Advanced Computer Networks
Professor Sameer Kulkarni
Department of Computer Science Engineering
Indian Institute of Technology, Gandhinagar
Lecture 49
Data Center Networking - Introduction Part 2

(Refer Slide Time: 00:17)

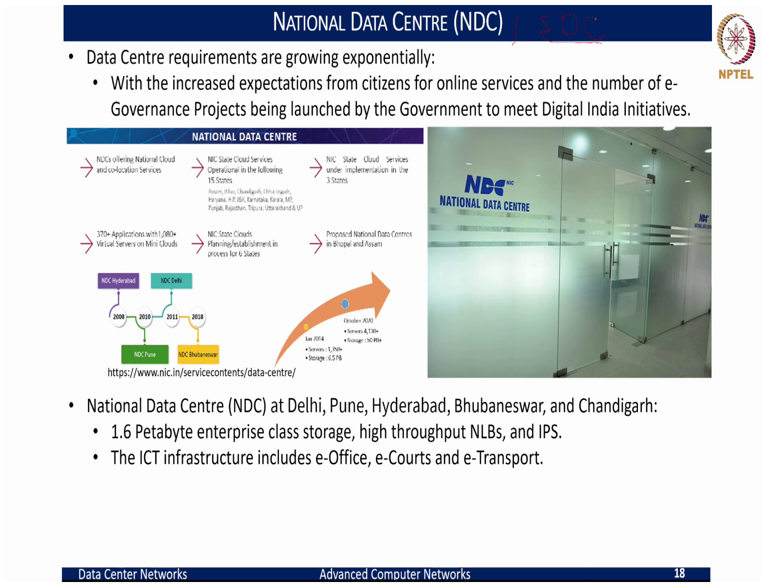

Let us now try to understand some physical characteristics and familiarize ourselves with some of the technical aspects and jargon that is used with these data centers.

(Refer Slide Time: 00:29)

So, first, let us ask ourselves, like how big are these data centers? And if we see the published details from various of the major companies like Facebook, Google, HP, IBM, and Microsoft about their data centers. Note that this is just the information for one of the data centers that you host as the host, several of the data centers across the US and worldwide. So, if we see the Facebook data center at Santa Clara, it has a square foot of 86,000 square foot space on which that Facebook data center is spread.

Likewise, Google in South Carolina has around 200 square feet of space, the infrastructure or a building encompasses this square feet area. And likewise, for HP is around 200 square feet area. And if you see the Microsoft Chicago, one that we saw earlier in the picture as well, the white building is around 700,000 square feet of area. And these are really very huge. So, if we have to put this in perspective, like what these numbers are.

And if we see in India, the recent match that happened the cricket match at the Ahmedabad, Motera Stadium, that is around 240,000 square feet area. And now put this in perspective, like when we say the data center of HP is 200,000 square feet, it is as big as this Motera Stadium or if you are a football fan, if you see the Wembley Stadium, that is like 172,000 square feet. And these data centers that we are speaking of are, in fact, much more bigger.

And also, you can see that some of them are big, some of them small, and it is not that all the data centers would be in the same size or same area. So, we have to understand how these data centers could be classified and what are the means to say, like, what are the required infrastructure requirements and, what are the computational requirements, and how these data centers can be put or classified.

So, for this, we can try to understand and say based on the amount of area that they will need, we may want to classify these data centers.

(Refer Slide Time: 02:44)

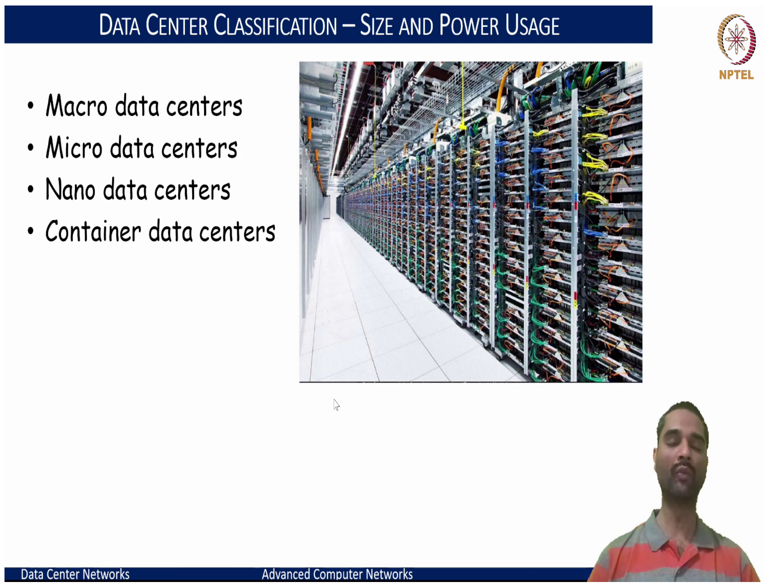

So, if I say that there are data centers, which are going to host so many of these servers, one, I can try to classify them based on the area requirement, and this area automatically translates to how many of the compute servers or storage servers or the compute nodes that I would put together the combination of compute and storage that I would be able to host in a given space.

