Cloud Computing
Prof. Soumya Kanti Ghosh
Department of Computer Science and Engineering
Indian Institute of Technology, Kharagpur

Lecture – 15

Open Stack

Hello, so we will continue our discussion on cloud computing. Today we will discuss and so a demo on a open source cloud which is open stack; one of the very popular open source cloud. So, we will initially we will have some brief overview of open stack then we will; so, a demo on open stack; that how way open stack can be configured VM can be provision. Primarily, we have been used these open stack for IAS type of cloud infrastructure as a service type of cloud; that we will demonstrate today.

So, if you see open stack one of the very popular open source cloud; which you can download and install in a particular hardware configuration even with couple of species with you can install your open stack and see the performance of IAS type of cloud.

(Refer Slide Time: 01:28)

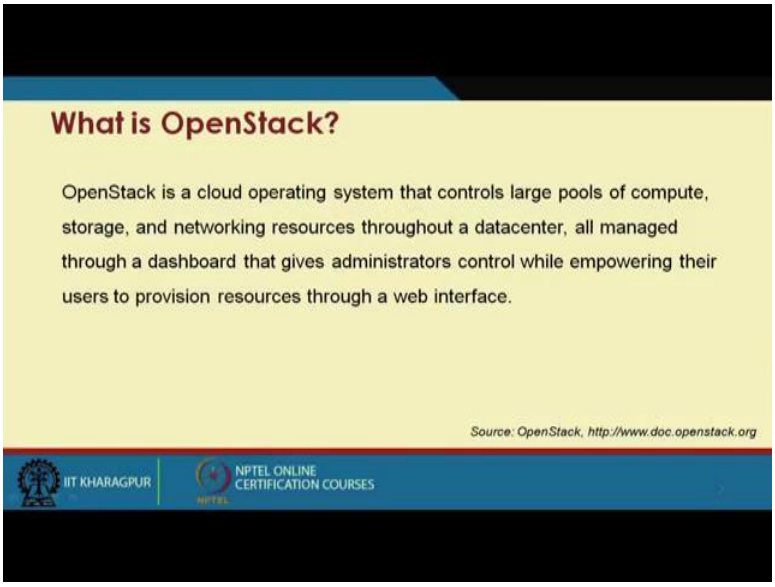

So, open stack is a cloud operating system that controls large pool of compute storage networking resources through a data centre; throughout a data centre, all manage through dashboard and gives administrator control and while empowering their users provisioning your resources through a wave interface.

So, what it says; its say you have a set of resources it is a layer above those base this bare metal resources and it gives the administrator to control these resources, and the user can basically provisions VMs out of it. So, it access a IAS type of cloud and as it is as we mentioned it is a open source. So, you can download and install and give it provision of the things. Incidentally in IIT Kharagpur, we have done experimental; we have made experimental cloud called Meghamala, which is based on open stack and which has been installed over blade server.

But never the less open stack you can install over PC's; here also we have tried with PC's etcetera for small-small student projects. So, it gives a feel that how a particular infrastructure as a cloud works and also how we can work on those things.

(Refer Slide Time: 02:57)

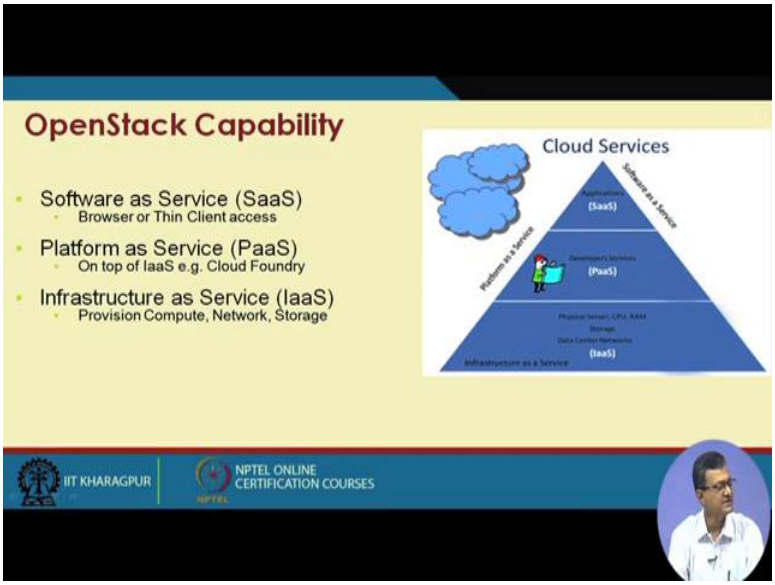

So, as for as the open stack capability is there; so, it has a capability of all the services though primarily we use more as the infrastructure as a service. So, at the infrastructure as a service provision come compute network storage at the pass level on top of IAS cloud foundry and over the as the SasS level that browser or thin client accesses. So, these are the whole pyramid of typical cloud services; which is accessible over internet. So, infrastructure as a service where you have physical serve as like CPU, RAM, storage, data center networks, etcetera; over that some developers service and over that some software as a service.

(Refer Slide Time: 03:44)

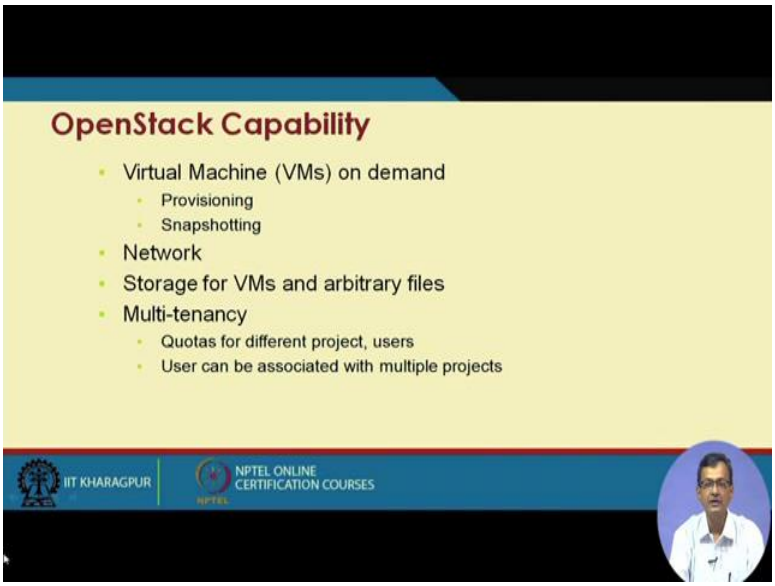

Now, if you look at the capability of open stack; so, primarily if you look at from point of view as infrastructure of service. So, it is a VMs, it can give VMs on demand; so, it is both provisioning and snapshotting is possible. So, it is virtual machine on demand, it has a provision for networking, it has storage for VMs and arbitrary files. So, you can have storage for virtual machines and other files and supports multi tenancy; quotas for different project users, user can be associated with multi projects and type of things.

