We have seen that virtualization provides us with opportunities to reduce computing costs, and this has encouraged the move to the cloud. In this section, I want you to consider one of the often-overlooked questions about the cloud, which is, "Who actually owns the data in the cloud?" IT hardware is expensive and complex, and most customers don't like complex. With the global scale of the cloud, we as a cloud service provider offer **Azure** services, or **Office 365**, that behave just like a utility company.

Here is another metaphor to comprehend the logistics of CLOUD Services. For anyone who has experienced a cruise liner vacation, imagine "abstraction" as the difference between the passenger and crew on a cruise ship. For example, there are 2,695 passengers on the Luxury Liner and 1253 crew, 18 decks, and four of these are dedicated to the crew. As a passenger, they are abstracted from most of the hidden workings of the ship and the crew. It all goes on behind the scenes. The cloud provider is the same. They are the cruise ship and crew, and you are the passenger.

When you subscribe to a cloud service, you can only manage the externally accessible parts that are directly made available to you. Below the surface, the cloud provider manages the hardware, the environment, security, billing, maintenance, and everything else. If you review the contracts of Office 365, Google, and Amazon Web Services, each of them is clear that you own your data and retain all rights : as a CLOUD SERVICE PROVIDER, we are strictly in the role of enabling storage and data access.

Computing is now a service. Something that you buy when you need it. In fact, everything is becoming a service, or service-orientated. For a long time, businesses have stated that they get frustrated by the complexity of IT, and just want it to work. Take a greengrocer, or a supermarket. We don't care about the logistics that go on behind the scenes to get the produce to the store: we just want the fruits and vegetables. The main services currently offered by the cloud providers include services such as communication as a service, compute as a service, infrastructure as a service, and network as a service.

Communications as a Service includes Voice over IP, such a Skype for Business, and secure VPNs (A **v**irtual **p**rivate **n**etwork extends a private network across a public network, and enables users to send and receive data across shared or public networks as if their computing devices were directly connected to the private network) that connect remote users. These services are possible without the customer needing to buy dedicated expensive hardware. With Computing as a Service customers can rent virtual servers to host their own workloads with an unlimited number of cores and RAM that is available on demand. Infrastructure as a Service allows you to rent virtualized infrastructure located within the datacenter to run your own cloud applications and servers. You rent the servers, routers, switches, storage, and firewalls. With Networking as a Service, you can use the backbone networking and firewalls and configure the subnets, private and public IP addresses, and all the routing and remote connectivity

that you need. The last three cloud services include platform as a service, software as a service, and monitoring as a service. Platform as a Service includes all the middleware tools that you need. You rent the preconfigured platform, which provides you with the tools to build your own applications on.

Examples of PAS include Azure and SQL Services. Software as a Service is where you consume the finished product or software and have no access to the underlying hardware or customization. Office 365 and SharePoint are examples of Software as a Service, or SAS. Monitoring as a Service allows you to geolocate monitoring tools around the globe that are used to monitor your cloud services and applications from different continents and keep you aware of the activity and any problems.

If we wanted to categorize cloud services in a different way, we could look at the type of services that organizations may receive from the cloud. Some of these would then be compute services, such as Azure or Amazon Web Services. Storage services, such as Azure Storage, or OneDrive for Business. Productivity services, such as Office 365, which allows users to collaborate, create, and share documents. Finally, search services, which can embed in-depth search functionality into data stores, such as databases and internet sites using Azure Search.

One of the primary reasons for using the Cloud is that the cost should be significantly lower when compared to traditional computing. In this section we'll introduce the various Cloud payment model options available. In a non-Cloud environment business typically incur a wide variety of on-premise IT infrastructure costs, including physical server hardware, associated storage, networking and onsite and offsite backup and archive. Each of these items is a real fixed cost. As their business grows they will need to increase their capacity by purchasing more of these assets.

There is also the cost relating to ensuring business continuity and disaster recovery, such as perhaps a second standby site, physical security and monitoring, and software licenses for all devices and servers. Finally, often the largest cost within the IT budget will be the staff costs. Some of these costs may fluctuate, such as licensing and staff costs, though most of these are fixed and discrete cost items that require planning and budget allocation. In business there are two types of funding model's available capital expenditure, which is often referred to as CapEx, and operational expenditure, which is known as OpEx.

We can see from the definition that a CapEx is when money is spent on fixed assets. This can include land, buildings, and equipment. Typically, a business will purchase IT equipment such as servers, the racks, and storage. This cost immediately leaves the bank account when the item is purchased, but the cost can be shown as an amortization or depreciation charge in the accounts of the business for several years. OpEx is the money spent on operational costs, such as wages, rent, utilities, and repairs.

