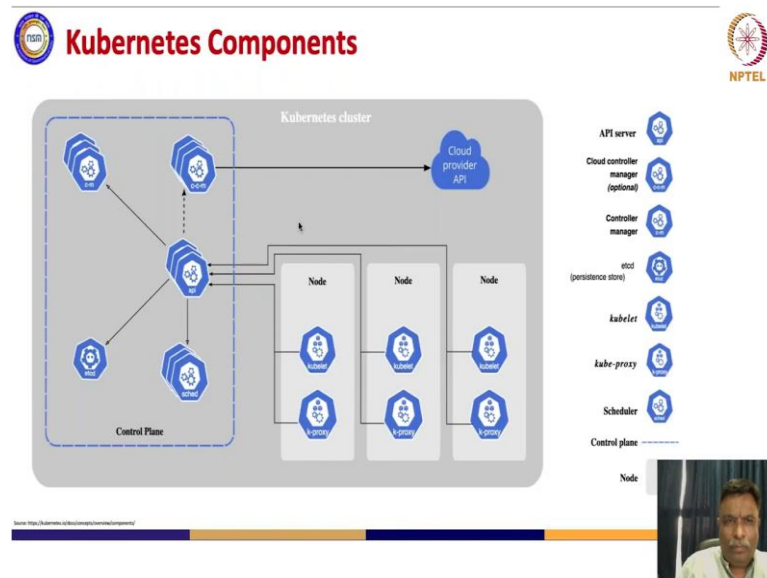


Applied Accelerated Artificial Intelligence
Prof. Satyadhyam Chickerur
School of Computer Science and Engineering
Indian Institute of Technology, Madras

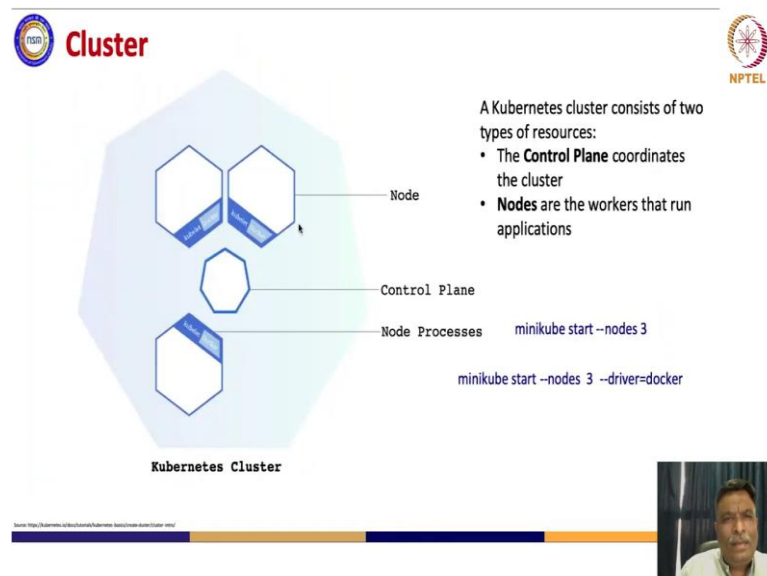
Lecture - 11

DeepOps: Deep Dive into Kubernetes with deployment of various AI based Services
Session I - Kubernetes Part - 2

(Refer Slide Time: 00:32)



(Refer Slide Time: 00:34)



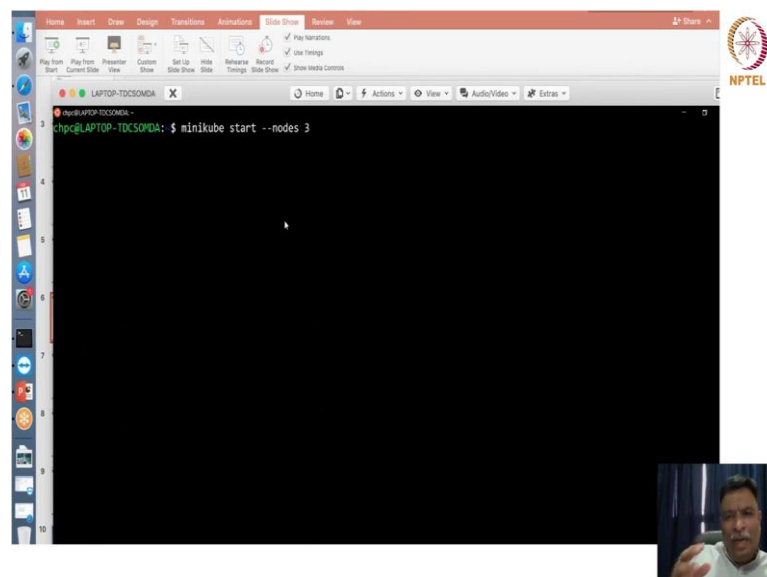
So, let us try to go into a bit of a detail of what a cluster is. If you see this diagram, ok if you see this diagram, this is let us say a Kubernetes Cluster. What we are trying to show here is there is a control plane here; this control plane is responsible for managing this cluster. These are various nodes. Each of these nodes are running certain docker containers and then there is a kubelet, ok.

So, this is a control plane. These are various nodes, this another way of representing a cluster, but this makes things a bit easier. So, I thought we will discuss this. So, this is a control plane. These are various nodes. Each of these nodes, ok will have certain containers.

So, if you technically connect to the previous slide we can assume a Kubernetes cluster will consist of two types of resources, control plane coordinates the cluster and nodes are the workers that run the application. So, this is a very simplistic definition of what a Kubernetes cluster is.

So, now, let us try to actually see, ok a Kubernetes cluster. Do a hands on thing and start understanding that, ok.

(Refer Slide Time: 02:10)

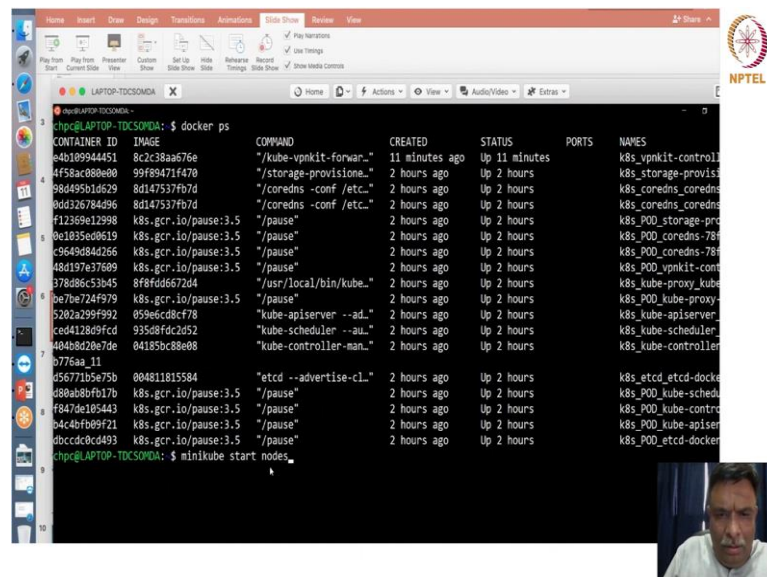


So, the first thing which we do is; so, I am using something which is called as minikube here. A minikube basically is a stripped down version of Kubernetes and it is a single node cluster, right for learning Kubernetes.

So, I am using that to make you people familiar with the concept of Kubernetes. And tomorrow we will see, the actual dashboard with certain specific things, right and then do DeepOps on it. And, then do certain things like the dashboard analysis and something like that, ok.

So, now, let us try to do this minikube is a stripped down version of Kubernetes, ok. So, we will start minikube start nodes 3. So, I hope this is visible; we will increase the size of this a bit, yes. So, when I say minikube start nodes 3 before that let me just go and check whether do I have any dockers running, ok.