And accordingly, the classifications were done in terms of their size. And the size also translates to how much of the power you would need to run them. So, one of the most commonly used classifications is based on the size and power usage requirements of the nodes that are hosted. And accordingly, we classify them as macro data centers, micro data centers, and containerized data centers, and we will also speak about these Nano data centers, which are typical data centers to think of, but as a means that were developed to ensure that we can have sustainability data center ecosystem in a distributed world.

(Refer Slide Time: 03:44)

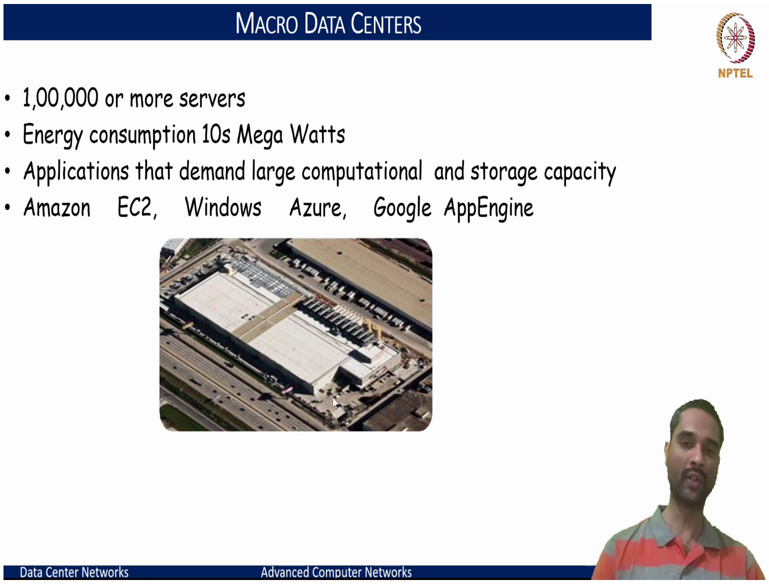

So, what do we mean by macro data centers? It is basically a space where you are hosting hundreds of thousands of these servers. So, think of servers ranging from one lakh to more than a lakh of servers put together in one place; you would call that as a macro data center. And to power up such a macro data center, you would need 10s to 100s of Megawatts power to drive these servers in that node.

And typically, this is where like most of the widely used critical applications on the internet, like Google, Windows, and Amazon, host such large data centers, and these data centers where you can see are much bigger than the biggest of the stadiums that we have, as we saw earlier. And these specifically also demand very large computation and storage capabilities. Any applications that cater to these would typically be hosted in these macro data centers.

(Refer Slide Time: 04:48)

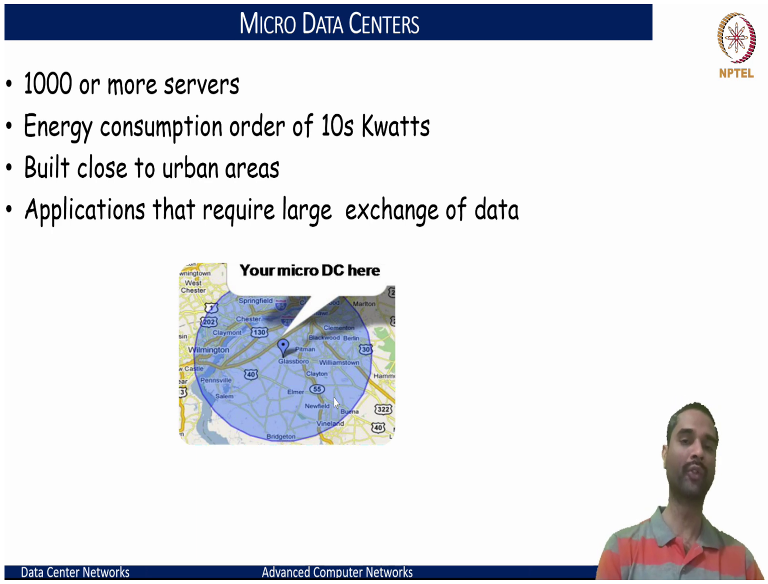

And the second of these are the micro data centers. And think of these micro data centers as a place where you would have thousands or more servers but less than a hundred thousand of them. And in essence, the power that you would need to run them ranges somewhere in 10s to 100s of kilowatts of power that you would need to run them.