So, in essence it gives you a full-fledged experience of a cloud and as it is a open source you are installing. So, you have the control both administrative and physical control over the whole thing, so the hardware is used and you are running on the things. So, it is a good thing or it is a very popular thing for individual user or group or lab to install and have a life experience and extremely useful for students with couple of PCs or even a couple of laptops to you can install open stack and see that how things are there.

As we mentioned earlier that there is a good amount of research going on; on resource management in cloud, resource management, power management and people are talking about green cloud and sort of things. So, this sort of lab scale implementation of open source clouds may help in having experience and experimentation of different type of parameters on the cloud. So, open stack is one of the very popular open source cloud.

(Refer Slide Time: 05:58)

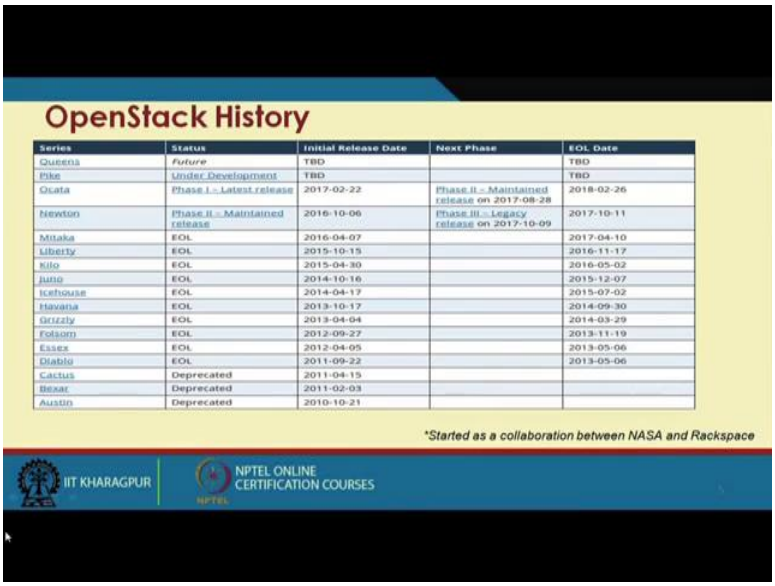

So, if you look at a history of open stack; it started with collaboration between NASA and Rackspace and over years, it has gone through different project mode and you can see that it started with around 2010 and we are having regular releases with over years.

So, if you see there is started with a project called Austin and go on going to that Newton and Pike and Queens which are to be declared.

(Refer Slide Time: 06:43)

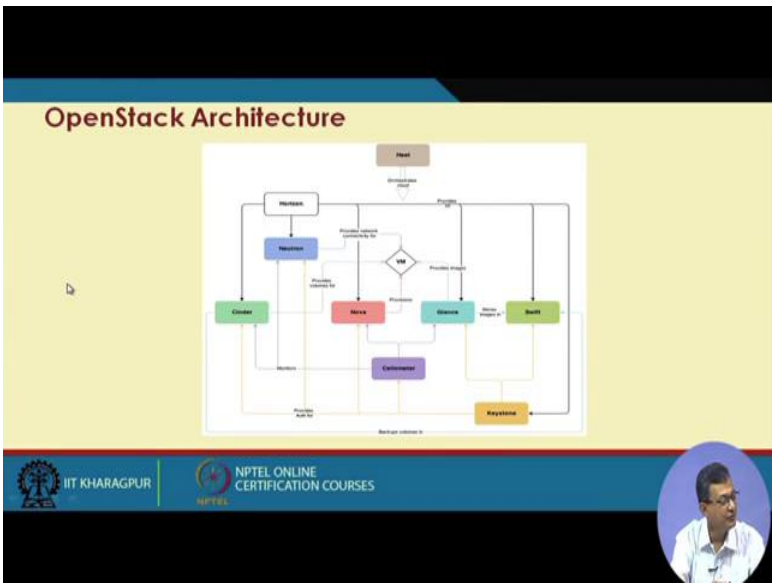

So, if we look at the overall architecture of open stack; there are a couple of components like one is that horizon which is primarily the dashboard; which invite the dashboard

project horizon. So, project newton which is the mostly look at the networking cinder which primary look at the block storage of open stack; nova is the compute, glance is the image services; the different type of different image services which can be hosted in the open stack is there.

Then we have other things like swift; which is object storage like we have cinder as block storage. So, object storage is swift; then ceilometer which is ceilometers, which is the telemetry services; keystone is the idea identity services. So, what we can see that this all different sort of services; which are required in a typical cloud are there in open stack.

So, though they developed in different project mode, but it comes as a bundle and we can have when we have installations, we have all these flavor into the things. So, it is good to have a feel of the cloud and you can have operational cloud using your own internal or in-house cloud using these type of open source. So, things it is less costly you have to pay for the infrastructure and if you have a one infrastructure as I mean already present or access infrastructure already present; which we can deploy it into the using open stack or any other open source cloud.

(Refer Slide Time: 08:44)

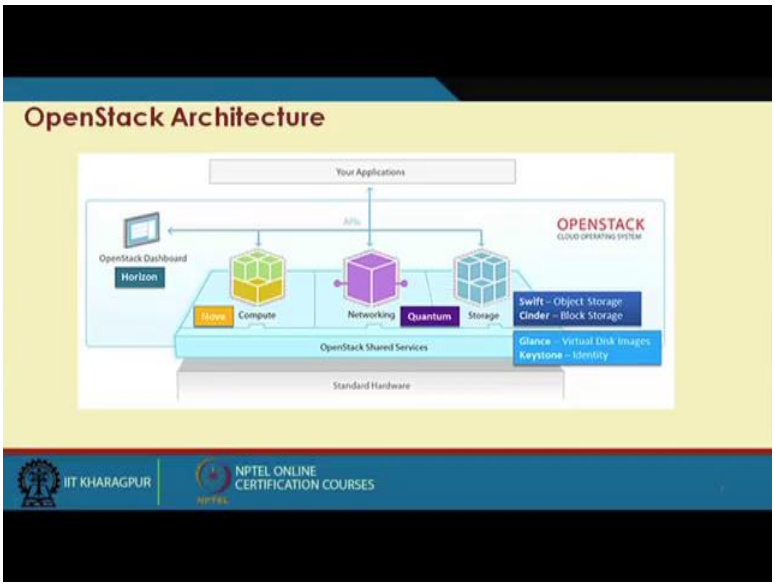

So, if you look at from the point of view; then you have the standard hardware at the backbone, then the open stack shear services and over that we have compute networking

storage services and so, this and then their open stack dashboard and the services we talked about.