(Refer Slide Time: 04:03)



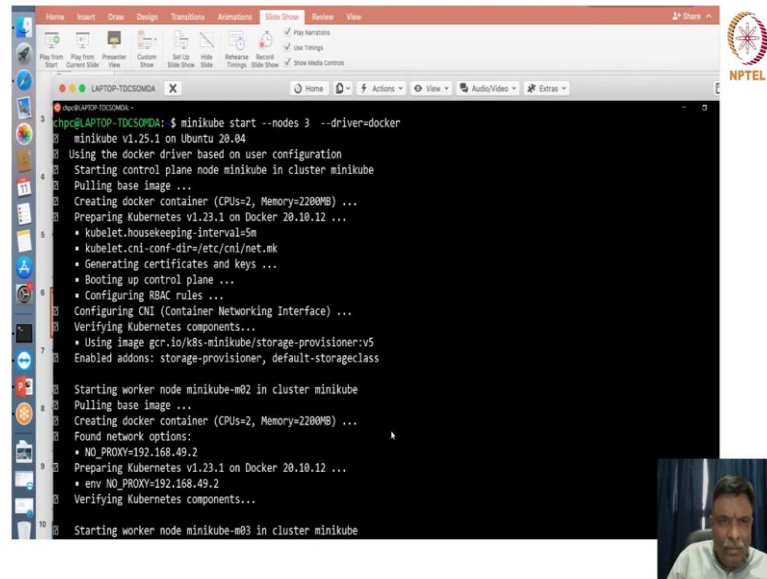
```
chpc@LAPTOP-TDCSOMDA: $ docker ps
```

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS	PORTS	NAMES
e4b109944451	8c2c38aa676e	"/kube-vpnkit-forwan..."	11 minutes ago	Up 11 minutes		k8s_vpnkit-control
4f58ac080e00	99f89471f470	"/storage-provisione..."	2 hours ago	Up 2 hours		k8s_storage-provis
98d49501d529	8d1475377b7d	"/coredns -conf /etc..."	2 hours ago	Up 2 hours		k8s_coredns_coredn
0d0326784d96	8d1475377b7d	"/coredns -conf /etc..."	2 hours ago	Up 2 hours		k8s_coredns_coredn
f12369e12998	k8s.gcr.io/pause:3.5	"/pause"	2 hours ago	Up 2 hours		k8s_POD_storage-pro
0e1035e00619	k8s.gcr.io/pause:3.5	"/pause"	2 hours ago	Up 2 hours		k8s_POD_coredns-78
e9649d044766	k8s.gcr.io/pause:3.5	"/pause"	2 hours ago	Up 2 hours		k8s_POD_coredns-78
40d137e37600	k8s.gcr.io/pause:3.5	"/pause"	2 hours ago	Up 2 hours		k8s_POD_vpnkit-cont
378d06c53b45	8f8fd6672d04	"/usr/local/bin/kube..."	2 hours ago	Up 2 hours		k8s_POD_kube-proxy
0e7be724f970	k8s.gcr.io/pause:3.5	"/pause"	2 hours ago	Up 2 hours		k8s_POD_kube-proxy
5202a39f9992	0596cd8c7f70	"kube-apiserver --ad..."	2 hours ago	Up 2 hours		k8s_kube-apiserver
c0841380fcd	93508fcd052	"kube-scheduler --au..."	2 hours ago	Up 2 hours		k8s_kube-scheduler
0a4b8d20e7de	04185bc8e08	"kube-controller-man..."	2 hours ago	Up 2 hours		k8s_kube-controller
0776aa11						
056771b5e75b	004811815584	"etcd --advertise-cl..."	2 hours ago	Up 2 hours		k8s_etcd-etcd-dock
080ab0bfb17b	k8s.gcr.io/pause:3.5	"/pause"	2 hours ago	Up 2 hours		k8s_POD_kube-sched
f847de185443	k8s.gcr.io/pause:3.5	"/pause"	2 hours ago	Up 2 hours		k8s_POD_kube-control
04c4bf089f21	k8s.gcr.io/pause:3.5	"/pause"	2 hours ago	Up 2 hours		k8s_POD_kube-apiser
dbccdc0c493	k8s.gcr.io/pause:3.5	"/pause"	2 hours ago	Up 2 hours		k8s_POD_etcd-docker

```
chpc@LAPTOP-TDCSOMDA: $ minikube start nodes 3
```

So, these are certain dockers which are running, ok which are related to some other issues, but for the time being assume that whatever we are thinking of in terms of minikube node, right those are not running. We will see that. So, I just wanted to show you that when we start our cluster; so let me clear this and then let me start, minikube start nodes 3.

(Refer Slide Time: 04:50)



```
hpc@LAPTOP-TDCSOMDA:~$ minikube start --nodes 3 --driver=docker
minikube v1.25.1 on Ubuntu 20.04
Using the docker driver based on user configuration
Starting control plane node minikube in cluster minikube
Pulling base image ...
Creating docker container (CPUs=2, Memory=2200MB) ...
Preparing Kubernetes v1.23.1 on Docker 20.10.12 ...
  • kubelet.housekeeping-interval=5m
  • kubelet.cni-conf-dir=/etc/cni/net.m
  • Generating certificates and keys ...
  • Booting up control plane ...
  • Configuring RBAC rules ...
Configuring CNI (Container Networking Interface) ...
Verifying Kubernetes components...
  • Using image gcr.io/k8s-minikube/storage-provisioner:v5
Enabled addons: storage-provisioner, default-storageclass

Starting worker node minikube-m02 in cluster minikube
Pulling base image ...
Creating docker container (CPUs=2, Memory=2200MB) ...
Found network options:
  • NO_PROXY=192.168.49.2
Preparing Kubernetes v1.23.1 on Docker 20.10.12 ...
  • env NO_PROXY=192.168.49.2
Verifying Kubernetes components...

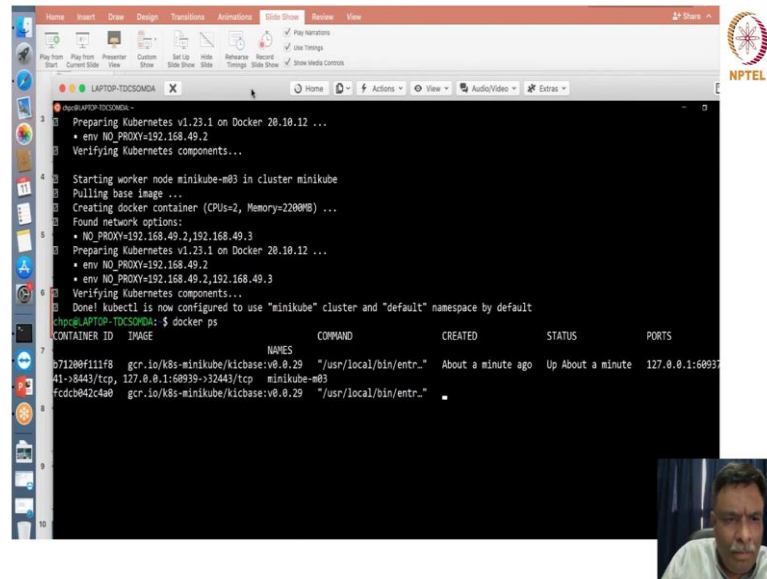
Starting worker node minikube-m03 in cluster minikube
```

So, what are we going to do here just try to understand this when this image is getting created. We are creating, right docker containers now, right. And what is that going to do? We are creating 3 nodes, ok. Each of these nodes are supposed to run, ok some particular workload as a docker. So, we are creating 3 nodes. Out of these 3 nodes we will have, 1 node as a master node and another 2 nodes as worker nodes.

So, if you see here we have created a minikube now, the name is minikube, the cluster is having one control plane node or a master node. We are creating a docker container which has got 2 CPUs with some memory, and then we are trying to start a worker node which is named as minikube-m02 which basically is the second node technically or the first worker node, ok.

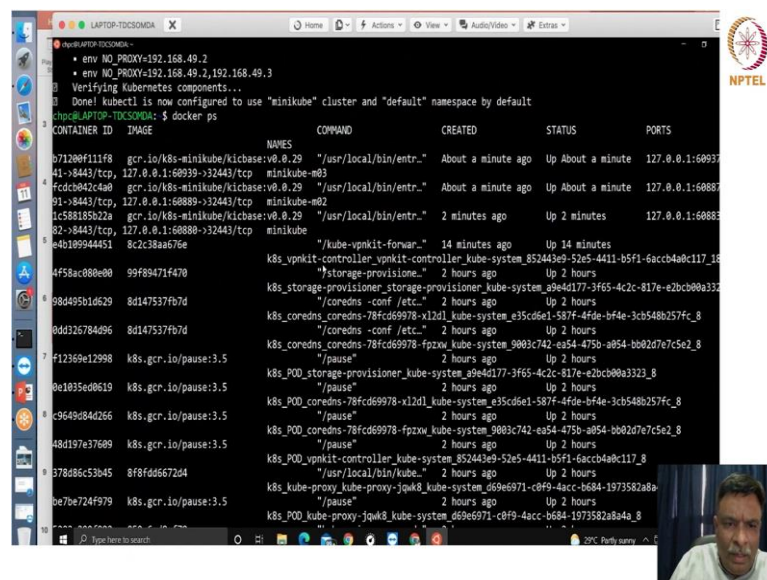
And if you see here, sorry; if you see here; what is this? Ok.