And these are specifically built so that they can be close to the urban areas where you would have a lot of data that you would want to exchange and minimize or localize the traffic that you would want to exchange or to place rather than spread them across the globe. And that is the concept of having lots of micro data centers.

So, if you think of all your telecom operators they would want to operate, they would host their data centers to manage their backplane to ensure the connectivity that you would have. And now, when we think of 5G, a lot of Edge data centers will also come with it basically in the scale of micro data centers that are going to be hosted closer to the user than at one place that is far off from the user.

And also, the emergence of like the State data centers that I mentioned earlier is the means to build these micro data centers, wherein they will you will replicate these servers across the state on small scales, while the national data centers are going to be at a macro scale that is going to be put, hosting lots of those servers.

(Refer Slide Time: 06:17)

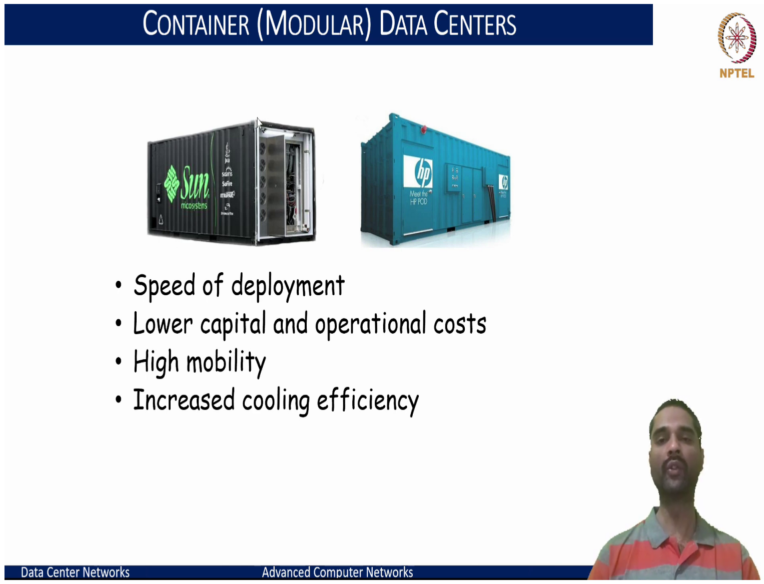

And the third of the model is what we call as a containerized or modular data center. And what these really look like is exactly the containers that we are showing here. Like any container that you see on the ships, it is exactly of the same size, and you are able to put in the computation and storage requirements that will be needed in such a container.

And typically, these are often used as a means to meet the requirements of the small-scale enterprises on the campus networks that you will want to run, and facilitate the local computation and storage requirements. And very importantly, compared to the micro or the macro data centers, the important aspect that modular data centers bring is the mobility that you can move them on the go, and as we ship the base, we can ship these containers and then have them made accessible at a local premise.

And also, bringing a lot of lower capital and operational expenses. And in fact, the entire containers can even be leased or managed by a third party for managing, and we could just lease a container to meet the requirements of the campus networks. And this is also because these are very small in number and hosted in a very small thing; the cooling efficiency is also very high because of much more flexibility in placing these containers at a place where we would want. And hence, these container or modular data centers are often seen for many of the enterprise networks, the startup companies that will want the campus networks that we want to grow, all of these typically apply these container data centers.

(Refer Slide Time: 08:02)

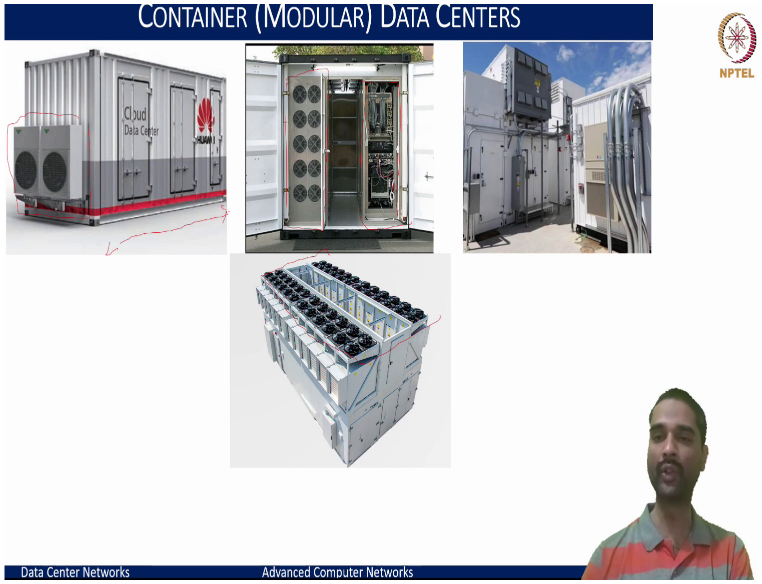

And they can be run by many of the different companies where you could really rent them on a need basis, as well as lease them over a period of time and deploy them. And to look in like this is how a container data center would look like where you have a container with a cooling on the external side and you could also have coolings on the top, as you see here, the cooling on the side, as you see here.

