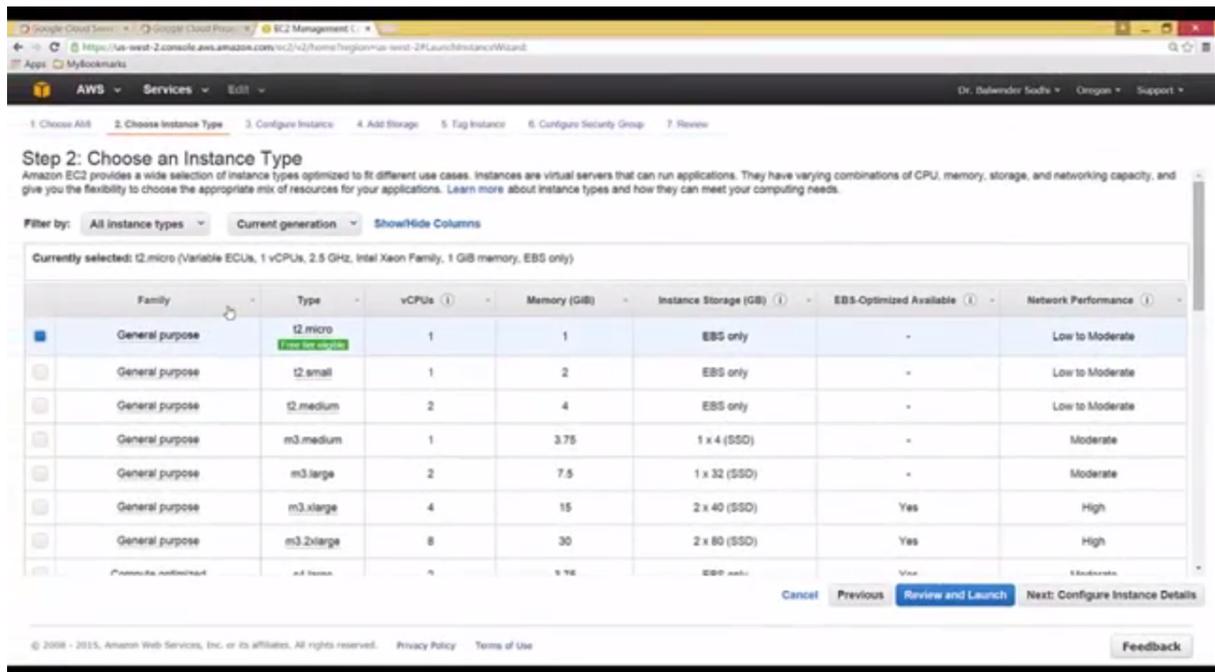
(Refer Slide Time 06:05)



So let's try to launch a new instance. So after I clicked on the launch new instance button, it takes me to multi -- it takes me through multiple steps right from choosing the Amazon Machine Image, AMI to choosing the instance type, then configuring the instance, adding some storage, putting any meta information such as tagging that machine and I'm putting some security setup on this. Finally, after I review all the settings, I can finally launch that particular virtual machine.

So let's pick some Amazon virtual machine image here. So there are a variety of them. So let's put a Ubuntu -- so let's try to launch a Ubuntu Server instance. I will select that.