Most businesses like to predict and manage their monthly OpEx cost closely so they can create accurate forecasts and budgets. If you lease an asset such as a new server this is paid by a monthly lease payment which is accounted for as OpEx. A new and very popular model of financing is called Pay-as-You-Go. Everything is becoming subscription based. You're probably familiar with the Pay-as-You-Go billing model from your cell phone of TV entertainment provider. You pay for the service and for what you consume. Pay-as-You-Go cloud computing is the same principle as we spin up more processing resources and consume more data bandwidth our monthly cost increases.

There is no commitment and you only pay for the usage of the Cloud each month. Pay-as-You-Go is then an OpEx type of expenditure. We already know that accountants love OpEx since it is predictable. In the next movie we will cover how the Cloud is able to expand capacity and scale to our needs and provide regeneracy and great availability across the globe.

You're probably looking to understand how using the cloud enables businesses to scale out their cloud presence and protect against outages. In this video, I want you to gain the awareness and understanding of the way in which the cloud allows you to expand capacity and consumption and ensure that your assets in the cloud are redundant and always available. One of my earlier clients is a gift marketing company, and they spent most of their year gearing up for the holiday period. During the months from November and December, they made over 80% of their money. Let's look at how they optimized their computing requirements. Imagine if they had to buy and then maintain service and rent data center space for the whole year. That's very expensive, since they would have a lot of underutilized service for most of the year. A better solution is to purchase a base level of computing resources and then rent the additional capacity for the busiest periods. This hybrid approach provided the business with rapid scalability and better economics by using the cloud. So that you can estimate the cloud usage and therefore cost, you need to perform some analysis.

This includes looking at the current demand on your resources. Capacity analysis provides you with the threshold limits that you expect for current and future utilization. You will need to track the actual performance utilized by application services such as the processes, the bandwidth, storage, and networking. This data will help you determine what level of cloud resources you'll start with. Scalability relates to the responsiveness of the system, how much work it can handle. Consider a horse pulling a carriage. Most of the time, the horse is walking, but it can speed up to a trot without too much delay.

How fast can the cloud resources react to the workload? You configure the cloud service to be elastic, meaning that once your current resources are at full capacity, additional resources can be added dynamically by the cloud provider. This time, we're adding more horses, which gives us more pulling power, scaling out the capacity. Our data center can be copied to multiple separate storage locations within a data center,

and stored in different data centers geo-located across the globe. Very few businesses could achieve geo-redundancy without the cloud because of the huge costs involved.

Enterprises require reliable IT systems and spend a lot of money to prevent downtime. When a system is reliable and protected against failure, it is called highly available. The amount of uptime that a cloud provider will advertise for a service such as Office 365 is usually shown as a percentage of the total uptime as shown in the table as several nines. As the availability percentage increases, the cost of providing the service also increases dramatically.

Most cloud providers will deliver a service level agreement which specifies an official commitment to the levels of service that they will strive for. If you're migrating mission-critical services to the cloud, you should check the SLA carefully. In the next movie, we'll cover how to compare cloud services with on-premise services.

You rent the servers, routers, switches, storage, and firewalls. With Networking as a Service, you can use the backbone networking and firewalls and configure the subnets, private and public IP addresses, and all the routing and remote connectivity that you need. The last three cloud services include platform as a service, software as a service, and monitoring as a service. Platform as a Service includes all the middleware tools that you need. You rent the preconfigured platform, which provides you with the tools to build your own applications on.

Examples of PAS include Azure and SQL Services. Software as a Service is where you consume the finished product or software, and have no access to the underlying hardware or customization. Office 365 and SharePoint are examples of Software as a Service, or SAS. Monitoring as a Service allows you to geolocate monitoring tools around the globe that are used to monitor your cloud services and applications from different continents and keep you aware of the activity and any problems.

If we wanted to categorize cloud services in a different way, we could look at the type of services that organizations may receive from the cloud. Some of these would then be compute services, such as Azure or Amazon Web Services. Storage services, such as Azure Storage, or OneDrive for Business. Productivity services, such as Office 365, which allows users to collaborate, create, and share documents. Finally, search services, which can embed in-depth search functionality into data stores, such as databases and internet sites using Azure Search.

Finally, for businesses, Windows 10 is now being offered to enterprises as a monthly subscription. In the next movie we will cover the cloud payment models that are available.