And even within a container, you can think of these are the racks that are being fit and a series of racks ranging around 10s to 20 of the racks that can fit in a container. And these containers themselves can be of different sizes; they can be small and large. And then you can place, and the powering also becomes a lot more easier so that you can plug kind of power, the hybrid mode of power that you would want to set up and use all of these flexibilities come with these containerized data centers.

(Refer Slide Time: 09:02)

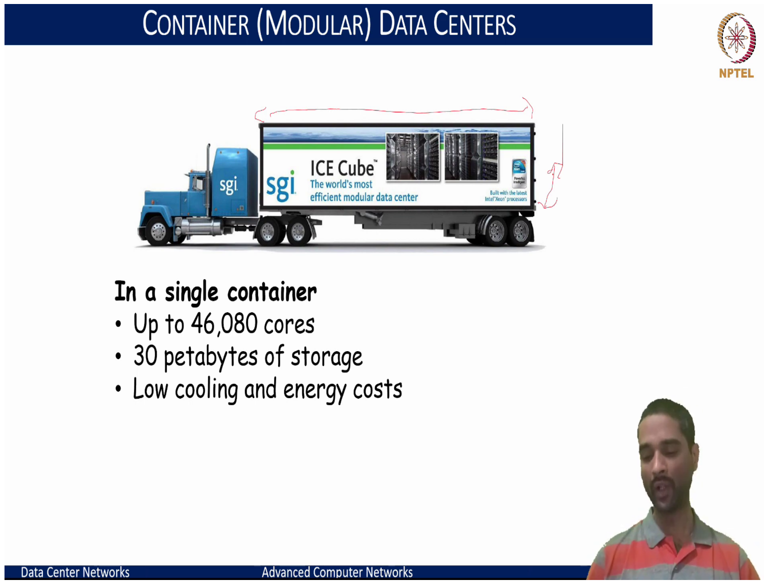

And such a single data center again, if you can consider this as a very large truck. And it can host almost around 50,000s of the cores. Or if you take racks maybe around 30 to 50 of the racks, which can basically host each 50 to 100 racks with a couple of aisles that you can set up. So, you are able to host several of these computers until you could also include around 30 to 50 petabytes of storage that can come within these containerized data centers.

But when all of this, what we end up seeing is there is very high power and very high infrastructural requirements that you see. And especially when you think of power for each of these ranging in several kilowatts, then the question comes, how can we minimize this, especially when we want to lower the carbon footprint, we would want to minimize the power necessary or the energy that you would need to run these data centers.

(Refer Slide Time: 10:12)

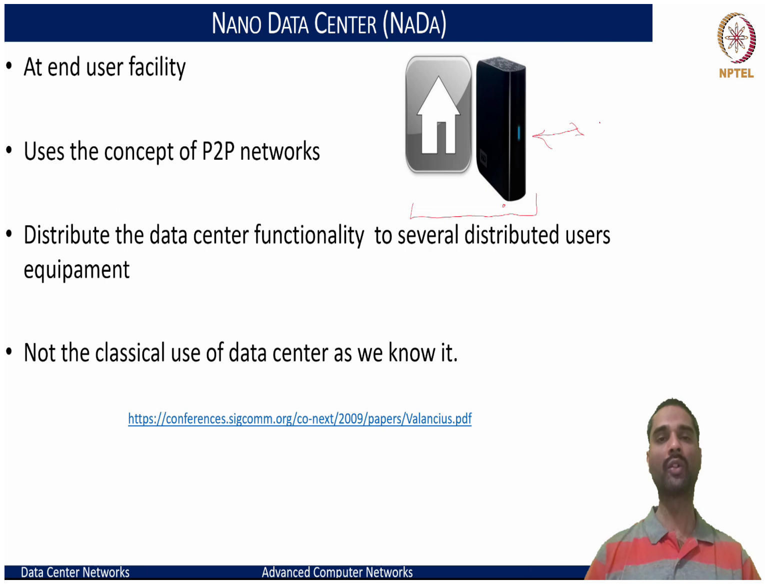

And researchers put together some thoughts and try to come up with what we call as these Nano data centers, or NADA. And I put reference to this paper was published in Co-next in 2009, that brought the concept of saying when a user connects to the internet, he is connected to his gateway device or a WiFi router. And every user has custom needs in terms of why you hook up to the internet and what kind of services you would use. In essence, if I am able to localize my requirements, I may not need all of the services that are going to be run everywhere, I will need a subset of them that are being run at my place. And you can bring those services much more closer to the users.