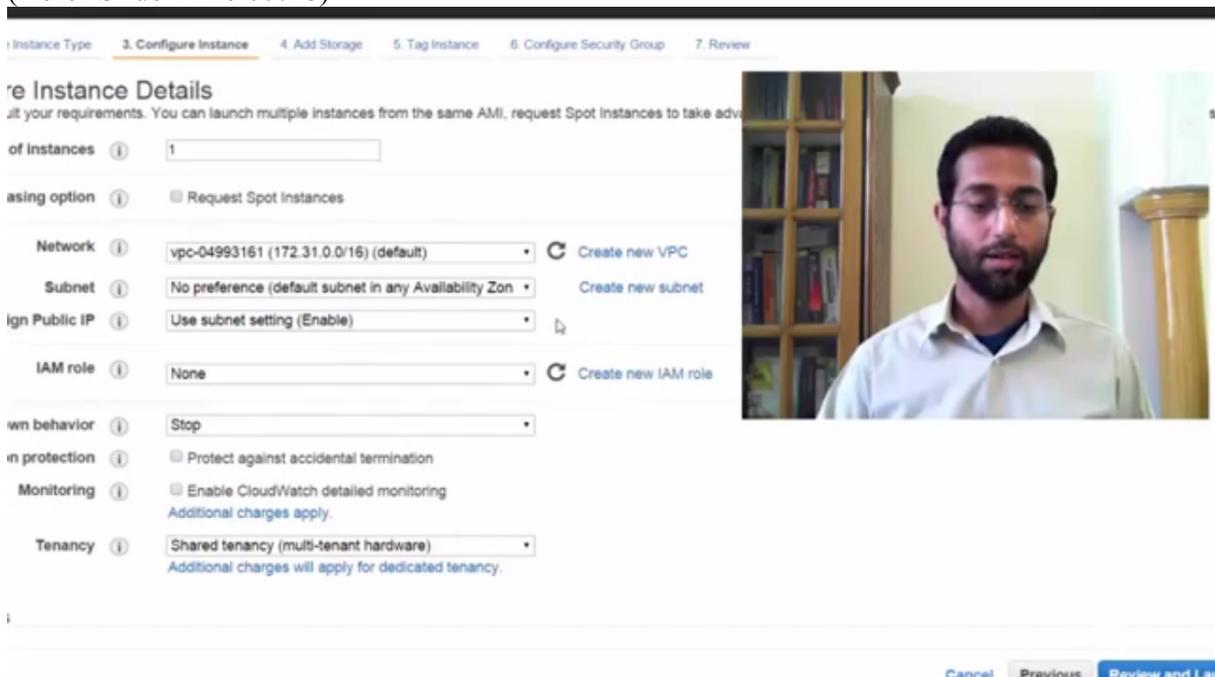
(Refer Slide Time 06:50)



So in this step it is asking me to choose the type of virtual hardware that my virtual machine should have. So here you see a variety of options such as micro, small, medium, large and so on and so forth. Each one of those differ in terms of how many virtual CPUs does it have, how many gigabytes of memory is attached to it, and what is the type of instance storage that is virtual machine's storage attached to that virtual machine and so on.

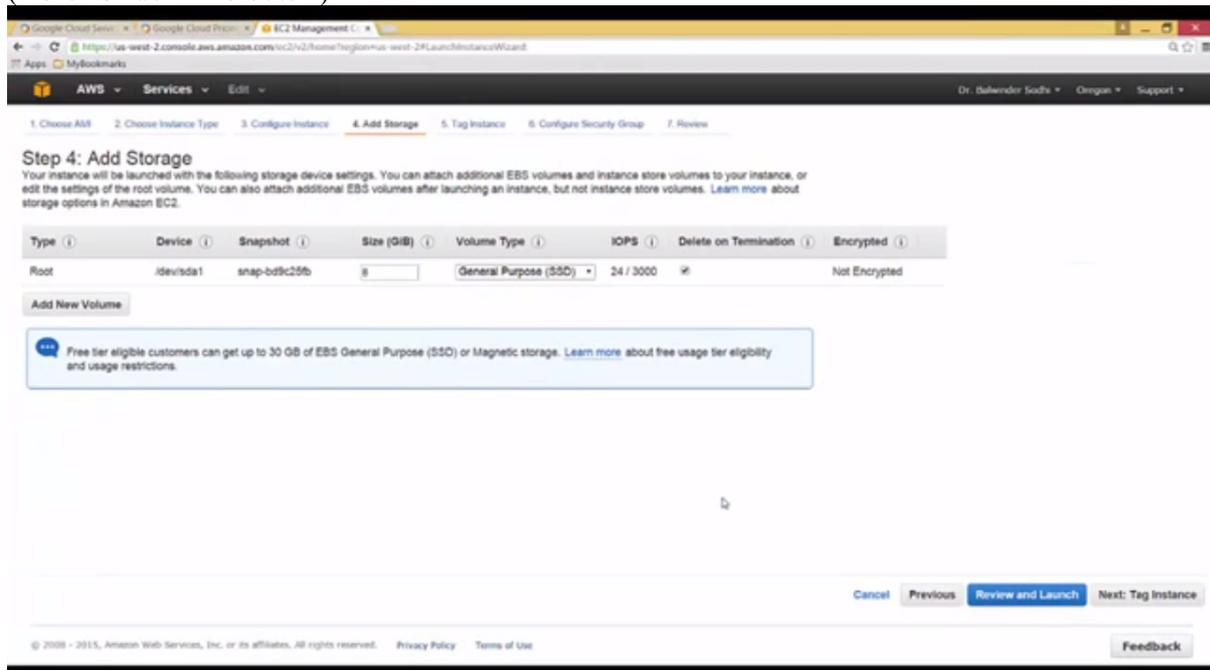
So let us go by the default. I will pick that micro instance. Let's go to the next step.