So, the user applications come from the top and it goes on to this overall inform APIs and utilizes this different services to execute particular or realize particular job. I would like to mention here that we have taken most of these resources from again open stack site and other resources; so, these are the things which you can get in right that also. So, if you look at the major components as we are mentioning here; if you just go back like horizon, newton, cinder, nova glance, swift, ceilometer and keystone; so, these are the major component.

(Refer Slide Time: 09:54)

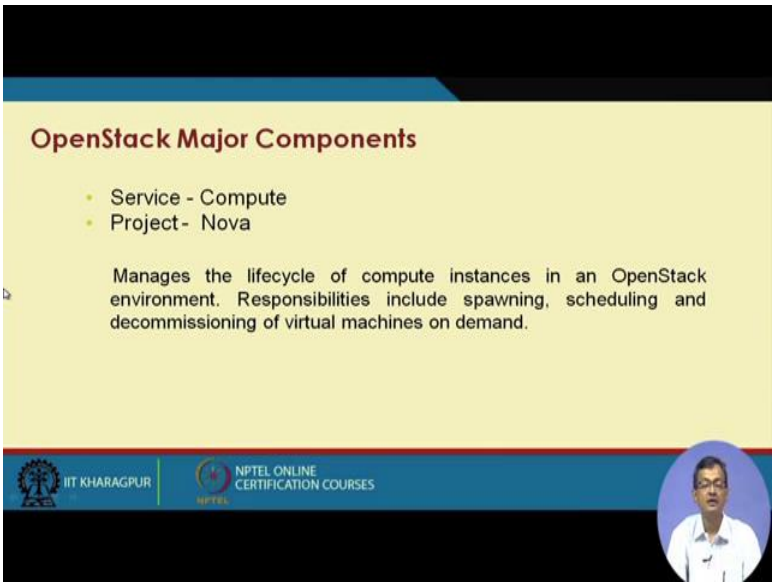

So, we will just have a brief overview of these components before we see a demo on those things. So, one of the critical components compute which the service is compute and the project is a nova project; so, manages the life cycle of compute instances in a open stack environment. It manages the life cycle of a compute instances in a open stack environment, responsibilities includes spawning, scheduling, decommissioning of virtual machine on demand.

So, it primarily work with this VM scheduling, VM spawning and decommissioning or releasing the virtual machine on all on demand. So, if you look at it manages the whole compute cycle, so if you see that there are different type of flavor for VMs with different configuration in typical cloud environment; here also we can have different flavor of

VMs. So, I will user request that through the dashboard; the manager or the controller of the cloud can allocate different VMs to the user based on the requirement.

So, it comes with a storage; the ephemeral storage and a persistent storage which can be associated with this VMs; so, these are possible using this compute or nova services.

(Refer Slide Time: 11:44)

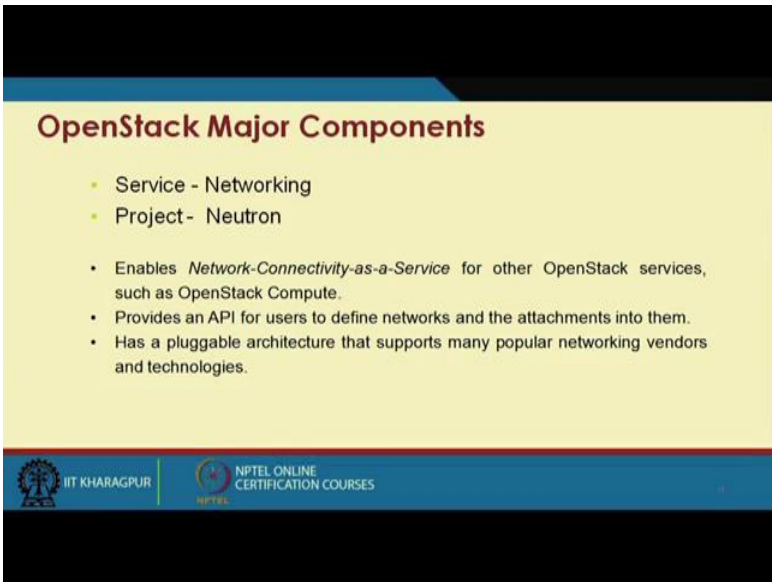

Another aspects is the networking, another major service is the networking; which is under the neutron thing, enables network connectivity as a service for other open stack service is such as open stack compute another things.

So, if you have other open stack surface like compute or nova storage services and so on. So, this neutron or the networking provides a networking as a service to this different component of open stack. So, provides an API for users to define network and attachment to them. So, it provides API to the users to define networks and how to attach the things. So, if you look at any cloud infrastructure; so, two type of networks are prominent there; one is the internal network, which is internal to the cloud and there is a external network to which is external to the cloud.

So, like if as you are talking about that we have a open source cloud in our institute that what we call it Meghamala; the experimental open source using open stack with a very only experimental cloud which have being used by faculty and research scholars for their

computing primarily or computing needs, it comes with different flavor; we will show that example and how we do.

So, it has a internal network for the cloud whereas the external network for that cloud is basically the IIT network. So, as it is a in house cloud; so, this is not accessible from the external world; however, the cloud itself has a internal thing which is for communicating between this different component and providing services and a external link which gives a connectivity to the external one.

So, this it based on this neutron type of services; if you are using open stack. So, it has a pluggable architecture that supports many popular networking vendors and technology. So, it is interoperable between different networking vendors and technologies so that is feasible in open stack.