And now, if you scale the users into a community like a campus, where students have specific requirements, faculty and staff have specific requirements. So, you could bring in those kinds of services to a closer vicinity and closer location for users to access them. And this is where the rise of what we see it in a Nano data centers as a thought experiment and as a means to build this infrastructure came up.

And it leverages the concepts of P2P networks, where you think of your gateways as a very Nano data center. So, your gateway can host specific services on its own. And it can basically also communicate with other gateway devices and the local, and you can have a hierarchy of these gateway devices that are being communicated to provide distributed functionalities to provide a means to distribute the functionality across multiple of these data centers and ensure that users can still connect to them.

And this way, you will have, basically, several of the distributed users having their own equipment acting as a peer connected with every other peer's network to facilitate a large data center. And still, the power requirements for each of these will be within the typical power of wattage of the home networks. And that is less than 100 watts, 40 watts to 60 watts is what you typically use and ensure that now you are able to facilitate the data center services to the greater mass with less power constraints. But, like, it goes like this is a good thought experiment. But making them feasible is a lot more challenging itself. But nonetheless, this is one of the things that I wanted to introduce to say how things can really be taken and addressed for different aspects.

(Refer Slide Time: 12:56)

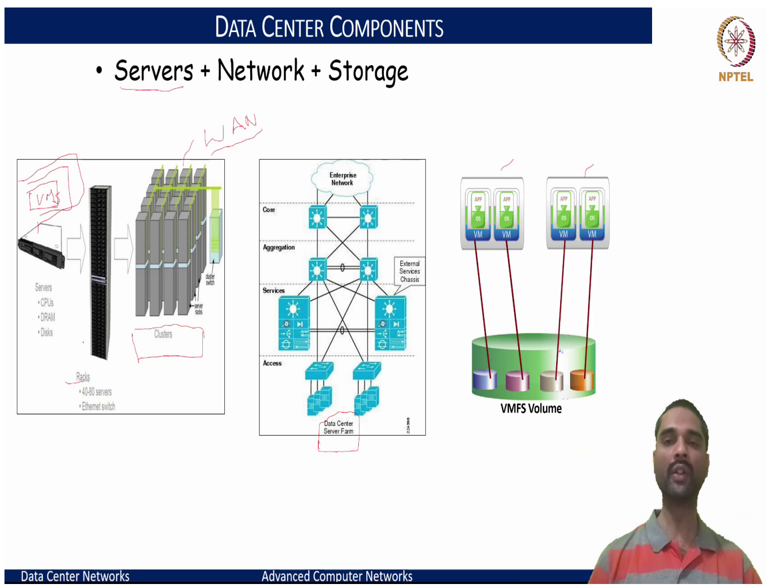

So, besides this classification, what we can see a generality in all of them is that these data centers, in essence, constitute servers and storage, and servers when we mean compute servers and storage as we mean storage devices, and both combined together with the network. So, the key components, if we have to think of in a data center, would be the storage servers and network. And the servers would basically encompass the number of cores or the CPUs that you would want. Also, local storage in terms of the disks and the memory, the data that you would run. And all of this put in the form of a rack. And these rack servers are all going to be connected using data center switches.

And a typical enterprise network would also have lots of these racks that are being placed. And all of these rack servers will also be interconnected. And that is what we see as a cluster of these interconnected racks. And these clustered interconnected racks also need to be connected. If you have campuses across different locations, they may have to be connected through a WAN and connected through the internet if you have to facilitate internet connectivity.

So, it is going to be public as well as private wide area networks to which you may be able to connect and access these services. And when we want to build such an infrastructure, we need to have a specific architecture with which we would want to build. And that is what we call as a Server Farm and architecture for Server Farms, we will look into that in a bit, but this is typically the model that is being used for building these enterprise networks. And as we saw the growth of virtualization, what you would also see is not just a physical server used as a server as it is. The server may host many of the VMs that could be running on the same physical server and be able to facilitate the required compute and storage requirements.

So, we could have compute VMs, as well as storage service that is also being virtualized to provide virtualized storage, and all of these being interconnected using the network. So, if we have to understand this, we need to understand what the servers are, what the racks are, and what is the way in which these data centers are going to be built.