(Refer Slide Time 07:28)



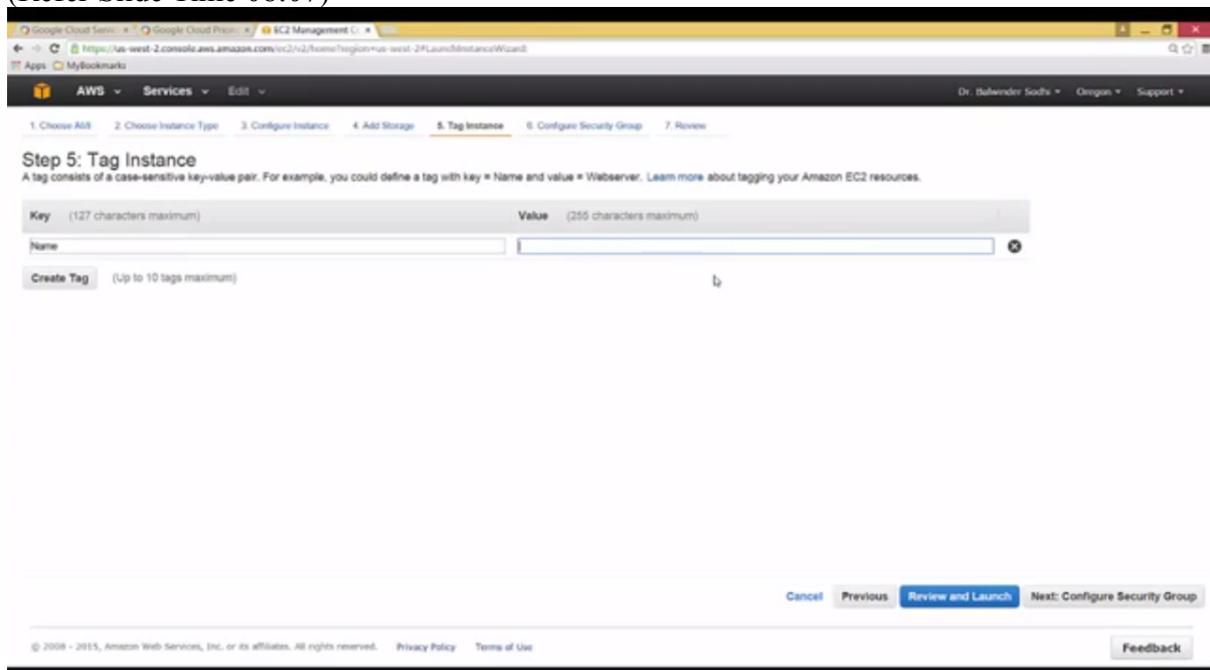
Here I can further configure, further specify various details such as number of instances I want to launch. I just want to launch one instance at this time and then there are various roles, security specific things, what should happen on the shutdown and so on and so forth.

