

Lecture 3

Fundamental Cloud Architectures Virtualization Technology



Contents :

- 1- Workload Distribution Architecture
- 2- Resource Pooling Architecture
- 3- Dynamic Scalability Architecture
- 4- Elastic Resource Capacity Architecture
- 5- Service Load Balancing Architecture
- 6- Cloud Bursting Architecture
- 7- Elastic Disk Provisioning Architecture

Workload Distribution Architecture

IT resources can be horizontally scaled through the addition of one or more identical IT resources, and a load balancer that provides runtime logic capable of evenly distributing the workload among the available IT resources (Figure 1).

The resulting workload distribution architecture reduces both IT resource over-utilization and under-utilization to an extent dependent upon the sophistication of the load balancing algorithms and runtime logic.

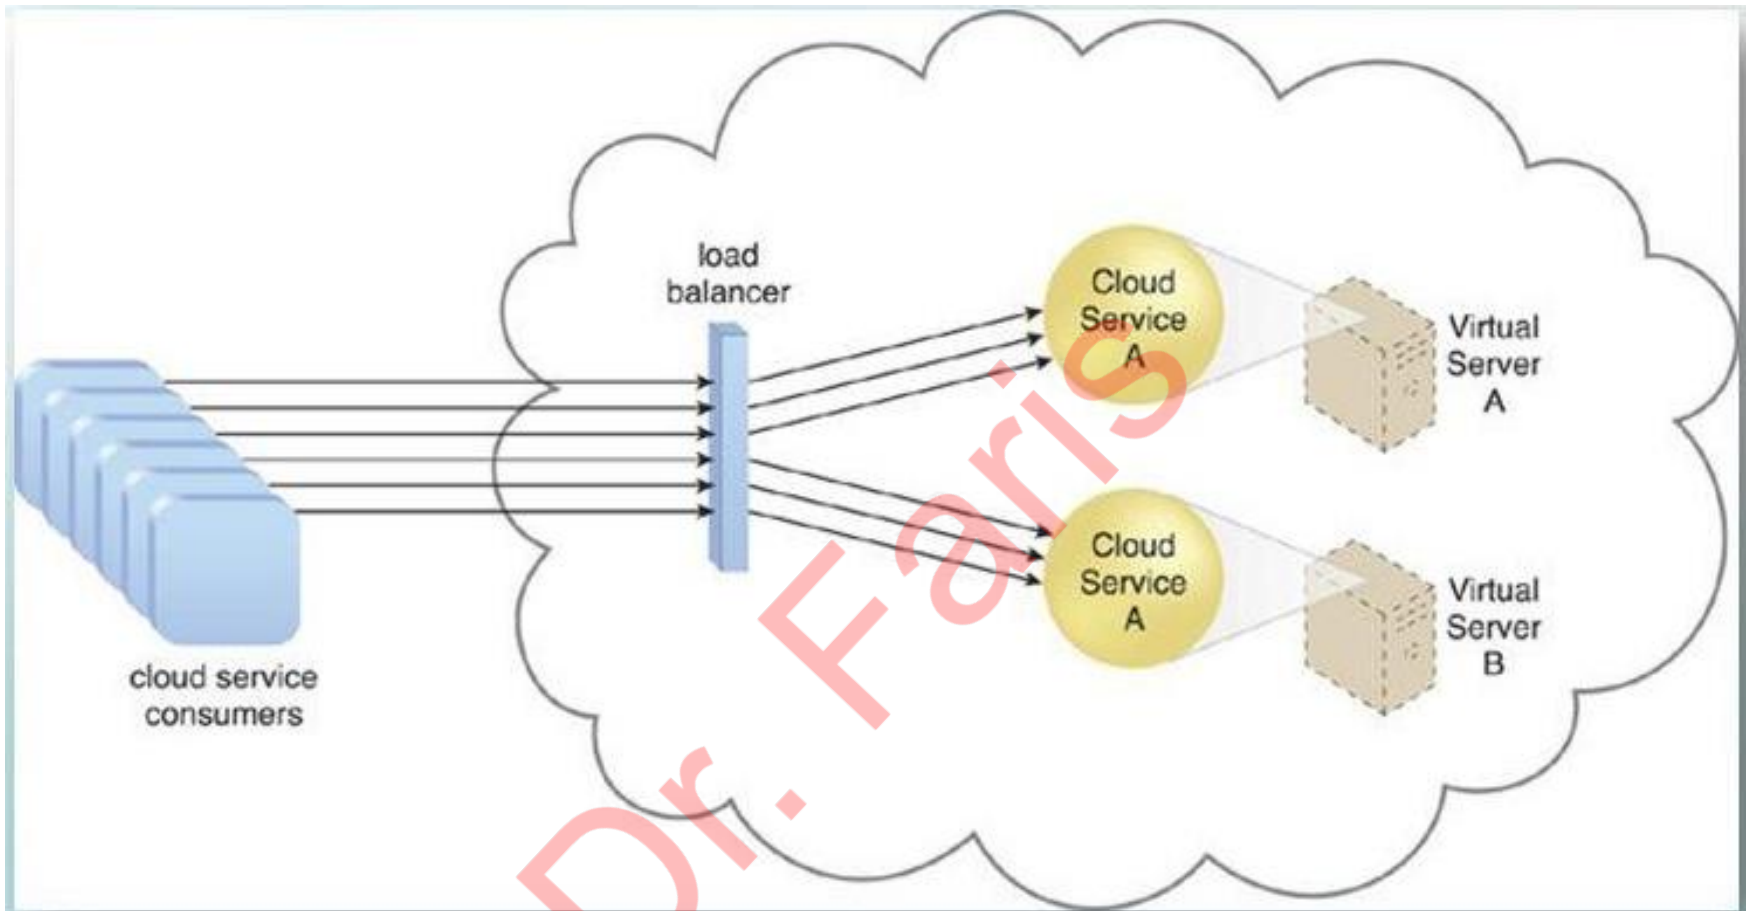


Figure 1. A redundant copy of Cloud Service A is implemented on Virtual Server B. The load balancer intercepts cloud service consumer requests and directs them to both Virtual Servers A and 3 to ensure even workload distribution.

Workload Distribution Architecture

This fundamental architectural model can be applied to any IT resource, with workload distribution commonly carried out in support of distributed virtual servers, cloud storage devices, and cloud services.

Load balancing systems applied to specific IT resources usually produce specialized variations of this architecture that incorporate aspects of load balancing.

Workload Distribution Architecture

The following mechanisms can also be part of this cloud architecture:

Audit Monitor — When distributing runtime workloads, the type and geographical location of the IT resources that process the data can determine whether monitoring is necessary to fulfill legal and regulatory requirements.

Cloud Usage Monitor — Various monitors can be involved to carry out runtime workload tracking and data processing.