(Refer Slide Time: 15:28)

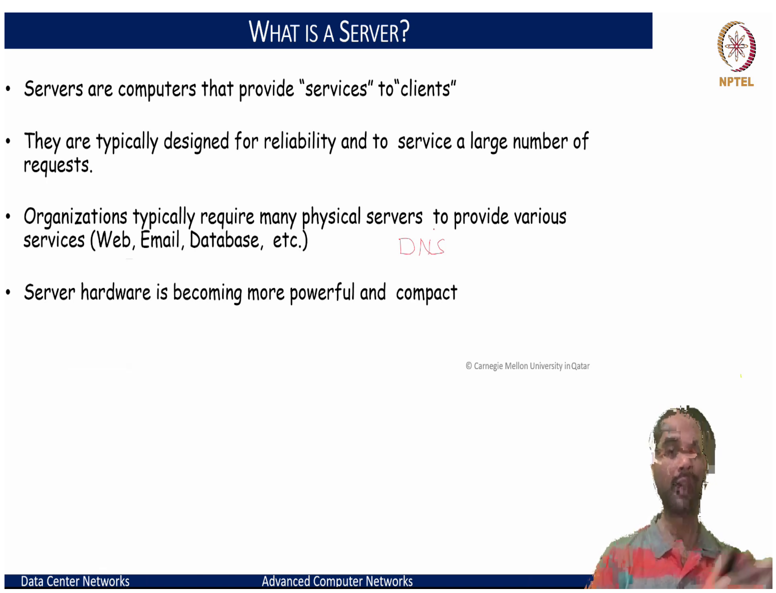

So, think of a server as a high-end computer that provides services to the clients. And because we have to have a lot of clients to be provided services at the same time, they have to be scalable in terms of the services that they can run. And these are the compute devices that are like we expect the services to be run 24x7.

And reliability has to be topmost in terms of ensuring that they are available for a majority of the time, and are able to run continuously. And any organization typically, if we think of all the servers, we will see the web servers, the email servers, database servers. And if we think of the internet, the most crucial is the DNS servers, without which the address translation and mapping of which IP to connect to would not happen.

So, all of these are the kinds of typical services that are going to be run as servers. And the server hardware, as we saw, has evolved from having around 8 to 32 cores nowadays, all the way supporting around 128 cores, hosting around several gigabytes, hundreds to thousands of gigabytes of onboard memory, and up to several terabytes of storage being supported locally.

(Refer Slide Time: 16:49)

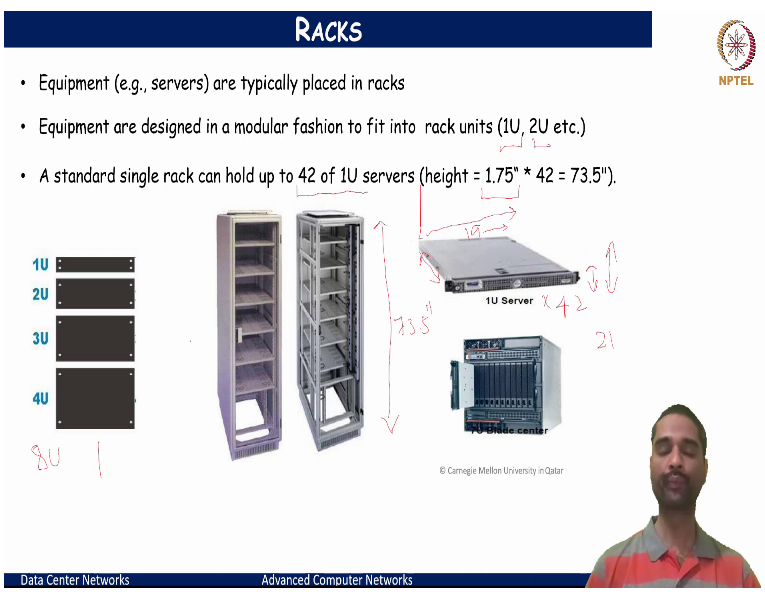

And these servers when we spoke, we said we would typically want to use these rack servers. And when you say these racks, we have to have a standard because if there are multiple vendors, if they do not follow the same standard; they may not be able to mix and match different servers.

And the standardization was also done in terms of what a rack should be in terms of the dimensions and how a rack should be aggregated and set up so that we can have multiple rack servers that can be hosted within a chassis that we can see of this model. And these racks allow us to build this stack of these servers in a modular fashion.

And the standards typically what we use as a 1U or 2U, one rack unit, or two rack units wherein each unit corresponds to 1.75 inches of height. And this dimension of 1.75", and that followed by around 19 inches on the sides, would allow us to define what we call as a one-rack unit that can be hosted in a rack.

And a typical Compute Engines or storage engines that we get may be of different sizes in terms of the change in terms of the height unit. And this could change from 1U to 2U to 3U, 4U, and we may even have 8U units of devices that can be placed within these racks. So, when I refer to a device as a 1U, we are referring to the height of it being 1.75 inches.