(Refer Slide Time: 14:04)

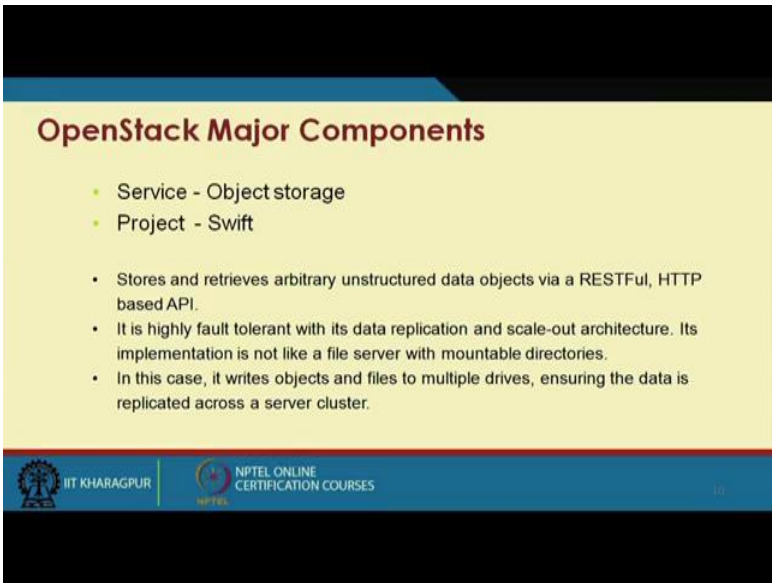

The next one is the storage service; which comes under project swift, store and retrieves arbitrary structure unstructured data objects as a RESTFul, HTTP based API.

So, it is a stores and retrieves unstructured data objects as a RESTFul; I believe that you know you understand about RESTFul services; RESTFul and HTTP based API. So, it is highly fault tolerant with data replication and scale out architecture; its implementation is not like file server with mountable directories etcetera. So, it is a basically a fault tolerant and system with a scale out architecture, it is not they simple file server. So, it is much

more than that. And in this case, it writes objects and files to multiple devices ensuring the data is replicated across the server clusters.

So, in order to make fault tolerant it writes into the things; as if you remember when we are talking about data on cloud or data services on the cloud, in our lecture or when we discussed about those. So, the replication is the scenario if discuss with the 3 replications.

So, it replicate the data into 3 places, so whenever there is a right operations; so all these right will be having should sink with all the replicas. Whereas, in case of read over sense in any of the replica can response to the things; in never the less this storage service of swift also provide such structures to means fault tolerant and services.

(Refer Slide Time: 15:50)

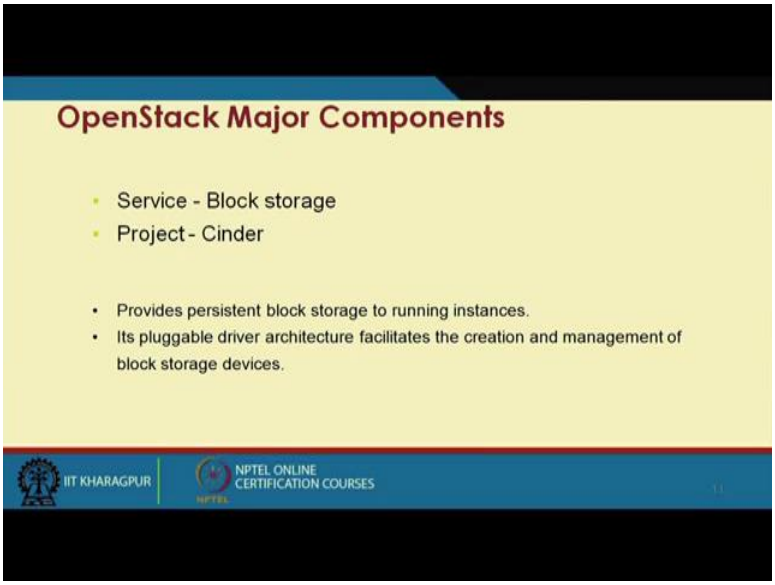

So, another type of storage service is your block storage, which is provided by cinder provide persistent block storage to running instances. So, it is a persistent block storage to the running instances; it is pluggable driver architecture facilitate the creation and management our block storage. So, it is a; again it is a persistent storage and persistent block storage and provides a pluggable driver architecture to facilitate creation and management of block storage things.

(Refer Slide Time: 16:25)

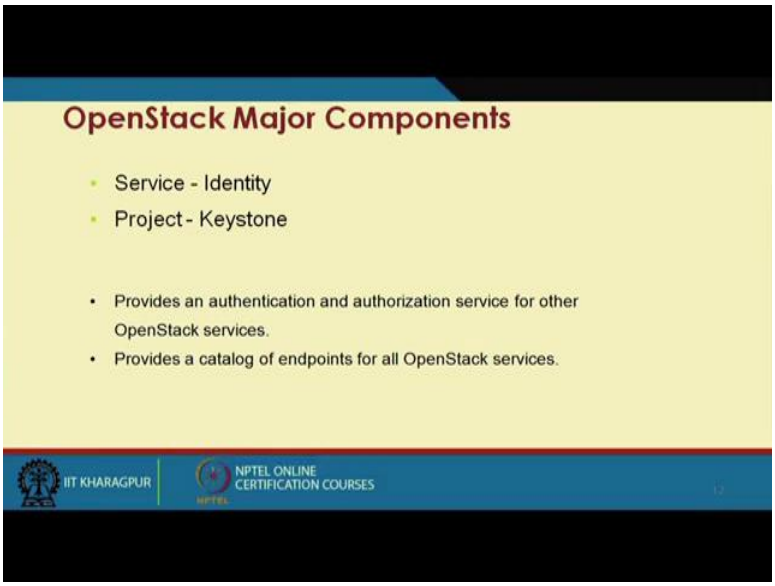

Then another component we are having which is not directly under compute or storage, but plays a vital role in realizing the cloud is the identity service or what you say keystone under the keystone project.

Provides authentication and authorization service for other open stack services. So, it is a identity service and provides authorization and authentication of services for other things. Provides a catalogue of end points for all open stack services; so, also it along with that it provides a catalogue of end points of open stack services that how those services are defined and type of things.

(Refer Slide Time: 17:10)

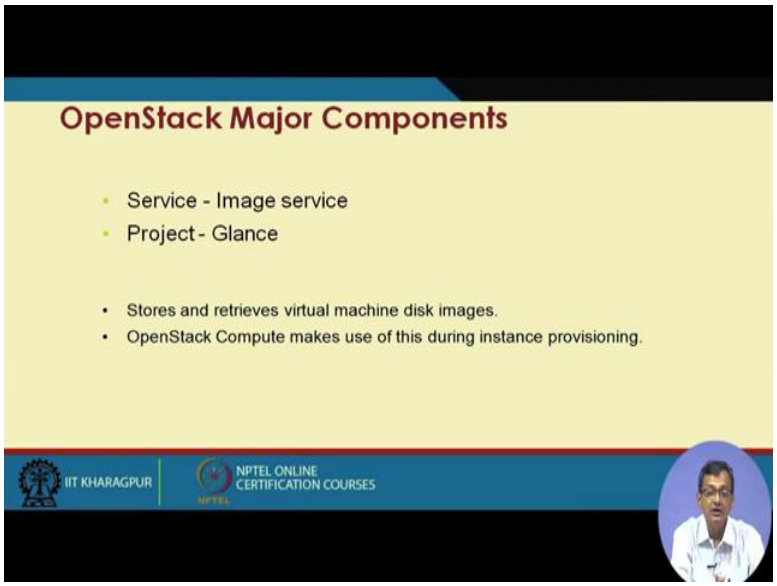

So, the next component is glance or what we say that image services. So, it basically open sack support separate type of images or image services where which can be loaded or which can be instantiated into that different VMs. So, I can have different VM and different images of say operating systems and other things can be instantiated on the difference VMs based on the user needs and requirement.