(Refer Slide Time 07:52)



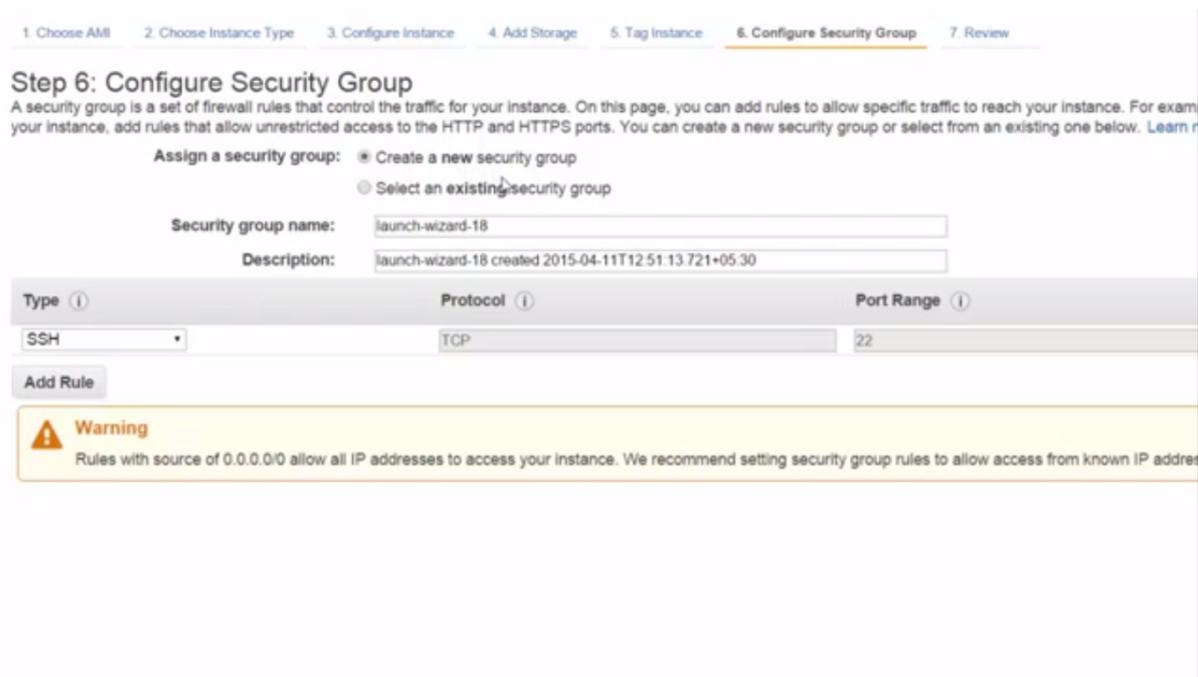
Then I go to the next step. I review the storage attached to it. So we just keep the default, which is 8 GB of virtual hard disk. It's the general-purpose SSD and put some meta information.

(Refer Slide Time 08:07)