And if it is 2U, then it is going to be the double of this 1.75" and likewise. So, we will see multiple of the devices come in these form factors. And when we make a rack, even the racks could come in different heights. Typical most standard rack that we typically use is a rack which can host 42 of the 1U servers or a height of 73.5 inches and this is able to host 42 of these 1U servers. And if there are 2U servers, then 21 of those, and you can have a mix and match of 1U, 2U, 3U, and 4U servers that you would be hosting within a container or within this rack-enclosed infrastructure that is built.

(Refer Slide Time: 19:34)

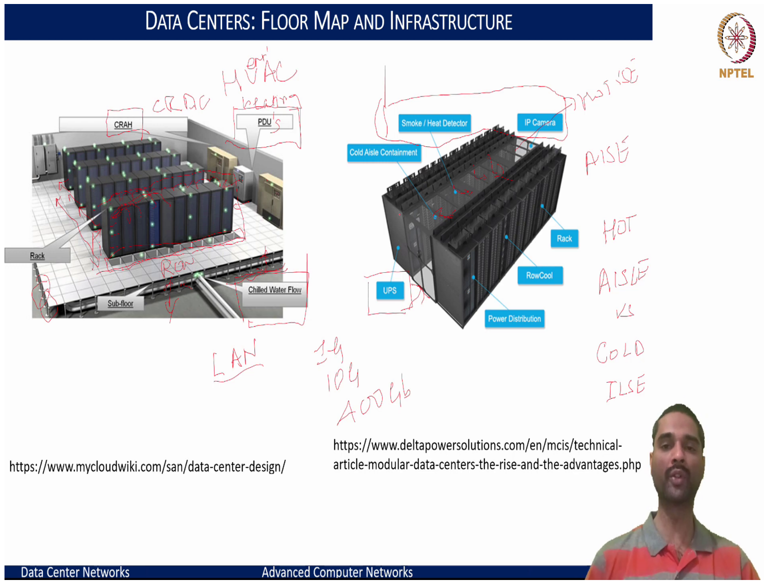

And these racks are again going to be placed and a series or what we call these as rows of racks. And these racks must be at the fundamental units of scale that we would want to replicate. So, all of these are basically the racks that are being placed in a row in a series. And these racks become a fundamental unit of scale that we can deploy.

And if we look at a cross-section of how these data centers floor maps would look, what are all the key components that would be built within a data center, you will see that they require a typical floor map for the plan and a typical set of components to be setup for building such data centers. And the base unit is a rack that is going to be scaled and put up, and as we scale and set them up, this is called a row of racks that are being set up together in this row depends on the size of the infrastructure where it is going to be hosted like in if it is a typical campus network, maybe 5 to 10 of those would be hosted together to form a row or if it is like the data centers, like we saw us, Microsoft or Google, there may be hundreds of them that are being put in a single row.

And the flooring for all of these data centers also has unique characteristics, and typically what we call as these raised floors or you have a Sub-floor on top of which these data center racks are being placed. And this is an artificial floor that is being made so that you can ensure that can be removed and moved as and how they are required. So, you can notice that these are artificially being placed over where the bars are being raised.

And you can put this makes the movement of these data center services, the entire racks from one end to the other, a lot more easier. And most importantly, this is done to ensure that there can be the circulation of air that can be done below so that you can cool them much more effectively. And these organizations have a floor that is lit up and has a row of servers, and having the cooling can be placed right in the vicinity so that you are able to cool these servers.

And you can see now that we are also setting up these rows of servers adjacent to each other with some partitioning or space that is being set up in between. This is necessary if we have to go manage each of the devices so that an IT personnel can go wire up the things, look for changing of any of the equipments, and any updates or things that you would want to do; all require a person to go update and manage these devices.

So, you need a space where the IT operators, network operators can go and look up for these devices and do specific operations. And also like these devices, each of these maybe we would want to power up from different sources like a typical enterprise network would have at least two kinds of different power units on which you will be able to power up these racks.

So, each rack, in a way, would have a power distribution model that is necessary to set up this, and power is the most important component to ensure that these devices can run without failure. And it is not just for the servers that you would need the power if I have the cooling system, the air conditioners, their coolers, their filters, all of these also require the power, and hence, this power is a very crucial aspect in ordering in running these data centers.

And typically, this is managed by what we call as the Power Distribution Units. And like I said, there may be multiple of the PDUs that are being set to facilitate or powering up this infrastructure. And the next critical component is cooling. Like we said like our computers dissipate a lot of heat, and if the heat is not being taken out, the efficiency of the devices goes down when we are talking of thousands of servers that are being placed, these are going to have a lot of heat that is going to be generated and we need the mechanism to cool them up. And this is where the infrastructure what we call as CRAH or even the CRAC is being set up. And what the CRAH means is it is basically a Computer Room Air Handler or CRAC we mean Computer Room Air Conditioners and these units are responsible for ensuring that whatever the heat that is being produced by the equipments is dissipated out and ensure that the room air temperatures are being maintained.