So, these sorts of services are provided by that; it is glance project or the glance service glancing service. So, it stores retrieve virtual machine disk images, so the virtual VM disk images are stores and retrieves by this glance. And open stack compute makes use of during the instance provisioning; as we are mentioning, whenever the open stack have instance provisioning; then they use this storage services.

Like you can have this image services; so, like you can have different images say for operating systems. So, you can have different flavor or images or other operating systems and when you instance; based on the requirement of the user, those are instantiated by this compute services; so, this image repositories are there in the open stack.

(Refer Slide Time: 18:54)

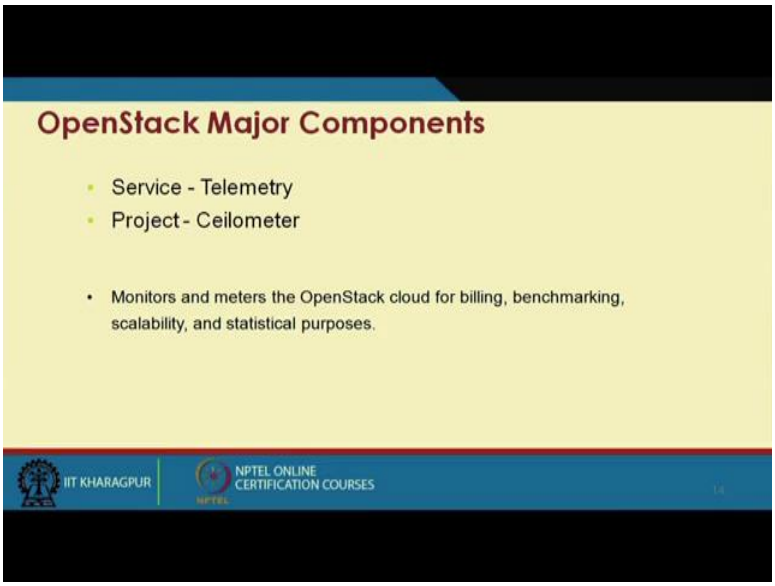

This is another service of telemetry; which is ceilometers monitors and meters the open stack cloud for billing, benchmarkings, scalability, statistical purposes.

So, this is important for overall metering of the cloud services as we have discussed in our initial lecture, as we understand that cloud is a meter service; that means, whatever uses and type of things are meter. So, that means this particular ceilometer or the telemetry services is in cloud, in open stack helps us in monitoring and metering the open stack cloud for billing, benchmarking, scalability, statistical purposes; it is not only for meeting, it is required that as you are having this measured services.

So, you can do benchmarking the address; the scalability resource and statistical analysis; like you want to know that what is the loading, how it be a, type of things. So, those type of things are available; so it is a: though it may not be directly contributing to the compute or storage or networking; which use anything that is the major component which allows us to work on those, allows us in this aspect.

But never the less it plays a vital role in having, in helping in realizing these metered service of the things along with statistical analysis or statistical measurement of things to charge the performance of the things, which not only help in building of the resource uses, also it helps in understanding the future requirement and how the infrastructure at the backbone need to be increased and type of things, it help in realizing those requirements also.

And maybe the train, maybe the periodical requirement; I can say that every one day this is the load; weekends load is less; however, some time; it is some particular time etcetera. So, this overall for analyzing over all the overall this performance of the things; what we required is more about different sort of data, which this telemetry service provides us.

(Refer Slide Time: 21:37)

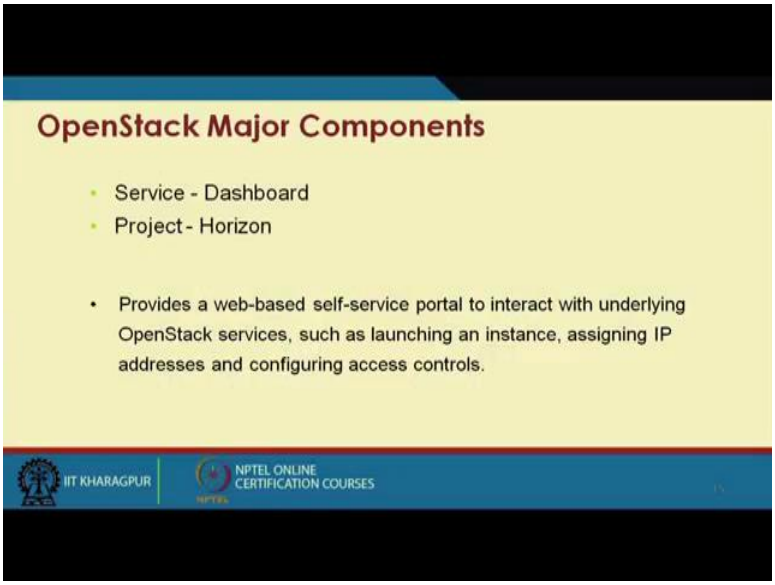

Then we have the component of open stack dashboard; that is what we say under project horizon. So, its provides a web base self service portal to interact with underlining open stack services; such as launching an instance, assigning IP addresses and configuring access control.

So, it is a dashboard under project horizon provides a well web base self service portal; to interact with underline in open stack services such as launching, instance launching, assigning IP address and configuring; actually we will show in our demo that how we are using this open stack dashboard to user management, to resource management. Like we want to assign a VM or assign IP address or if an means configured the access controls, loading images etcetera all can be done using this dashboard.

So, it is also important aspects and it is the; what we can say frontend interface for the administrator to manage the cloud. So, it is extensively used for management of the cloud.