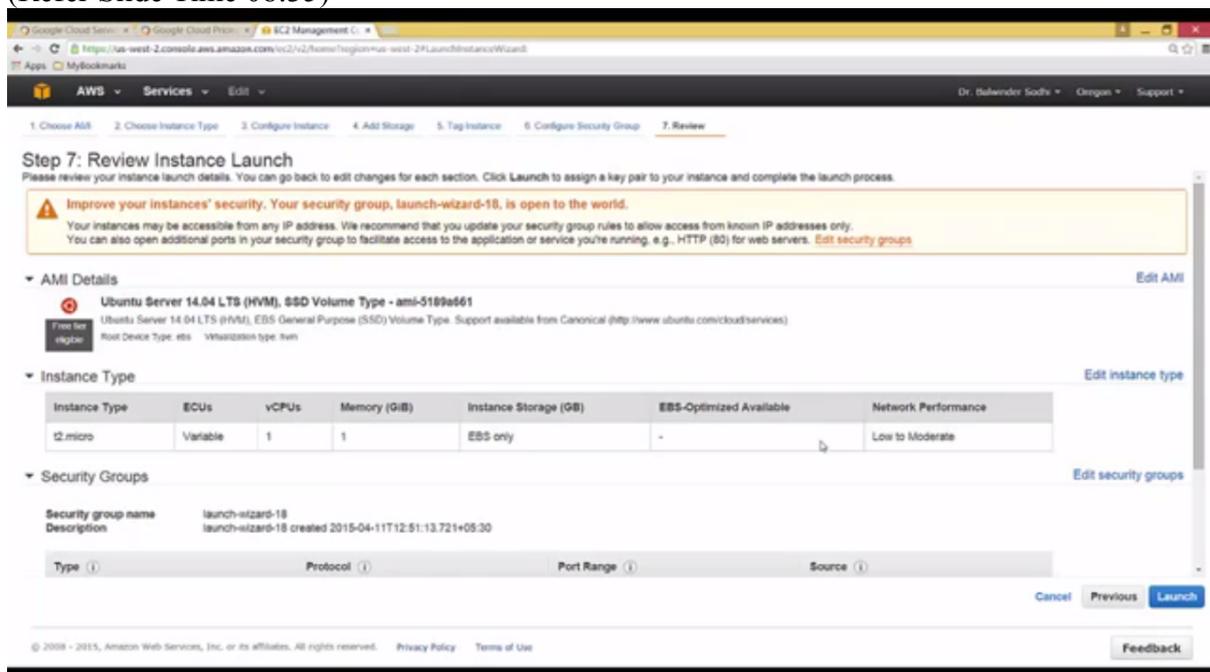
Let's say we put it as Cloud Demo. Go to the next step.

(Refer Slide Time 08:19)



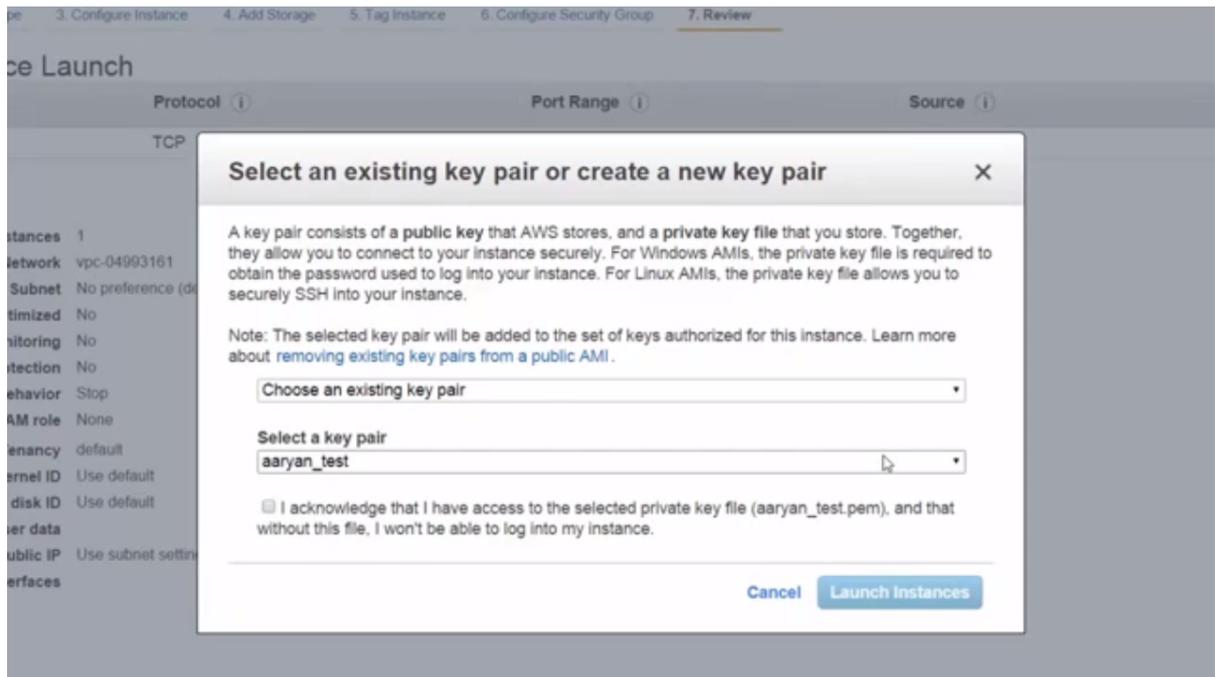
So here I can specify a new security group basically telling how your machine should be accessible over the network. So I can create a new security group or I can use an existing one. Let's just create a new one.