(Refer Slide Time: 06:59)



If you see here, we have another worker node which is minikube-m03, which basically is again a worker node in that particular cluster which is treated as minikube, right. So, I hope this is clear.

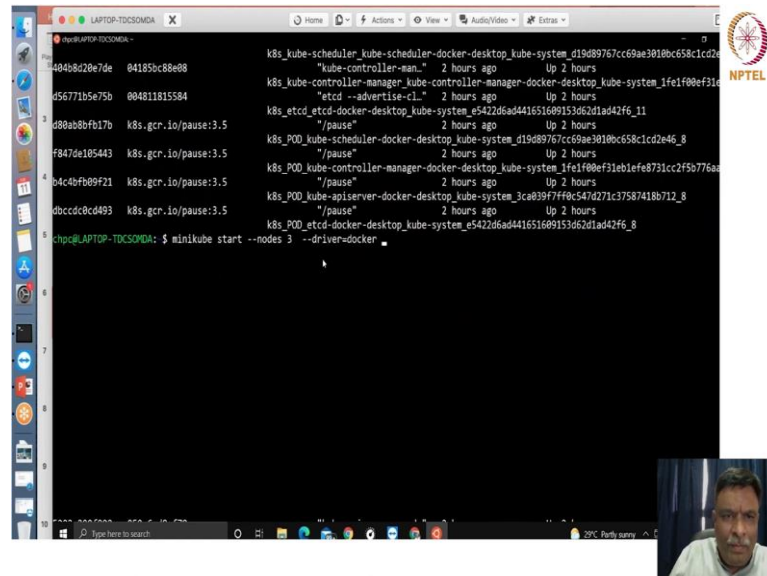
(Refer Slide Time: 07:36)



So, if you now do `docker ps`, if you see `docker ps` there are 3 things which are important for us to understand and see. We are now running 3 nodes, right, one is `minikube`, one is `minikube-m02`, and another one is `minikube-m03`, right. So, we have started `minikube`

with 3 nodes. And then; let where did it go, ok. So, minikube-m02, minikube-m03. So, I hope this is clear, right. So, we have created now.

(Refer Slide Time: 08:32)



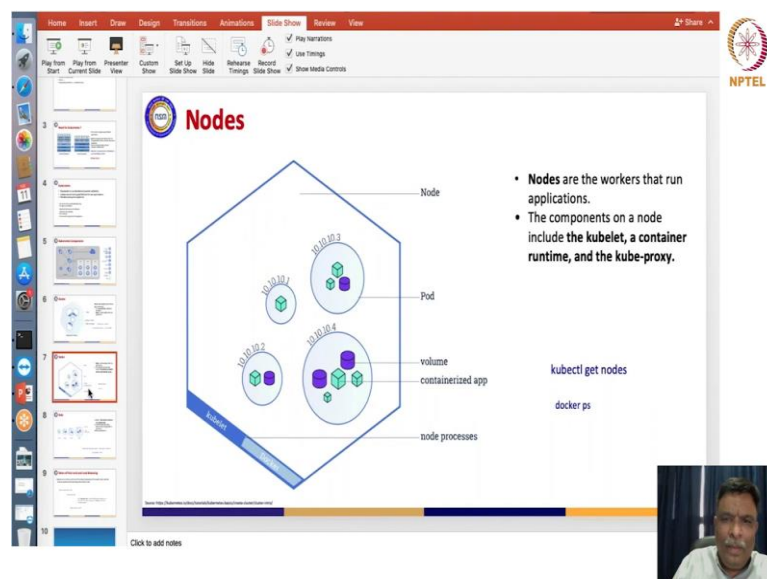
```

k8s_kube-scheduler_kube-scheduler-docker-desktop_kube-system_d19d89767cc69ae3010bc658c1cd2e
"kube-controller-man..." 2 hours ago Up 2 hours
k8s_kube-controller-manager_kube-controller-manager-docker-desktop_kube-system_1fe1f00ef31e
"etcd -advertise-cl..." 2 hours ago Up 2 hours
k8s_etcd_etcd-docker-desktop_kube-system_e542d6ad441651609153062d1ad42f6_11
"/pause" 2 hours ago Up 2 hours
k8s_POD_kube-scheduler-docker-desktop_kube-system_d19d89767cc69ae3010bc658c1cd2e46_8
"/pause" 2 hours ago Up 2 hours
k8s_POD_kube-controller-manager-docker-desktop_kube-system_1fe1f00ef31e1ef8731cc2f5b776a
"/pause" 2 hours ago Up 2 hours
k8s_POD_kube-apiserver-docker-desktop_kube-system_3ca039f7ff0c547d271c37587418b712_8
"/pause" 2 hours ago Up 2 hours
k8s_POD_etcd-docker-desktop_kube-system_e542d6ad441651609153062d1ad42f6_8
chp@LAPTOP-TDCSOMDA: $ minikube start --nodes 3 --driver=docker
  
```

Now, next thing is; why is this display giving me the issues.

Next thing is we will try to check, ok, the nodes.

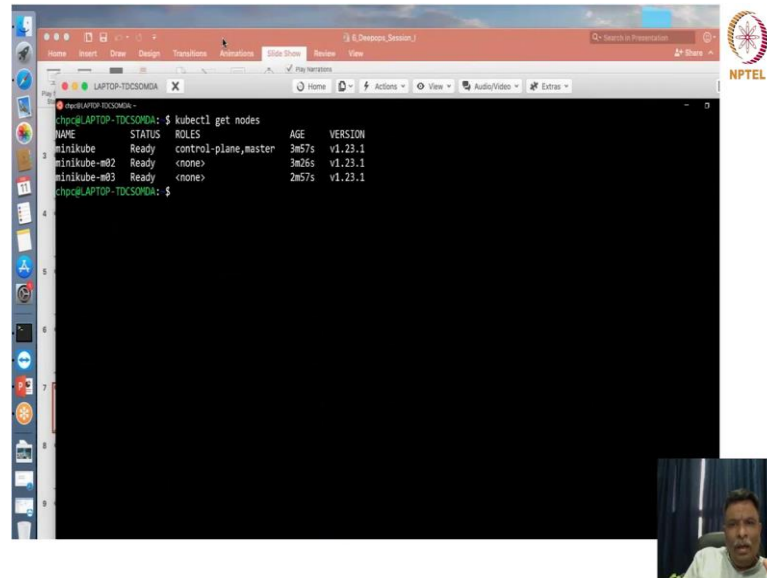
(Refer Slide Time: 08:48)



So, if you see the diagram here, we have got the node process which is a running as a docker, it has got various pods running with certain specific IP addresses.

So, nodes are the workers that are going to run the applications and we know the components on node will include kubelet, the runtime and the kube-proxy, right. So, this is a brief thing which we know.

(Refer Slide Time: 09:26)



```
chpc@LAPTOP-TDCSONDA:~$ kubectl get nodes
NAME              STATUS    ROLES          AGE   VERSION
minikube           Ready     control-plane,master   3m57s   v1.23.1
minikube-m02       Ready     <none>          3m26s   v1.23.1
minikube-m03       Ready     <none>          2m57s   v1.23.1
chpc@LAPTOP-TDCSONDA:~$
```

So, now, let us try to actually use this command which is kube control get nodes. So, if you see kubectl get nodes, what we get is we are getting what does our minikube cluster has, right. It has got a master which has a control plane which is having a name minikube.

Now, this particular minikube-m02 and minikube-m03 are two worker nodes, right; two worker nodes. So, they are ready for us to use, ok.

(Refer Slide Time: 10:24)

Pods

- A pod is the smallest execution unit in Kubernetes.
- A Pod represents a single instance of a running process in your cluster.
- Pods are ephemeral.

