27. Cloud Computing

Apparently, this topic is only part of the separate 200-105 exam objectives and not included in the composite (CCENT+CCNA) 200-125 objectives.

SERVERS

Physical Characteristics of a Server

- No KVM (Keyboard, Video, Mouse)—Servers live a lonely existence in cold, noisy, and windy data centers. Often lacking even light, they have evolved without webcams
- Rack mountable—19 inches wide, and some multiple of 1.75 inches (1U) high.

UCS (Unified Computing System)—Often nothing more than a re-branded Hewlett-Packard server, these are sold and, more importantly, supported by Cisco. "Unified" is just a marketing buzzword that Cisco applies to everything from computers to telephones.

Physical servers are often networked using a combination of ToR (Top of Rack) access switches in each rack cabinet (typically 6 feet tall and holding 42U of equipment) and EoR (End of Row) distribution switches. It is routine for each server to have more than one NIC.

VIRTUALIZATION

VM (Virtual Machine)—One Operating System (OS) environment for software to run in. Several can run on a single physical computer (host).

Hypervisor—A special-purpose OS for virtualization that runs directly on the physical host and allocates resources (RAM, disk, CPU cores) to the VMs. Normal operating systems run within the individual VMs and are fooled into thinking they have an entire computer at their disposal.

Each VM is given at least one virtual Network Interface Card (vNIC) to which it can assign an IP address. The concept of a virtual switch can be used to handle traffic to, from, and between multiple VMs on a single physical host.

CLOUD SERVICES

Cloud services are a highly dynamic use of virtualization. Special Publication 800-146 of the US NIST (National Institute of Standards and Technology) defines 5 essential characteristics:

- On-Demand Self-Serve—A client can use a web interface to create and pay for a VM
- Broad Network Access—VMs are available over a network and accessed in an ordinary way that lends itself to varied clients (computers, cell phones, tablets, etc.)
- Resource Pooling—Actual servers become part of a pile of resources from which VMs are created, without concern for which customers end up on which hardware
- Rapid Elasticity—requests are filled quickly from a seemingly endless supply. This differentiates cloud services from the age-old system of researchers booking time on a supercomputer next week
- Measured Service—resource usage is measurable and billable

The big concept is that the customer doesn't buy anything tangible that they get to keep. Everything is sold as a service. The same NIST document defines 3 service models:

- Infrastructure as a Service (IaaS)—In the web interface shown below (Alibaba Cloud), pricing is based on the number of CPU cores as well as the amount of RAM, disk, and internet traffic. Changing the operating system from Linux to Windows can double the price on the low end, while asking for the Mainland China region can quadruple it.