(Refer Slide Time 08:35)



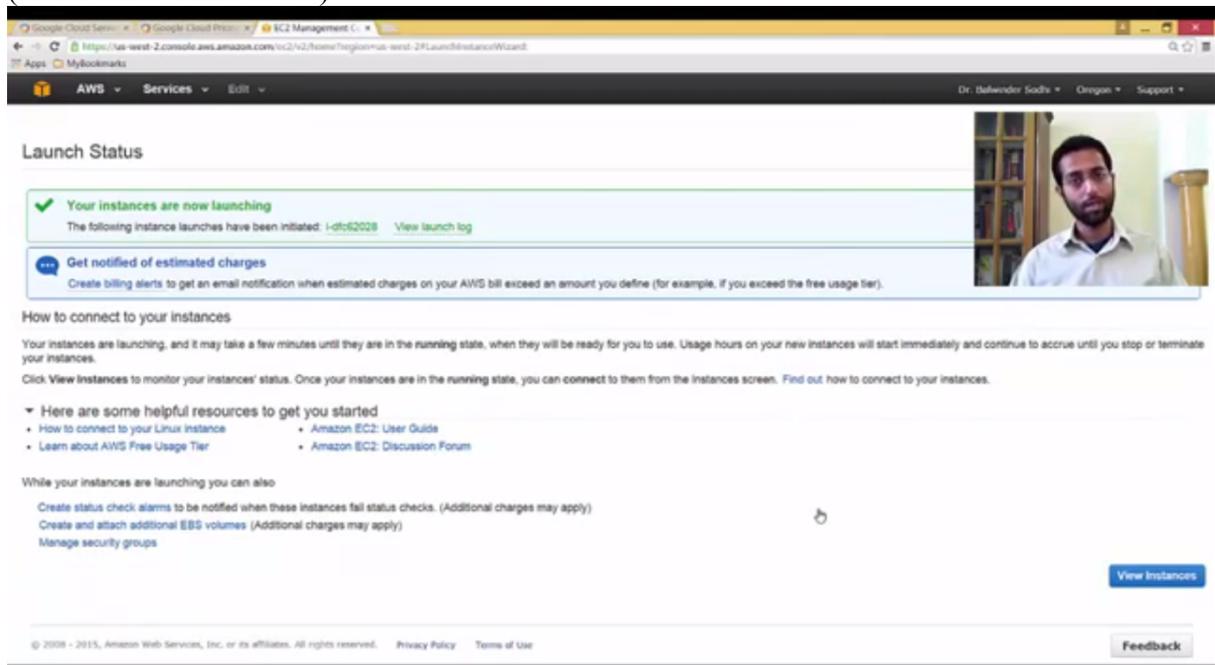
Now at this time it is showing me the overall configuration that I have selected in the previous steps. I can expand these details and look at more information as I have selected. So I will just click on launch.

(Refer Slide Time 08:53)



Now at this time, it will ask me whether I should get a new key pair, whether I should ask them to give me a new encryption key pair using which I can remotely log on to that virtual machine or I should use an existing one. So I already had created a key pair earlier. So I selected that and I put a check box here that I have -- I acknowledge that I have access to the selected key. If you don't have that then you will not be able to get into that virtual machine remotely.

(Refer Slide Time 09:35)



So let's just launch it. So it will perform various steps and after doing some basic checks, your virtual machine will be up and running.