Pod 1 Pod 2 Pod 3 Pod 4

IP address
Initiator
containerized app

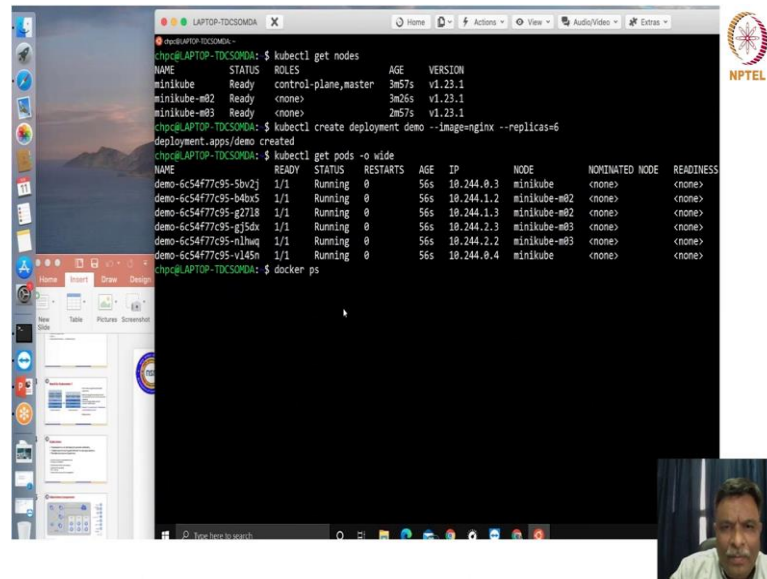
```
kubectl create deployment demo --image=nginx --replicas=6
```

```
kubectl get pods -o wide
```

Now, once you get this let us try to understand and deploy certain workload on it, ok. So, what we have tried to do is, we have tried to create two virtual machines which are treated as these worker nodes because this is a single node cluster which we are trying to have. We have a master node; we have got two virtual machines which act as nodes.

So, these two nodes technically can be your node 1, node 2 which are physical servers also. But for us to understand the concept, we are trying to have 2 worker nodes and 1 master node. And now we will use this command kube control create deployment demo with images of NGINX web server and will replicate it into 6 deployment pods, ok.

(Refer Slide Time: 11:33)



```
chpc@LAPTOP-TDCSONDA:~$ kubectl get nodes
NAME              STATUS    ROLES    AGE   VERSION
minikube          Ready    control-plane/master   3m57s   v1.23.1
minikube-m02      Ready    <none>    3m26s   v1.23.1
minikube-m03      Ready    <none>    2m57s   v1.23.1
chpc@LAPTOP-TDCSONDA:~$ kubectl create deployment demo --image=nginx --replicas=6
deployment.apps/demo created
chpc@LAPTOP-TDCSONDA:~$ kubectl get pods -o wide
NAME              READY   STATUS    RESTARTS   AGE   IP              NODE              NOMINATED NODE   READINESS
demo-6c54f77c95-5bv2j  1/1     Running   0          56s   10.244.0.3      minikube          <none>            <none>
demo-6c54f77c95-b4bx5  1/1     Running   0          56s   10.244.1.2      minikube-m02      <none>            <none>
demo-6c54f77c95-g2718  1/1     Running   0          56s   10.244.1.3      minikube-m02      <none>            <none>
demo-6c54f77c95-gj5dx  1/1     Running   0          56s   10.244.2.3      minikube-m03      <none>            <none>
demo-6c54f77c95-n1huq  1/1     Running   0          56s   10.244.2.2      minikube-m03      <none>            <none>
demo-6c54f77c95-v145n  1/1     Running   0          56s   10.244.0.4      minikube          <none>            <none>
chpc@LAPTOP-TDCSONDA:~$ docker ps
```

So, let us try to see that. What we do is we will do not going to type, no, yes.

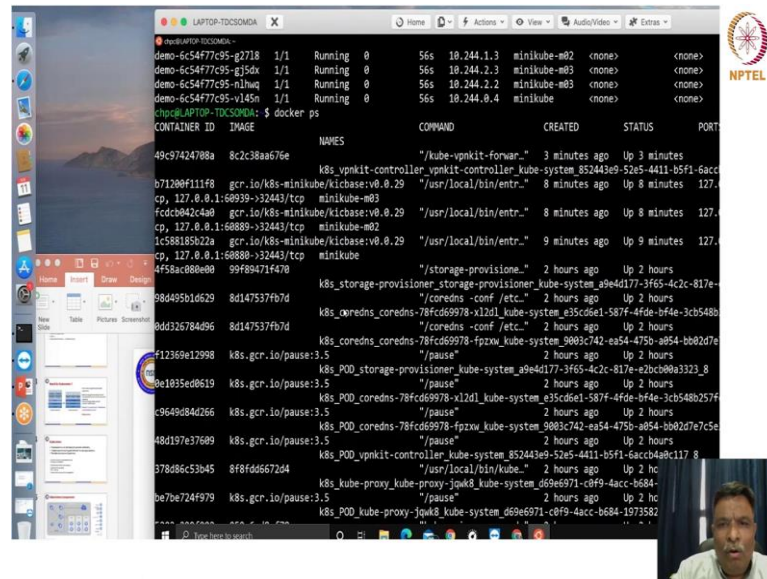
So, we are going to create, ok a deployment demo, ok and then we are going to actually use this in NGINX image, ok for just trying to what to say, trying to create, ok 6 such replicas, ok. So, let us try to do it. We have created, ok.

Now, once we have created let us try to again see whether those pods, ok are visible or not to us. And we use this command which is here which is kube control get pods -o wide. So, let us try to do this command and understand what is that we have created for demo, ok.

So, we have that command. So, do not want to type it again. So, the command is kube control get pods -o wide, ok. When you see this, what has happened actually? What effectively has happened is we have created 6 pods, which are of similar type, ok because you have created replicas of them.

And what is going to happen is it has now got allocated, ok to this cluster of ours. Here minikube is one master node it has got two pods running, minikube-m02, yes, has got two pods, ok and then minikube-m03 which is another node has got two pods. So, we created 6 pods. Those automatically got scheduled, ok on these 3 nodes automatically.

(Refer Slide Time: 14:37)