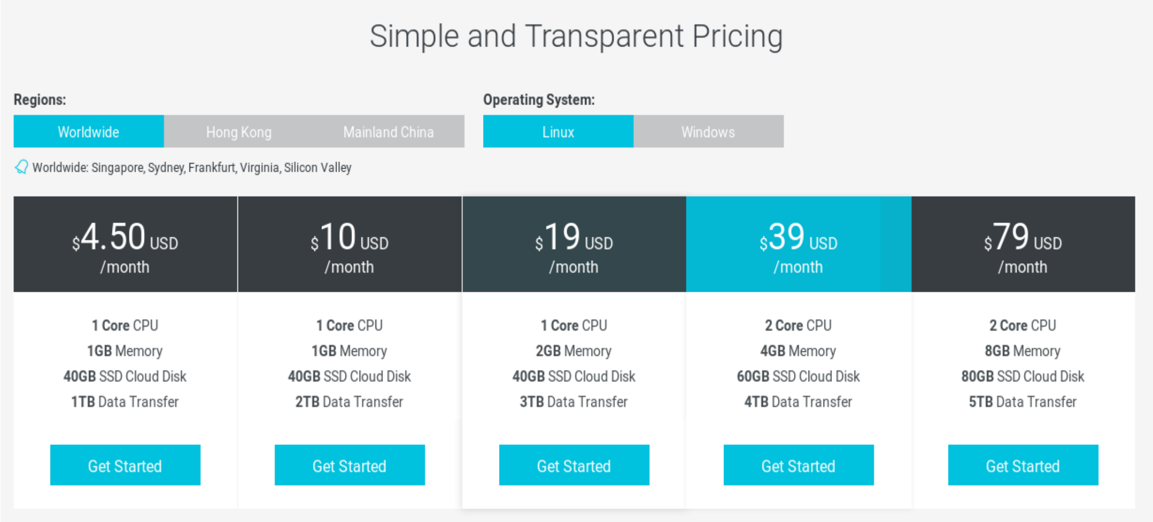

Simple and Transparent Pricing

Regions: Worldwide

Operating System: Linux

Worldwide: Singapore, Sydney, Frankfurt, Virginia, Silicon Valley

$4.50 USD /month	$10 USD /month	$19 USD /month	$39 USD /month	$79 USD /month
1 Core CPU	1 Core CPU	1 Core CPU	2 Core CPU	2 Core CPU
1GB Memory	1GB Memory	2GB Memory	4GB Memory	8GB Memory
40GB SSD Cloud Disk	40GB SSD Cloud Disk	40GB SSD Cloud Disk	60GB SSD Cloud Disk	80GB SSD Cloud Disk
1TB Data Transfer	2TB Data Transfer	3TB Data Transfer	4TB Data Transfer	5TB Data Transfer
Get Started	Get Started	Get Started	Get Started	Get Started

- Software as a Service (SaaS)—The customer is paying for the use of a software package, e.g. an e-mail server. The infrastructure (CPU, disk, network) required to make that happen is incidental and not directly specified or billed
- Platform as a Service (PaaS)—This is a higher level of abstraction than IaaS in that a customer is buying a platform to deploy applications without regard for or control over infrastructure details, like RAM, disk, networking, etc. Mr. Odom uses the example of a software development platform, complete with a suite of software development tools running on the cloud. Such a service would allow a software developer to be added to a project without the expense and delay of buying and configuring a desktop computer for them with all the necessary software. The new developer would still need a desktop computer, but it could be much more generic, needing only to access the cloud, where it would find all of the special tools needed for the job

Private Cloud—By adding a "cloud services catalog" (like the above screenshot) and rapid provisioning (order fulfillment) to its own internal virtualized datacenter, a company can essentially have its own internal cloud services.

Public Cloud—The customer can interact with their VMs using the internet, a VPN or private WAN link.

Terminology—The cloud adds an extra layer between a company and its customers—a cloud provider. Mr. Odom copes with the semantics as follows

- Provider—The cloud company selling you VMs
- Consumer—"You," the company that is a consumer of cloud services.
- Users / Clients—People (perhaps even the public) using the services on the VMs

You can contract a traditional WAN connection to your cloud provider(s), simply use the internet, or use a VPN over the internet. Of course, you'll always have to plan sufficient capacity, even if you're just adding cloud traffic to your existing internet connection. For that, remember that parts of your infrastructure may remain internal but need to communicate with the cloud. For example, you might keep your authentication servers internal.

The real decision between a traditional WAN and the internet options becomes an issue of security, speed guarantees, how quickly you can start using the cloud, and how easily you can change cloud providers. The following is adapted from Mr. Odom's Table 27-2.

	Secure	QoS	Quick Start	Easy to Change Clouds
Internet			•	•
Internet + VPN	•		•	•
WAN	•	•		
Intercloud Exchange	•	•		•

Using the Internet as a WAN has several advantages.