(Refer Slide Time 09:46)

The screenshot shows the AWS Management Console interface. At the top, there are buttons for 'Launch Instance', 'Connect', and 'Actions'. Below is a search bar and a table of instances. The table has columns for Name, Instance ID, Instance Type, Availability Zone, Instance State, Status Checks, Alarm Status, and Public DNS. The instance 'balwinder-de...' (ID: i-0bf45bfd) is highlighted in blue and has a state of 'stopped'. Below the table, the details for instance 'i-0bf45bfd (balwinder-demo)' are shown, including its Private IP (172.31.27.42) and various configuration options like Instance ID, Instance state, Instance type, Private DNS, Private IPs, Secondary private IPs, VPC ID, Public DNS, Public IP, Elastic IP, Availability zone, Security groups, Scheduled events, and AMI ID.

So I can go into the view instances. So at that place I will be able to see what my machine is actually doing. So on this screen, I can see where my machine is, whether it has started or not.

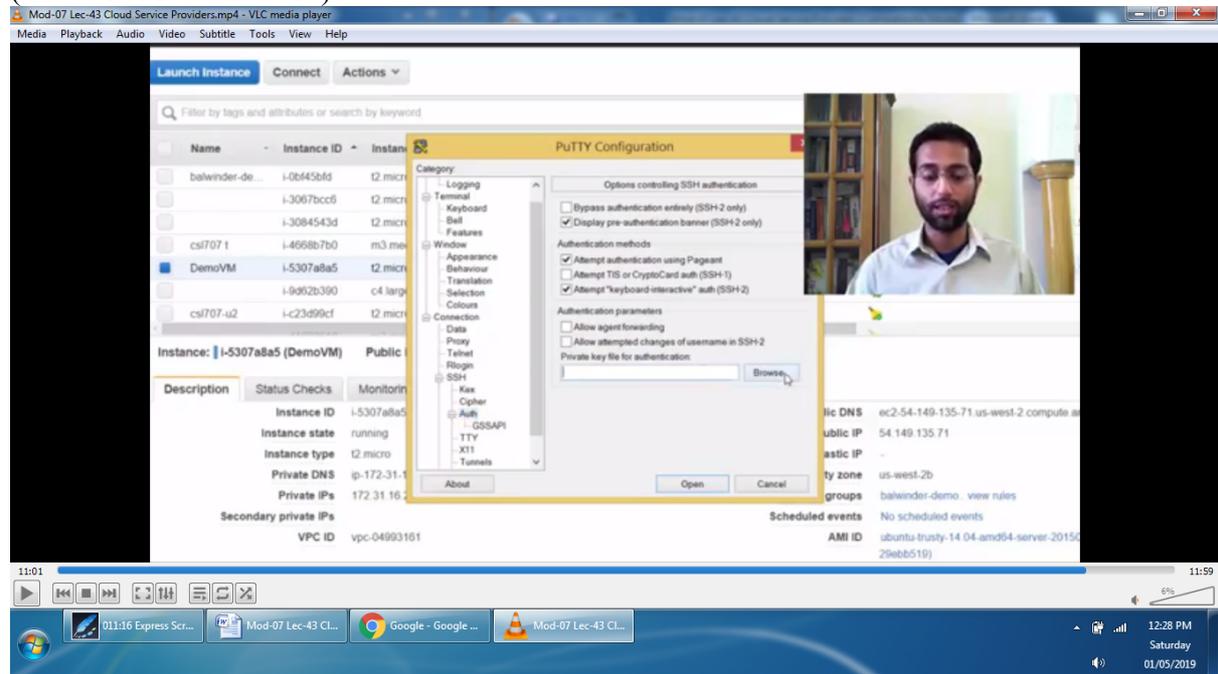
So this was the machine Cloud Demo and it is still in the initialization stage. Once it is in the running stage, we will see the appropriate status message here and the details about that machine are actually given over here. Okay. So all these status checks are complete in past. We can see the description of our machine. So it has assigned a public IP to this machine and also a public DNS is available using which we can SSH remotely log onto that machine. Let's try to do that. I'll copy this IP address and log on using a client tool.