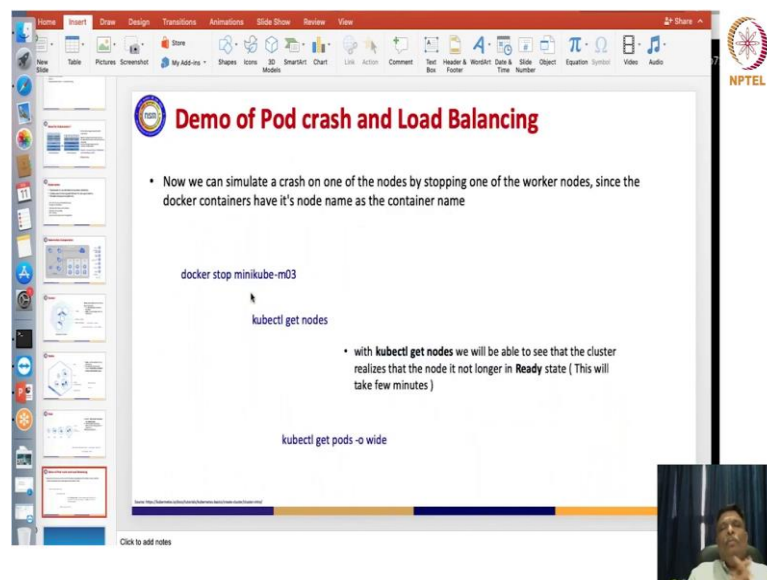


The screenshot shows a terminal window with the following output:

```
demo-6c54f77c95-g2718 1/1 Running 0 56s 10.244.1.3 minikube-m02 <none> <none>
demo-6c54f77c95-gj5dx 1/1 Running 0 56s 10.244.2.3 minikube-m03 <none> <none>
demo-6c54f77c95-nlhuq 1/1 Running 0 56s 10.244.2.2 minikube-m03 <none> <none>
demo-6c54f77c95-vl45n 1/1 Running 0 56s 10.244.0.4 minikube <none> <none>
chpc@LAPTOP-TDCSDQMA: $ docker ps
CONTAINER ID   IMAGE                                COMMAND                  CREATED        STATUS        PORTS
49c97424788a   8c2c38aa676e                        "/kube-vpnkit-forwar..." 3 minutes ago  Up 3 minutes
k8s_vpnkit-controller_vpnkit-controller_kube-system_852443e9-52e5-4411-b5f1-6acc
b71209f111f8   gcr.io/k8s-minikube/kicbase:v0.0.29 "/usr/local/bin/entr..." 8 minutes ago  Up 8 minutes  127.
cp, 127.0.0.1:60939->32443/tcp minikube-m03
fcdcb042c4a0   gcr.io/k8s-minikube/kicbase:v0.0.29 "/usr/local/bin/entr..." 8 minutes ago  Up 8 minutes  127.
cp, 127.0.0.1:60889->32443/tcp minikube-m02
1c588185b22a   gcr.io/k8s-minikube/kicbase:v0.0.29 "/usr/local/bin/entr..." 9 minutes ago  Up 9 minutes  127.
cp, 127.0.0.1:60880->32443/tcp minikube
4f58ac080e80   99f89471f47d                        "/storage-provisione..." 2 hours ago    Up 2 hours
k8s_storage-provisioner_storage-provisioner_kube-system_a9e4d177-3f65-4c2c-817e-
98d495b16629   8d147537f97d                        "/coredns -conf /etc..." 2 hours ago    Up 2 hours
k8s_coredns_coredns-78fcd69978-vl2d1_kube-system_e35c6e1-587f-4fde-bf4e-3cb548b
8dd326784d96   8d147537f97d                        "/coredns -conf /etc..." 2 hours ago    Up 2 hours
k8s_coredns_coredns-78fcd69978-fpzw_kube-system_9083c742-ea54-475b-a954-bb02d7e
f12369e12998   k8s.gcr.io/pause:3.5                "/pause"            2 hours ago    Up 2 hours
k8s_POD_storage-provisioner_kube-system_a9e4d177-3f65-4c2c-817e-e2bcb08a323_8
de1035e0619   k8s.gcr.io/pause:3.5                "/pause"            2 hours ago    Up 2 hours
k8s_POD_coredns-78fcd69978-vl2d1_kube-system_e35c6e1-587f-4fde-bf4e-3cb548b257f
c9649b842b6   k8s.gcr.io/pause:3.5                "/pause"            2 hours ago    Up 2 hours
k8s_POD_coredns-78fcd69978-fpzw_kube-system_9083c742-ea54-475b-a954-bb02d7e7c5e
48d197e37689   k8s.gcr.io/pause:3.5                "/pause"            2 hours ago    Up 2 hours
k8s_POD_vpnkit-controller_kube-system_852443e9-52e5-4411-b5f1-6accb4a8c117_8
378d86c53b45   8f8fd667264                        "/usr/local/bin/kube..." 2 hours ago    Up 2 hc
k8s_kube-proxy-jwqk8_kube-system_d69e6971-c8f9-4acc-b684-b684
be70e724f979   k8s.gcr.io/pause:3.5                "/pause"            2 hours ago    Up 2 hc
k8s_POD_kube-proxy-jwqk8_kube-system_d69e6971-c8f9-4acc-b684-b684-1973582
```

Now, we will try to do a experiment. And in the meantime, we can also see if you do docker ps, there is no information about these pods, ok running on as containers from the docker perspective correct. So, I just wanted to be show you this.

(Refer Slide Time: 15:00)



The slide is titled "Demo of Pod crash and Load Balancing". It contains the following text:

- Now we can simulate a crash on one of the nodes by stopping one of the worker nodes, since the docker containers have it's node name as the container name

```
docker stop minikube-m03
```

With `kubectl get nodes` we will be able to see that the cluster realizes that the node is no longer in Ready state (This will take few minutes)

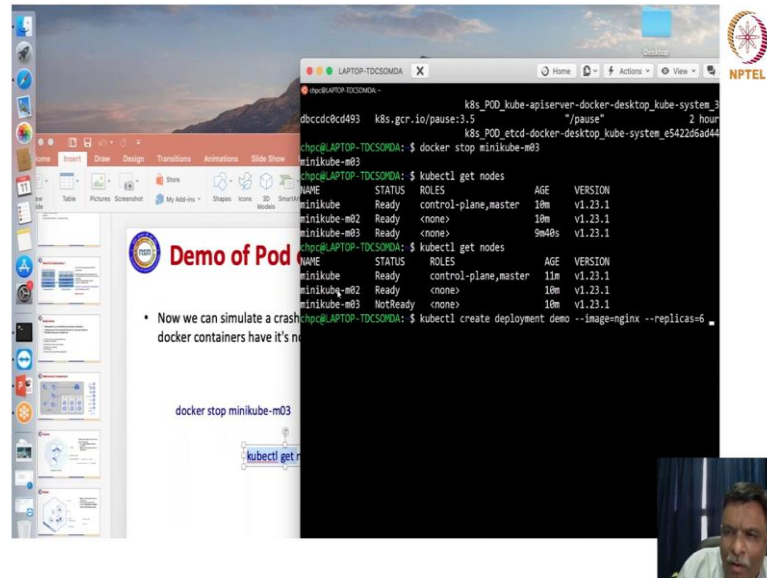
```
kubectl get pods -o wide
```

But now what we are going to do is we are going to actually do this demo of crashing a pod, ok crashing a pod and then doing some load balancing. What does that mean? That means that we can simulate a crash on one of the nodes by stopping one of the worker

nodes, ok as written here in the slide. And since the docker containers have its node name as the container name, it is very easy for us to understand.

So, let us try to do this and stop the dockers which are running on minikube-m03, ok. So, let me just do that. Let me just do docker stop, ok.

(Refer Slide Time: 15:43)

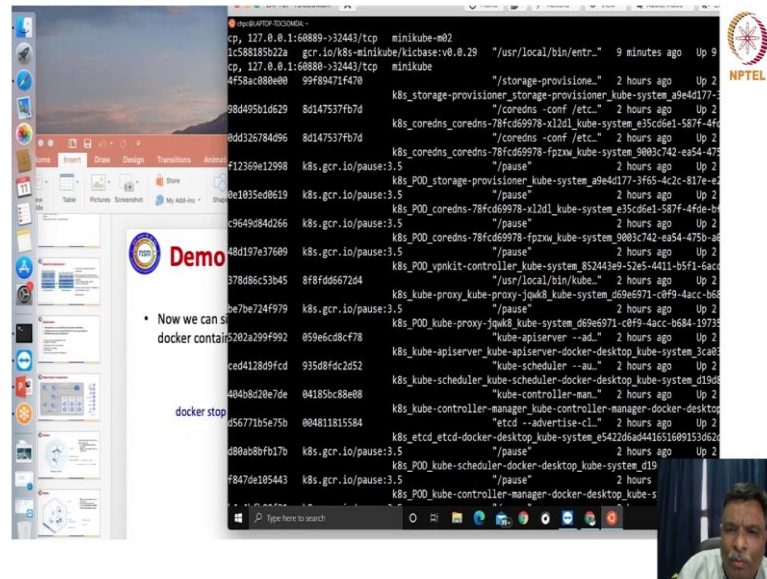


It is minikube-m03. So, once we do this it has stopped. So, now, let us try to again see this command, ok kubectl get nodes, right.

So, now, see the idea is that we have stopped the minikube-m03 docker, ok. When I say docker minikube-m03 has stopped, my node has to stop actually. So, it will take few minutes for getting that result because there has to have some synchronization, orchestration and the message passing to happen. So, let us try to see this.

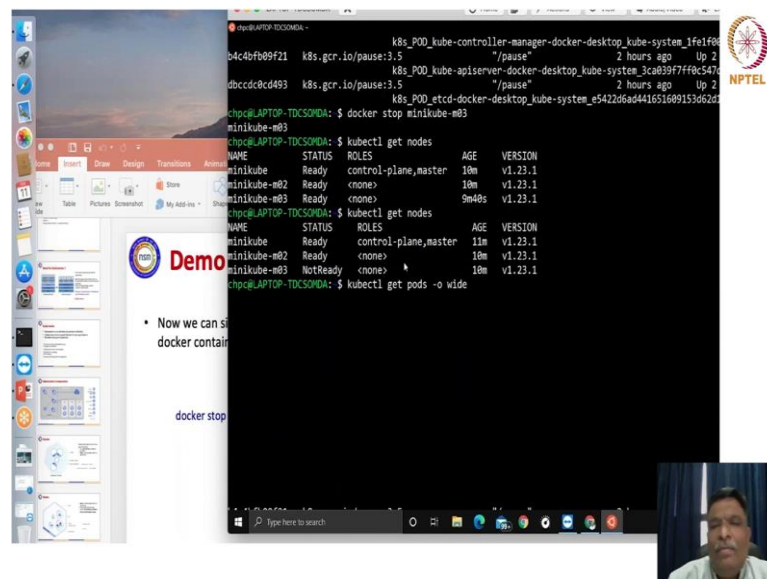
And now you have got the status as not ready. So, now, what we have tried to simulate is one of our nodes has failed. So, minikube-m03 node has failed. Now, we are only left with two nodes. So, what technically should happen? Since, Kubernetes has to do load balancing, it has to ensure, ok it has to ensure that the pods which were running on this node minikube-m03 should actually be relocated to or balanced to these two nodes which is minikube and minikube-m02, ok. So, let us try to do that also.

(Refer Slide Time: 18:15)



And what we are going to do is we are going to; sorry.