(Refer Slide Time: 23:09)

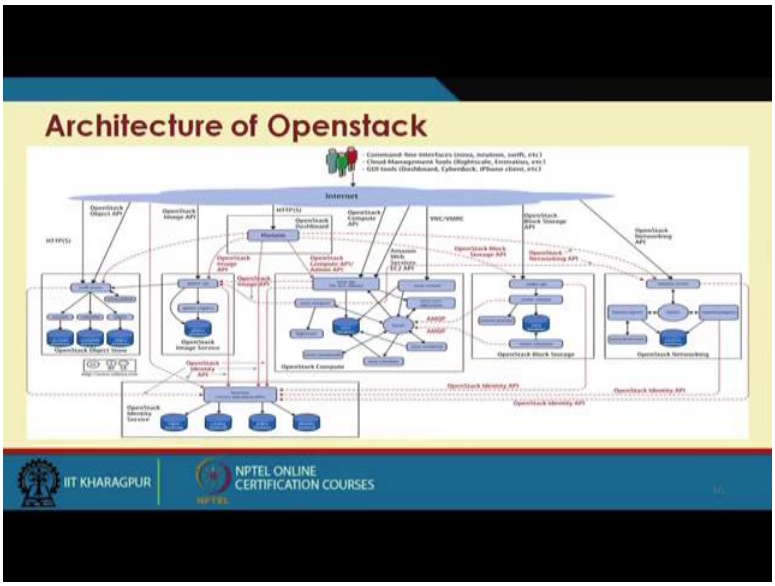

So, if we look at the overall architecture then what we see that we have different components. So, this is open stack object storage as we have discussed that is which is realized to swift; that open stack image services which is used to glance, open stack compute services which is realized to nova; realized by nova, there are block storage which is cinder; open stack networking; which is the neutron.

And open stack identity services which is keystone. So, there are different type of services and if we have seen that horizon which allows us to look at the dashboard. So, this that open stack dashboard and this is the internet that user comes on the command line interfaces; nova neutron swift etcetera, cloud management tools like right scale; other tools GUI tool dashboard, cyber duck and so on and so forth.

So, there is the user interface for the things and then it basically; based on its requirement it words into this. So, it shows that over how this overall at the top level; open stack modules are interconnected, how the process flows goes on into this open stack module. So, this is the; a overview of the thing; there are few more site of the individual component, we will not go to the much nitty-gritty of the thing, but nitty-gritty of the material, but try to see that what are the different type components.

(Refer Slide Time: 24:55)

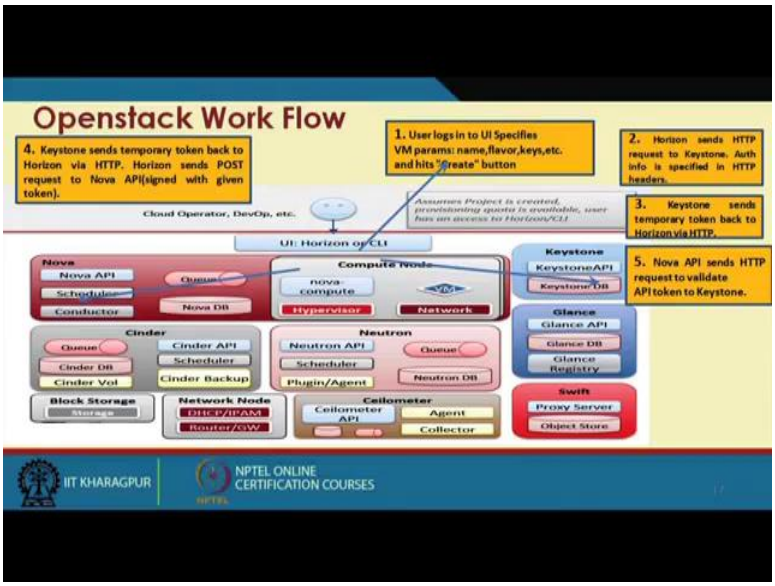

Like, if you look at the open stack workflow; so, these are the different components are there nova for compute, cinder for block storage; these are the compute nodes in the nova. There are other components within the nova, there is a cinder also have its own component. Neutron is the networking part; networking aspects; so, neutron API scheduler plug in, queues etcetera these are all networking related component.

We have ceilometers; which is primarily that the metering service or the telemetric component and there are keystone for the identity services. So, user log scene to UI specific VM parameters like name, flavor, keys etcetera and hits the create button. So, they do it on the means when user log scene with the particular creative VM. Then the horizon sends the HTTP request to keystone; authenticate information as specified in the HTTP header.

So, horizon or the dashboard service sends the information to this; I will re service then we are at the keystone sends the temporary token back to horizon via HTTP and nova API sends HTTP request to; now after that keystone sends the HTTP, then the keystone says the temporary token back to the horizon, via HTTP; horizon sends the post request to the nova API and send them. So, once that authentication is there then the horizon sends the things to the compute server; which in turned as it other operations. Now, finally the nova API sends HTTP equates to validate API token to case 2; so, this way the whole thing goes on.

(Refer Slide Time: 26:53)

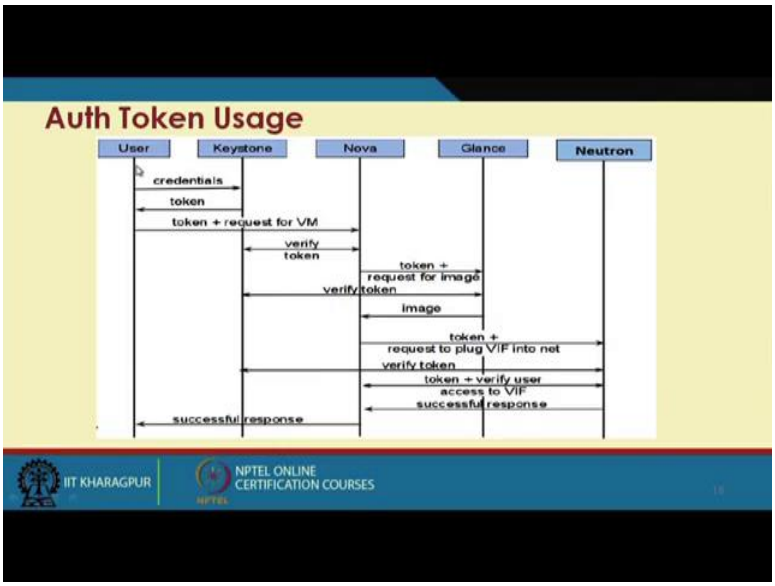

And if we look at the token exchange mechanism; so, the user to the keystone it is; for the identification or authentication, then it goes to the nova, then the glance and the neutron. So, this is what we say; if you look at that process flow of that particular authentication token uses.