And as a facility to this, what we typically use is also the HVAC is to ensure the ventilation, the air conditioning and heat management for all of these devices to be managed. So, HVAC as it is it is the heating and this is ventilation, and the air conditioning systems are what constitutes to how we can build the infrastructure for cooling up these devices.

So, you could also have chilled water or chilled water flow, which is basically a circulation of the chilled water to ensure that you can remove the heats. And typically, the CRAH would regulate the temperatures for these chilled water for air circulation, all of this to ensure that you are able to run these devices in an efficient fashion.

And likewise, we also when we have these connects, the other most important aspect is interconnecting them. And this is where structured cabling becomes very essential in terms of how we connect the servers within each of these racks. And how we connect the servers that are adjacent within a row and how we connect the row of servers with each other because all like we say, how to work together and there may be a means wherein some have to be connected in some LAN some maybe on another adjacent LANs, how we partition them, how we build this mechanism of these communications becomes a very critical aspect.

And the amount of infrastructure that you would need to cable these is also massive. So, we cannot cut corners when it comes to interconnecting them because that is pivotal for providing good service. And also, the interconnection infrastructure could also be scaled out as well as be adapted and replaced as you need for higher when we transition from a 1 gig-ethernet to 10 gig.

Now, we are almost on 400 gigabit-ethernet that you want to plug. So, all of these aspects become important to be considered. And when we have such a rack of series of rows of servers, we typically call the intermediate as an aisle that is being set up for management purposes and also for air ventilation.

And accordingly, we will see that if the computers are on the one side, where the heat is being generated, we may have what we call as a hot aisle versus what the other side, where you will see the front end of the devices as the cold aisle. So, if you have a device where continuously their cooling fans are being run and this is where the hot air is going to be generated, making this hot aisle, if these ends are continuously being cold, this is going to be like a cold aisle and likewise. Besides this kind of infrastructure, we would also see the need for several of the mechanisms like IP cameras, smoke and heat detectors, and likewise.

So, what this means is building a data center does not just involve taking into account the compute and storage requirements; we have to map that back to what is the corresponding floor map, what is the means to build the power for them to the what we call the uninterrupted power supplies, what kinds of PDU is that you will need to facilitate the uninterrupted power supply, what is the kind of cooling system that you will need to ensure that these devices run in an efficient manner.

(Refer Slide Time: 28:33)

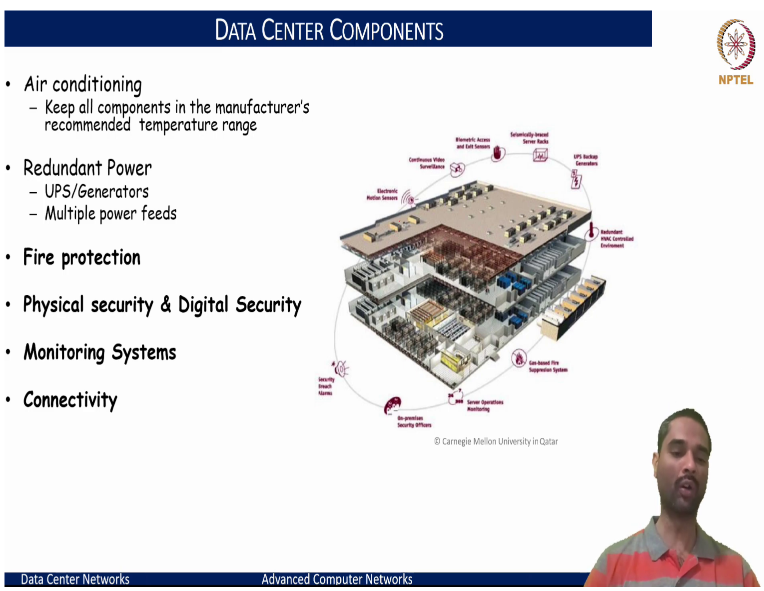

And also, because these are very expensive and very critical infrastructures, you would want to build the monitoring mechanisms to ensure that you are able to safeguard and ensure the security both physical as well as digital security for these data center infrastructures. And also, because you are having a lot of heat that is going to be generated, you also want to build a monitor for these for fire protection, ensuring that there is a means to protect you in the case of small fires. And that is where a need for better monitoring systems, including the digital surveillance becomes essential.

So, we have so far tried to cover about what are the core data center networks, the types, the classifications the basics in terms of what constitutes a data center from the outside in terms of non-technical aspects, and we have familiarized ourselves with respect to the key jargons of technical terms that we may be using later on. Next, we will try to look into the core technical challenges with respect to data center networking and what are the aspects that we need to be focusing on.