(Refer Slide Time: 18:20)



We are going to now see this information again. Please understand, we have not scheduled or we have not transferred the pods running on minikube-m03 to m02 and minikube. So, let us try to see where are those pods running.

(Refer Slide Time: 18:39)

```
chpc@APTOP-TDCSOWDA: ~  
$ kubectl get nodes  
NAME             STATUS    ROLES    AGE   VERSION  
minikube          Ready     control-plane,master   10m   v1.23.1  
minikube-m02      Ready     <none>    10m   v1.23.1  
minikube-m03      Ready     <none>    9m48s v1.23.1  
chpc@APTOP-TDCSOWDA: ~  
$ kubectl get pods -o wide  
NAME             READY   STATUS    RESTARTS   AGE   IP            NODE           NOMINATED  
demo-6c54f77c95-5bv2j 1/1     Running   0           6m28s 10.244.0.3    minikube       <none>  
demo-6c54f77c95-b4bx5 1/1     Running   0           6m28s 10.244.1.2    minikube-m02   <none>  
demo-6c54f77c95-g2j18 1/1     Running   0           6m28s 10.244.1.3    minikube-m02   <none>  
demo-6c54f77c95-g35dx 1/1     Running   0           6m28s 10.244.2.3    minikube-m03   <none>  
demo-6c54f77c95-nlhq  1/1     Running   0           6m28s 10.244.2.2    minikube-m03   <none>  
demo-6c54f77c95-vl45n 1/1     Running   0           6m28s 10.244.0.4    minikube       <none>  
chpc@APTOP-TDCSOWDA: ~  
$ kubectl get pods -o wide
```

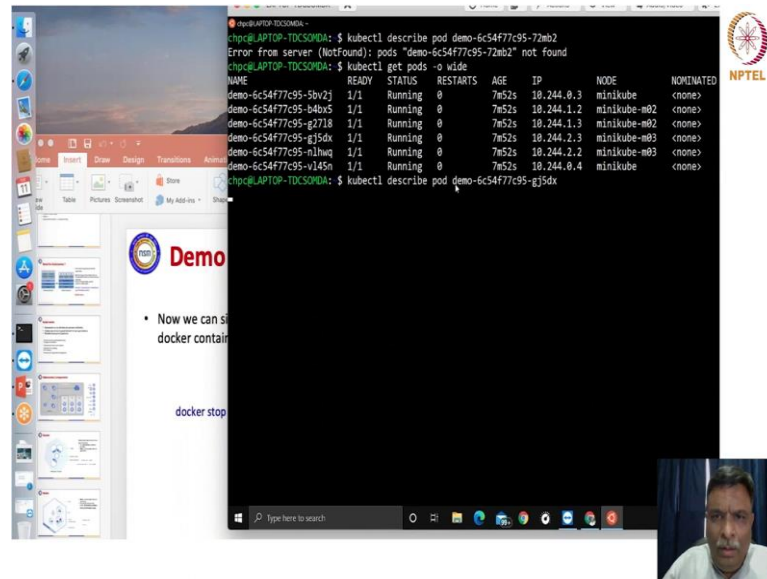
So, still if you see minikube-m03 is running this pod and minikube-m03 is running this pod and each of these nodes are running two pods each, right.

Since, I told you it will take a few minutes for us to get back that information, let us wait for a few minutes. In the meantime, let us try to see, ok the description of each pod, ok. So, why is this.

(Refer Slide Time: 19:22)

```
chpc@APTOP-TDCSOWDA: ~  
$ docker stop minikube-m03  
minikube-m03  
chpc@APTOP-TDCSOWDA: ~  
$ kubectl get nodes  
NAME             STATUS    ROLES    AGE   VERSION  
minikube          Ready     control-plane,master   10m   v1.23.1  
minikube-m02      Ready     <none>    10m   v1.23.1  
minikube-m03      Ready     <none>    9m48s v1.23.1  
chpc@APTOP-TDCSOWDA: ~  
$ kubectl get pods -o wide  
NAME             READY   STATUS    RESTARTS   AGE   IP            NODE           NOMINATED  
k8s_coredns-coredns-78fcd69978-fpznw_kube-system_9803c742-ea54-475b-a6... 1/1     Running   0           2 hours ago  Up 2  
k8s_gcr.io/pause:3.5 1/1     Running   0           2 hours ago  Up 2  
k8s_POD_storage-provisioner_kube-system_a9e4d177-3f65-4a2c-817c-47... 1/1     Running   0           2 hours ago  Up 2  
k8s_POD_coredns-78fcd69978-xl2d1_kube-system_e35c6e1-587f-4fde-b... 1/1     Running   0           2 hours ago  Up 2  
k8s_POD_coredns-78fcd69978-fpznw_kube-system_9803c742-ea54-475b-a6... 1/1     Running   0           2 hours ago  Up 2  
k8s_gcr.io/pause:3.5 1/1     Running   0           2 hours ago  Up 2  
k8s_POD_vpnkit-controller_kube-system_852443e9-52e5-4411-b5f1-6ac... 1/1     Running   0           2 hours ago  Up 2  
k8s_kube-proxy_kube-proxy-jowk8_kube-system_d69e6971-c0f9-4acc-b684-1973... 1/1     Running   0           2 hours ago  Up 2  
k8s_POD_kube-proxy-jowk8_kube-system_d69e6971-c0f9-4acc-b684-1973... 1/1     Running   0           2 hours ago  Up 2  
k8s_kube-apiserver_kube-apiserver-docker-desktop_kube-system_3ca8... 1/1     Running   0           2 hours ago  Up 2  
k8s_kube-scheduler_kube-scheduler-docker-desktop_kube-system_d19d... 1/1     Running   0           2 hours ago  Up 2  
k8s_kube-controller-manager_kube-controller-manager-docker-desktop_kube-system_3ca8... 1/1     Running   0           2 hours ago  Up 2  
k8s_etcd_etcd-docker-desktop_kube-system_e5422d6a441651609153d62... 1/1     Running   0           2 hours ago  Up 2  
k8s_gcr.io/pause:3.5 1/1     Running   0           2 hours ago  Up 2  
k8s_POD_kube-scheduler-docker-desktop_kube-system_d19d89767cc69ae... 1/1     Running   0           2 hours ago  Up 2  
k8s_POD_kube-controller-manager-docker-desktop_kube-system_3f1ef8... 1/1     Running   0           2 hours ago  Up 2  
k8s_gcr.io/pause:3.5 1/1     Running   0           2 hours ago  Up 2  
k8s_POD_kube-apiserver-docker-desktop_kube-system_3ca83947ff0c547c... 1/1     Running   0           2 hours ago  Up 2  
k8s_gcr.io/pause:3.5 1/1     Running   0           2 hours ago  Up 2  
k8s_POD_etcd-docker-desktop_kube-system_e5422d6a441651609153d62... 1/1     Running   0           2 hours ago  Up 2
```

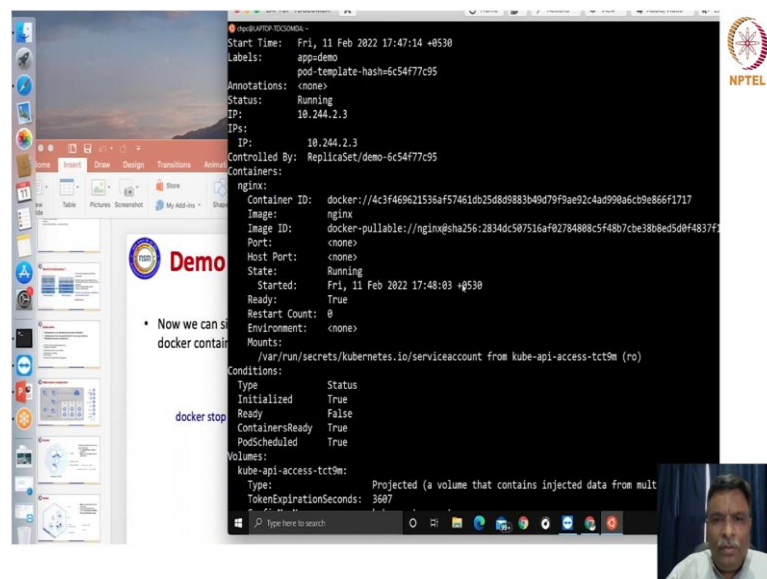

(Refer Slide Time: 19:33)