(Refer Slide Time: 27:16)

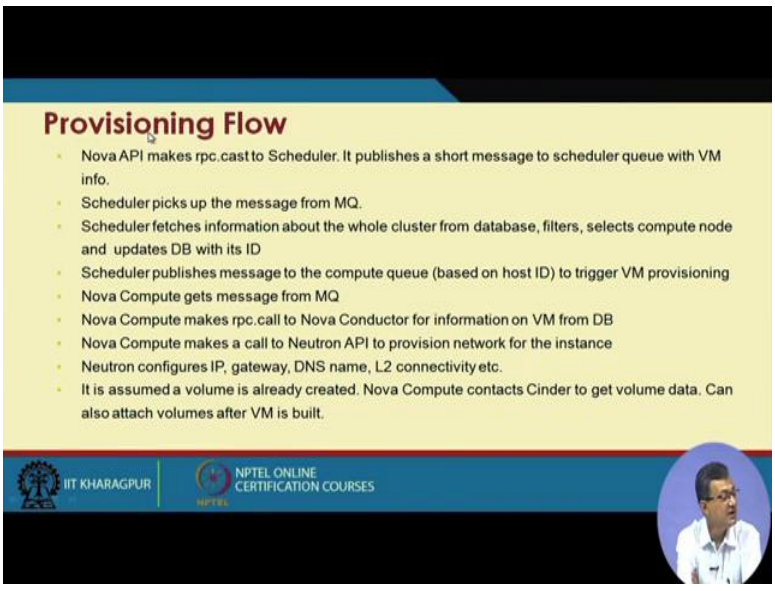

Similarly, there is a provisioning flow nova appearing makes a rpc cast to scheduler; scheduler picks up the message from the message queue, scheduler fetches information about the whole plaster from the databases, data base filters; selects compute nodes

scheduler, publishes message to the compute queues. Nova compute gets message from the nova message queue MQ and nova compute makes rpc called to nova conductor and so on and so forth. So, that is how that provisioning is made into the open stack.

(Refer Slide Time: 27:54)

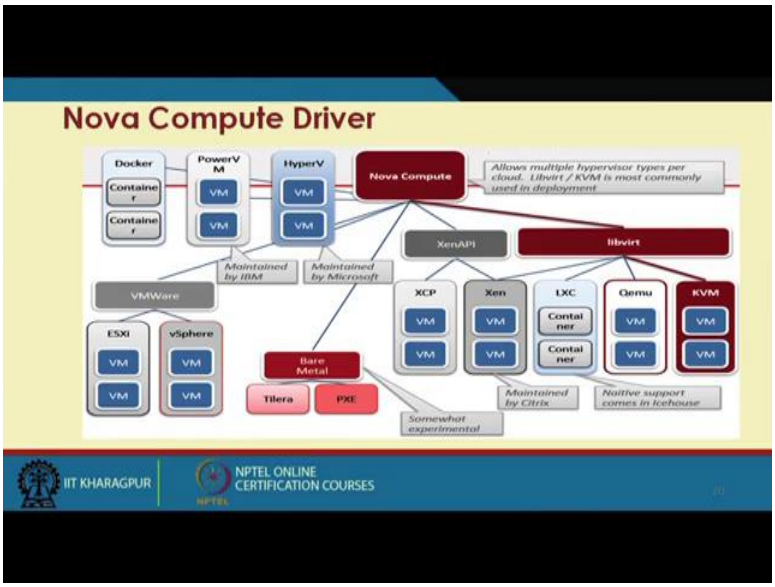

And then we have individual component like compute driver; which are expansion of the things. We are not going detail into the open stack thing; those who are interested in particular things can refer to their resources.

Similarly, we have neutron architecture which is the networking architecture and as such neutron also have several component into; a under its folder.

(Refer Slide Time: 28:27)

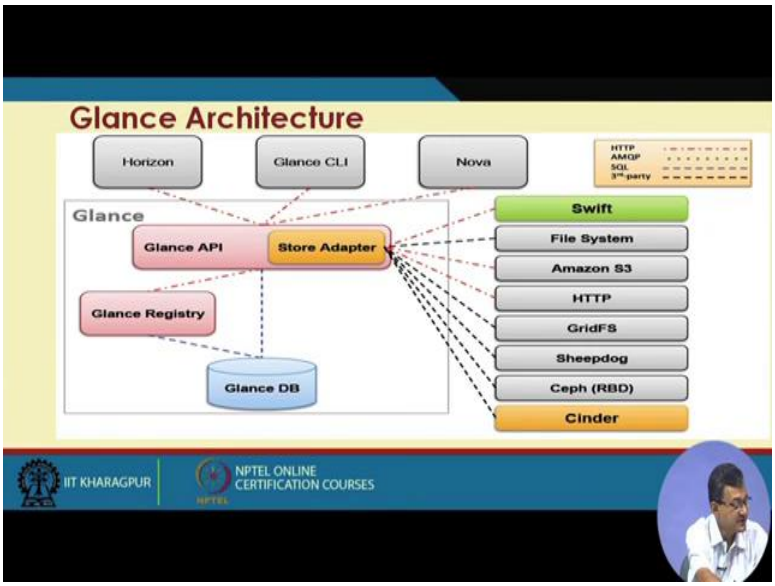

The glance which is for image service; so, it has glance databases, can registry that what sort of services is there; then glance API and storage adopter. So, we have again few components under the glance.

(Refer Slide Time: 28:52)

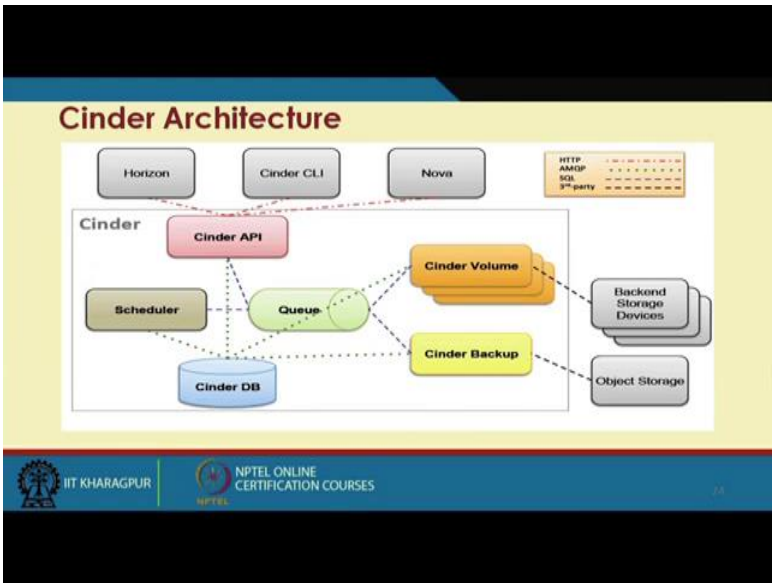

Similarly, cinder architecture which is cinder architecture which is called block storage; so these are the major component of the cinder; that is cinder databases, scheduler, API, volumes and a backup service of this cinder.