- Convenient for starting on the cloud—no waiting for a WAN to be setup
- Easy to switch cloud providers or use more than one, for the same reason.
- Good for geographically diverse workforce, especially traveling workers

Complications when using the internet to reach cloud applications include

- Security—"Man in the Middle" attacks are possible
- Capacity—May need to upgrade your internet connection for extra traffic. This may actually be a positive because it puts the excess capacity from two uses (cloud and normal internet usage) in the same basket and usable when either runs short.
- No QoS (Quality of Service)—The user experience may suffer when the internet is busy
- No SLA (Service Level Agreement)—The internet tends to be provided "as is" by an ISP. Dedicated WANs are more likely to guarantee speed and availability

For a VPN to your cloud services, you'll configure your end of the VPN. For the cloud end, you can either choose a cloud provider who sets that up for you or configure it yourself on a virtual router as one of the things that you're running on the cloud.

A true WAN connection to the cloud provider is private and likely has performance (QoS) guarantees, but takes time to setup, incurs an ongoing expense, and has the remote end firmly anchored at that one cloud provider, leaving you less flexible. When setting up the WAN, you may want to choose where in the cloud provider's network (which city) you connect.

Intercloud Exchange—A third company that sells private cloud networking as a service to overcome the inflexibility of a WAN to your cloud provider.. Your company WAN connects to the exchange and it connects to many cloud providers, allowing you to change cloud providers without changing your WAN.

Cisco Intercloud Fabric—A family of products that assist migration from one cloud provider to another—not an intercloud exchange.

Traffic Patterns

When a company with branch offices replaces an in-house server with a cloud server or service, they can either have the branch traffic flow directly to the cloud from each branches or via HQ, as before, doubling the traffic at HQ as it comes in from the branch and back out to the cloud.

You might add direct internet connections to the branch offices in spite of security concerns.

Common Services

DNS—When a function referenced by domain name is migrated to the cloud, you'll need to update its DNS entry. If the name is within your domain (*function*.yourdomain.com), your normal DNS server will still need to take the initial query. You can either update it to give out the cloud provider's public IP address for that function or update your DNS server to punt the request to the cloud provider's DNS server. The latter provides the cloud provider extra flexibility. You still have the first step in the lookup, so none of this affects your ability to change cloud providers.

Addressing—Cloud providers are likely to use a private address space internally to number their VMs and use NAT for any that a consumer wants directly reachable from the internet. The consumer won't be configuring DHCP to accomplish any of this. If the consumer wants to provide cloud services only for the internal use of their own employees, it makes sense to use the consumer's own internal addresses (perhaps private) and advertise those subnets internally so that their traffic is routed across the VPN or WAN link to the cloud provider.

NTP—A consumer who has a large number of VMs at a cloud provider may wish to run NTP on one of their VMs (perhaps on a virtual router) as a client and server, to learn the time as a client and then serve it to all the other VMs.

VIRTUAL NETWORK SERVICES

SLB (Server Load Balancing)—balance load between identical servers.

- The pool of balanced VMs might be represented by a single public IP address
- Individual VMs being balanced would each have a private IP address
- Incoming connections might be NAT-translated to a single private address, that of a load-balancing VM, sold as a service, then balanced across the private IP addresses of the servers
- The router performing NAT and the load-balancing are both sold as services (even if bundled for pricing purposes) and are not directly configured by the consumer (you)

VIRTUAL NETWORK FUNCTIONS

Virtual Network Infrastructure—Implementing traditional routers, switches, and firewalls as software in VMs.

VNF (Virtual Network Function)—The resulting virtual router or switch that can be installed in your cloud. This makes the most sense when you're leasing multiple IaaS (raw computing and storage) VMs and need to connect them with more control than the cloud provider offers

NFV (Network Functions Virtualization)—A term used by service providers for how they "virtualize network functions inside their network"

Cisco Offerings

CSR (Cloud Services Router)—A router running Cisco IOS XE, implemented as a VM. Can be used to provide the cloud end of a VPN if you want control over both ends.

Nexus 1000v—a "virtual switch" sold by Cisco that you can install on a virtualization server if you don't want to use the virtual switch provided by the hypervisor.

ASAv—an ASA firewall implemented as a VM.