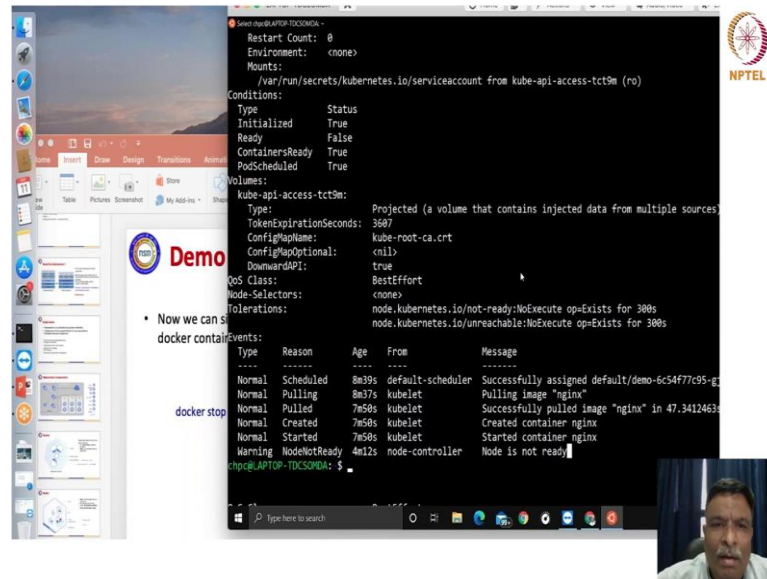
So, I will describe it using this command which is. See, described being a pod, why this pod is not available now, let us try to understand it. Because that pods name was a bit different. So, we will have to see some pod here like which was supposed to be running on minikube-m03 which we had stopped. So, let us try to see this pod, ok.

Let me just copy that information. Yes.

(Refer Slide Time: 20:51)



(Refer Slide Time: 20:54)



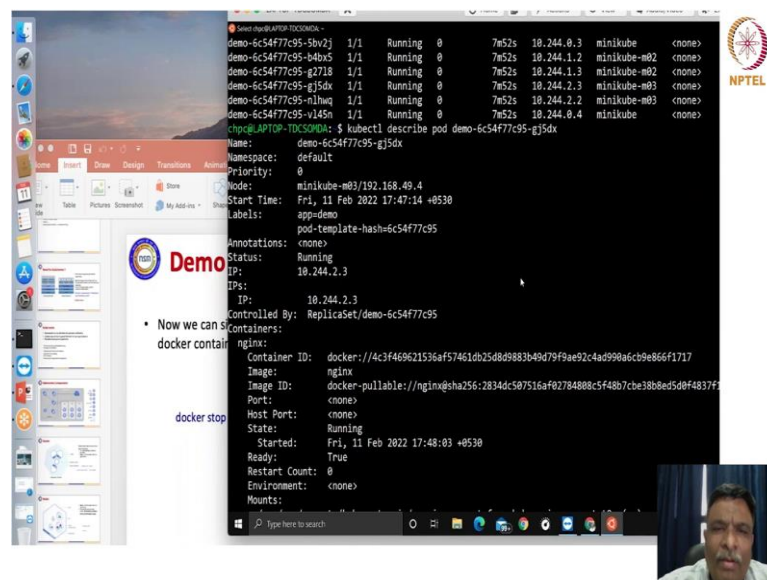
```
Restart Count: 0
Environment: <none>
Mounts:
  /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-tct9m (ro)
Conditions:
  Type             Status
  Initialized       True
  Ready            False
  ContainersReady  True
  PodScheduled     True
Volumes:
  kube-api-access-tct9m:
    Type: Projected (a volume that contains injected data from multiple sources)
    TokenExpirationSeconds: 3607
    ConfigMapName: kube-root-ca.crt
    ConfigMapOptional: <nil>
    DownwardAPI: true
QoS Class: BestEffort
Node-Selectors: <none>
Tolerations: node.kubernetes.io/not-ready:NoExecute op=Exists for 300s
              node.kubernetes.io/unreachable:NoExecute op=Exists for 300s
Events:
  Type     Reason      Age    From          Message
  ----     -
  Normal   Scheduled   8m39s  default-scheduler  Successfully assigned default/demo-6c54f77c95-gj5dx to minikube-m03
  Normal   Pulling     8m37s  kubelet        Pulling image "nginx"
  Normal   Pulled      7m50s  kubelet        Successfully pulled image "nginx" in 47.3412463s
  Normal   Created     7m50s  kubelet        Created container nginx
  Normal   Started     7m50s  kubelet        Started container nginx
  Warning  NodeNotReady 4m12s  node-controller  Node is not ready
```

Now we can stop the container

docker stop demo-6c54f77c95-gj5dx

So, if you see this information about the pod, ok it has done scheduling, pulling, created, everything happened, right started the container, created the container, but at one point in time it shows node is not ready. This is the time when we stopped the node, ok.

(Refer Slide Time: 21:16)

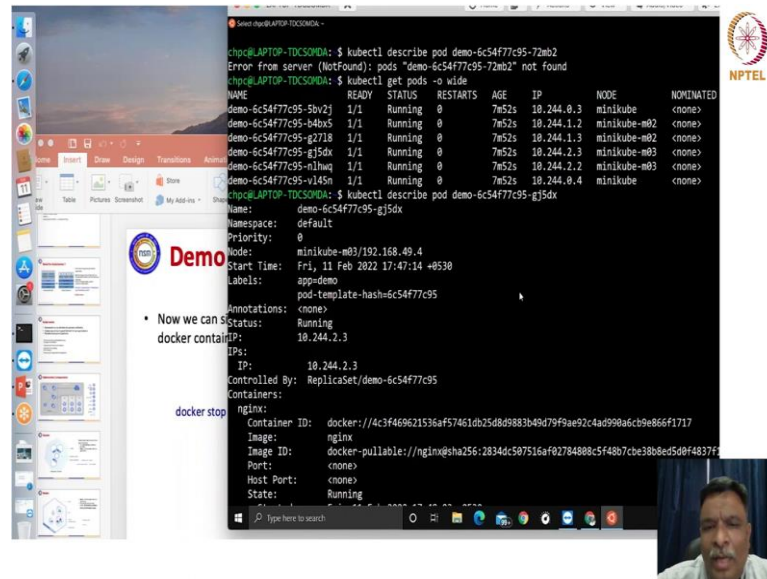


```
Name: demo-6c54f77c95-gj5dx
Namespace: default
Priority: 0
Node: minikube-m03/192.168.49.4
Start Time: Fri, 11 Feb 2022 17:47:14 +0530
Labels: app=demo
        pod-template-hash=6c54f77c95
Annotations: <none>
Status: Running
IP: 10.244.2.3
IPs:
  IP: 10.244.2.3
Controlled By: ReplicaSet/demo-6c54f77c95
Containers:
  nginx:
    Container ID: docker://Ac3f469621536af57461db2548d9883b49d79f9ae92c4ad999a6cb9e866f1717
    Image: nginx
    Image ID: docker-pullable://nginx@sha256:2834dc597516af02784888c5f48b7cbe38b8e5d8f4837f7
    Port: <none>
    Host Port: <none>
    State: Running
      Started: Fri, 11 Feb 2022 17:48:03 +0530
      Ready: True
      Restart Count: 0
      Environment: <none>
      Mounts: <none>
```