(Refer Slide Time 10:50)

This screenshot shows the same AWS Management Console interface as the previous one, but with a 'PuTTY Configuration' dialog box overlaid. The dialog box has a 'Category' list on the left and 'Basic options for your PuTTY session' on the right. The 'Host Name (or IP address)' field is filled with '54.149.135.71' and the 'Port' field is '22'. The 'Connection type' is set to 'SSH'. The 'Close window on exit' options are 'Always', 'Never', and 'Only on clean exit', with 'Only on clean exit' selected. The 'Open' button is highlighted. In the background, the instance details for 'i-5307a8a5 (DemoVM)' are visible, showing its Public IP as '54.149.135.71'.

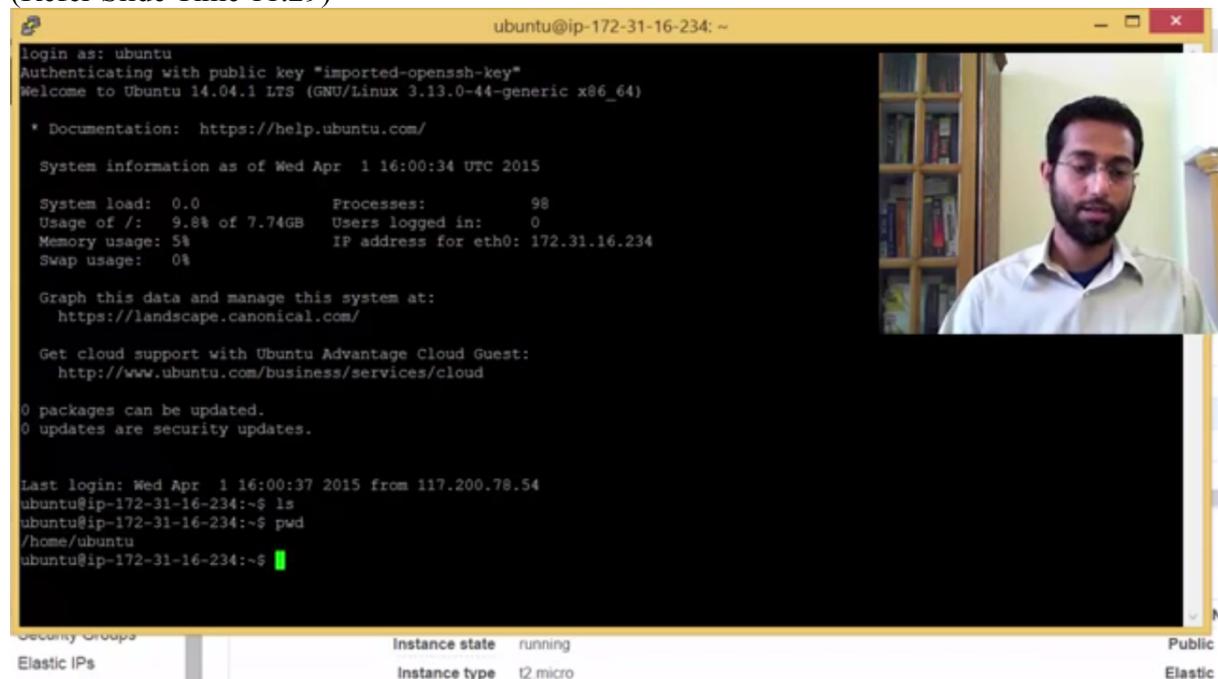
Put the IP address and under Auth, I have actually generated a key using -- I have already downloaded the key. So I will select that and all right.

(Refer Slide Time 11:01)



So we have logged on to the remote machine, which has been started in Amazon's data center. So I can list. There are no files. I can run some command.

(Refer Slide Time 11:29)



We have seen that launching a machine didn't take much time and we were quickly able to log on to that machine remotely.

(Refer Slide Time 11:39)



THANK YOU!



So what we have seen is different players that provide different services in cloud computing space. We have looked at Amazon Web Services and we have also looked at different services, which are provided by Google cloud platforms.

So that's it for this lecture. Thank you.