(Refer Slide Time: 29:10)

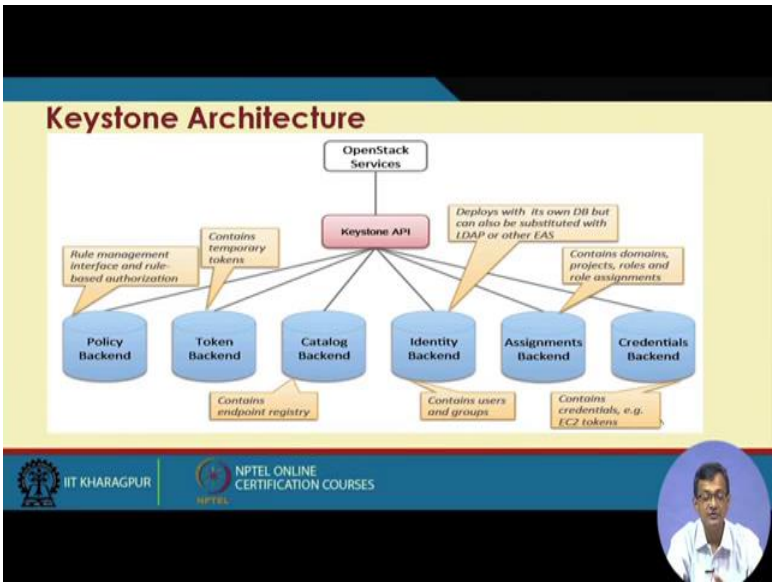

Keystone as we have already discussed is the identity type of service and it has different modules or different components. Policy back in token, tokenization, catalogue service, identity service, assignment backend and credentials backend, so these are the different type of component or services under this keystone identity service.

(Refer Slide Time: 29:38)

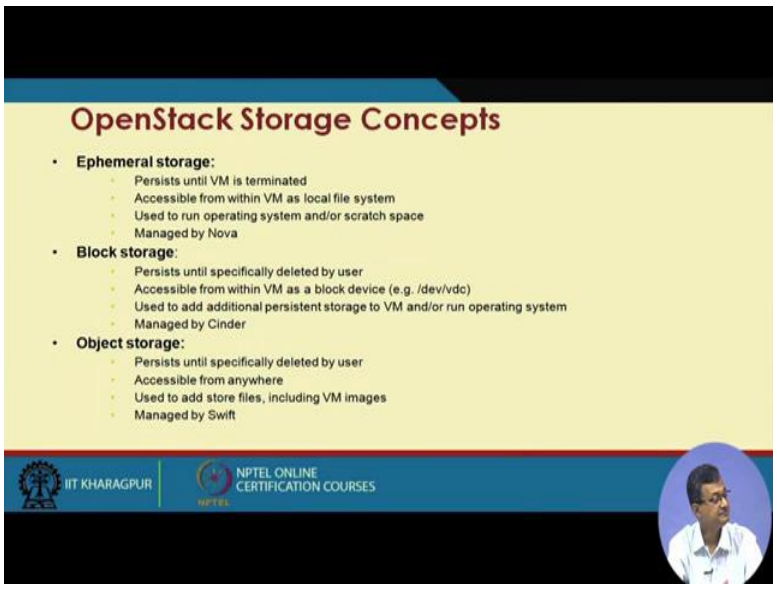

So, if you look at the open stack storage concepts; it is in the similar line with the other cloud storages; like ephemeral storage; persists until the VM terminated, accessible from within VM as local file systems; this is the ephemeral storage. So, once the VM is

terminated; it also goes off. Used to run operating system and or scratch space like as it is ephemeral, so it is primarily used to run operating system which are loaded as and when things are required. And it is also can be user a scratch space and managed by the nova; that we have seen block storage persist until specifically deleted by user.

So, it is a block storage, so it persists until specifically deleted by the user accessible within VM as a block service. Used to add additional persistent storage to VM and to operating systems, so it a used to add additional storage to the things; otherwise you are with the VM taking as storage only and it is managed by cinder.

Then we have a object storage, which is managed by swift persists until specifically deleted; accessible from anywhere, used to add and store files including VM images. So, this VM images which are managed by some of the images managed by the glance services are also used to add store files into the object storage and this and the object storage is managed by as we have discussed earlier is by swift.

(Refer Slide Time: 31:34)

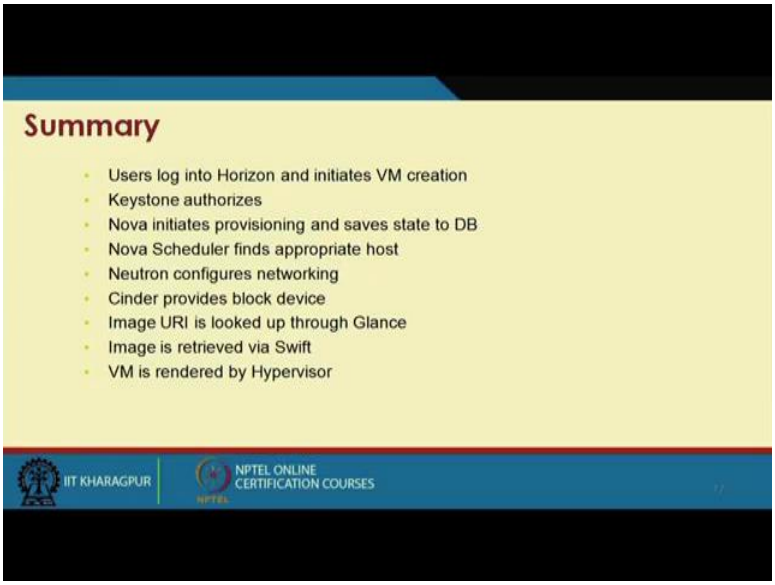

So, in summary if you have a quick overview of the whole this open stack, the user logs in horizon and initiates the VM creation. Keystone authorizes it, nova initiates provisioning and saves state to the database. Nova scheduler finds the appropriate host. Neutron configures the networking aspects, cinder provides block devices. Image URI is looked up through glance.

Image is retrieved via swift and VM is rendered by a hypervisor. So, this is in a overall that a brief overview of these open stack; open source cloud. So, what we will do next is a demo on these things; I will try to give a live demo of our open stack installation at IIT Kharagpur, as I mentioned and then we will see that how VM can be created and all those things.

Thank you.