Now we can stop the container

docker stop demo-6c54f77c95-gj5dx

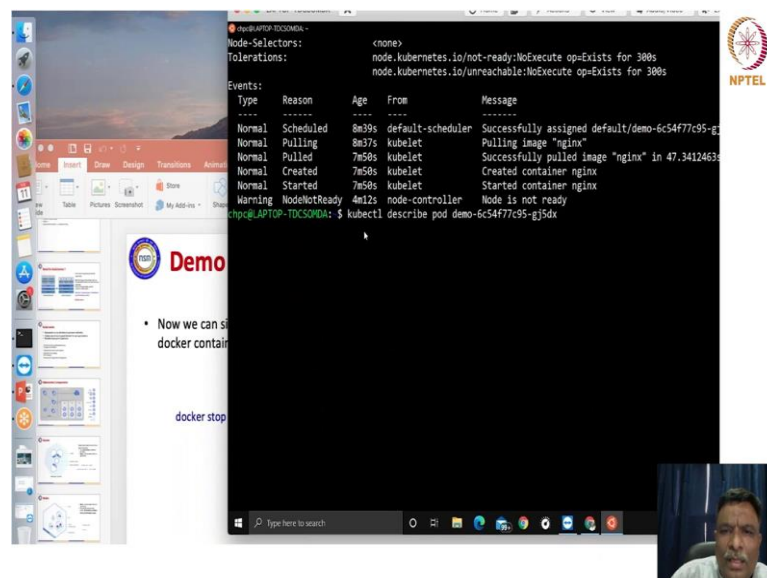
(Refer Slide Time: 21:17)



```
chpc@LAPTOP-TDCSOMDA: ~$ kubectl describe pod demo-6c54f77c95-72mb2
Error from server (NotFound): pods "demo-6c54f77c95-72mb2" not found
chpc@LAPTOP-TDCSOMDA: ~$ kubectl get pods -o wide
NAME                                READY    STATUS    RESTARTS   AGE   IP              NODE             NOMINATED
demo-6c54f77c95-5bv2j               1/1      Running   0           7s    10.244.0.3      minikube         <none>
demo-6c54f77c95-b4bx5               1/1      Running   0           7s    10.244.1.2      minikube-m02     <none>
demo-6c54f77c95-g27l8               1/1      Running   0           7s    10.244.1.3      minikube-m02     <none>
demo-6c54f77c95-gj5dx               1/1      Running   0           7s    10.244.2.3      minikube-m03     <none>
demo-6c54f77c95-nlhqj               1/1      Running   0           7s    10.244.2.2      minikube-m03     <none>
demo-6c54f77c95-vl45n               1/1      Running   0           7s    10.244.0.4      minikube         <none>
chpc@LAPTOP-TDCSOMDA: ~$ kubectl describe pod demo-6c54f77c95-gj5dx
Name:         demo-6c54f77c95-gj5dx
Namespace:    default
Priority:      0
Node:         minikube-m03/192.168.49.4
Start Time:   Fri, 11 Feb 2022 17:47:14 +0530
Labels:       app=demo
              pod-template-hash=6c54f77c95
Annotations:  <none>
Status:       Running
IP:           10.244.2.3
IPs:          <none>
Controlled By: ReplicaSet/demo-6c54f77c95
Containers:
  nginx:
    Container ID:  docker://4c3f469621536af57461db25d8d9883b49d79f9ae92c4ad998a6c9e866f1717
    Image:         nginx
    Image ID:      docker-pullable://nginx@sha256:2834dc507516af02784888c5f48b7cbe388ed5d0f4837f
    Port:         <none>
    Host Port:     <none>
    State:         Running
```

So, this has got a lot of information. We will discuss maybe in the days to come we will gradually understand, right what is this particular container ID, what is that image which you have put in, ok, what is the state of that particular pod actually, ok, like it was running, somewhere running, right, yeah state was running; at what time it started, right. So, how many times it has been restarted actually all of this. So, this is how you describe each pod.

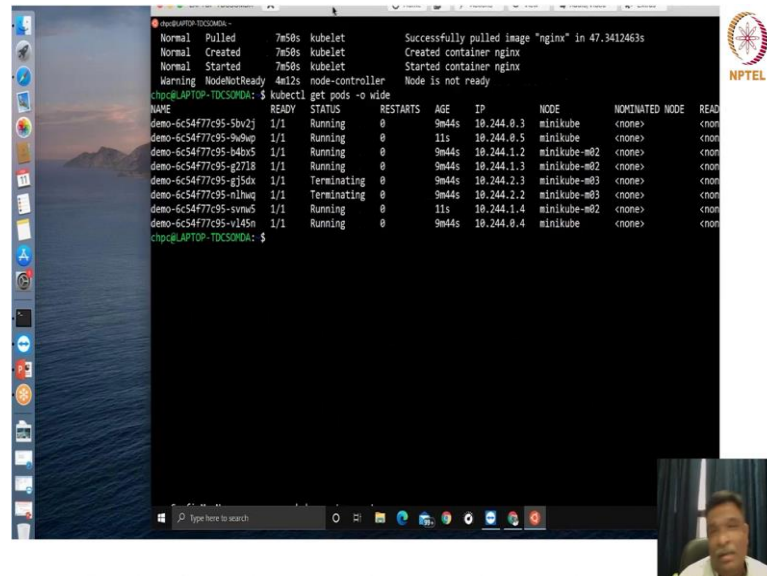
(Refer Slide Time: 21:47)



```
chpc@LAPTOP-TDCSOMDA: ~$ kubectl describe pod demo-6c54f77c95-gj5dx
Name:         demo-6c54f77c95-gj5dx
Namespace:    default
Priority:      0
Node:         minikube-m03/192.168.49.4
Start Time:   Fri, 11 Feb 2022 17:47:14 +0530
Labels:       app=demo
              pod-template-hash=6c54f77c95
Annotations:  <none>
Status:       Running
IP:           10.244.2.3
IPs:          <none>
Controlled By: ReplicaSet/demo-6c54f77c95
Containers:
  nginx:
    Container ID:  docker://4c3f469621536af57461db25d8d9883b49d79f9ae92c4ad998a6c9e866f1717
    Image:         nginx
    Image ID:      docker-pullable://nginx@sha256:2834dc507516af02784888c5f48b7cbe388ed5d0f4837f
    Port:         <none>
    Host Port:     <none>
    State:         Running
Tolerations:  <none>
node.kubernetes.io/not-ready:NoExecute op=Exists for 300s
node.kubernetes.io/unreachable:NoExecute op=Exists for 300s
Events:
  Type     Reason            Age   From          Message
  ----     ------            -
  Normal   Scheduled         8m39s default-scheduler Successfully assigned default/demo-6c54f77c95-gj5dx to minikube-m03
  Normal   Pulling           8m37s kubelet       Pulling image "nginx"
  Normal   Pulled            7m50s kubelet       Successfully pulled image "nginx" in 47.3412463s
  Normal   Created           7m50s kubelet       Created container nginx
  Normal   Started           7m50s kubelet       Started container nginx
  Warning  NodeNotReady      4m12s node-controller Node is not ready
```

So, let me just again go to and see whether we have actually been able to; yes check here now.

(Refer Slide Time: 21:51)



The screenshot shows a terminal window with the following output:

```
Normal Pulled 7m50s kubelet Successfully pulled image "nginx" in 47.3412463s
Normal Created 7m50s kubelet Created container nginx
Normal Started 7m50s kubelet Started container nginx
Warning NodeNotReady 4m12s node-controller Node is not ready

chpc@LAPTOP-TDCSOWDA:~$ kubectl get pods -o wide
NAME READY STATUS RESTARTS AGE IP NODE NOMINATED NODE READ
demo-6c54f77c95-5bv2j 1/1 Running 0 9m44s 10.244.0.3 minikube <none> <none>
demo-6c54f77c95-9w9up 1/1 Running 0 11s 10.244.0.5 minikube <none> <none>
demo-6c54f77c95-b4bx5 1/1 Running 0 9m44s 10.244.1.2 minikube-m02 <none> <none>
demo-6c54f77c95-g2718 1/1 Running 0 9m44s 10.244.1.3 minikube-m02 <none> <none>
demo-6c54f77c95-gj5dx 1/1 Terminating 0 9m44s 10.244.2.3 minikube-m03 <none> <none>
demo-6c54f77c95-nlhwg 1/1 Terminating 0 9m44s 10.244.2.2 minikube-m03 <none> <none>
demo-6c54f77c95-svm5d 1/1 Running 0 11s 10.244.1.4 minikube-m02 <none> <none>
demo-6c54f77c95-vl45o 1/1 Running 0 9m44s 10.244.0.4 minikube <none> <none>

chpc@LAPTOP-TDCSOWDA:~$
```

A small video call inset in the bottom right corner shows a man speaking.

We initially had two pods, ok running on each of these nodes, we had 3 nodes. Then, we simulated a crash by stopping one of the nodes, and then we told that, ok there are only two nodes available. So, now, what has happened? If you see minikube-m03 node, the pod which was running here, ok has terminated, this pod also has terminated.

And if you see this has been shifted, ok to one particular pod to this particular minikube node and another one to minikube-m02 node. So, if you see this way, there are again 6 pods running, but those 6 pods are actually split or distributed among only two nodes, which is the master node and one worker node, because another worker mode actually is crashed or its not available.

Now, in this situation, since the master node is also running a worker node sorry a pod, it is not necessary that the pod could be scheduled on the master node. Generally, it does not happen, people do not try to do that or when the Kubernetes cluster is set up it is not done, it is not advisable. But here we are just showing there is a option of no schedule. So, if you do not schedule it, it goes to some other node instead of the master, right.

So, this is what actually we had thought of talking today. And in the next session, we will link this particular concept of scheduling, orchestration, ok to something which is

very specific to AI workloads which is DeepOps. We will see how we can see the dashboards, ok and try to see the Kubernetes dashboard, try to see how the visualization happens of the CPU, the memory, various pods, ok all of this in the next session.

So, I suppose today is this thing we are done.