Hypervisor — Workloads between hypervisors and the virtual servers that they host may require distribution.

Logical Network Perimeter — The logical network perimeter isolates cloud consumer network boundaries in relation to how and where workloads are distributed.


Resource Cluster — Clustered IT resources in active/active mode are commonly used to support workload balancing between different cluster nodes.

Resource Replication — This mechanism can generate new instances of virtualized IT resources in response to runtime workload distribution demands.

Resource Pooling Architecture


A resource pooling architecture is based on -the use of one or more resource pools, in which identical IT resources are grouped and maintained by a system that automatically ensures that they remain synchronized.

Provided here are common examples of resource pools:




Resource Pooling Architecture

Physical server pools are composed of networked servers that have been installed with operating systems and other necessary programs and/or applications and are ready for immediate use.




Virtual server pools are usually configured using one of several available templates chosen by the cloud consumer during provisioning. For example, a cloud consumer can set up a pool of mid-tier Windows servers with 4 GB of RAM or a pool of low-tier Ubuntu servers with 2 GB of RAM.

Resource Pooling Architecture




Storage pools, or cloud storage device pools, consist of file-based or block-based storage structures that contain empty and/or filled cloud storage devices.




Network pools (or interconnect pools) are composed of different preconfigured network connectivity devices. For example, a pool of virtual firewall devices or physical network switches can be created for redundant connectivity, load balancing, or link aggregation

Resource Pooling Architecture



Pools of physical RAM can be used in newly provisioned physical servers or to vertically scale physical servers



CPU pools are ready to be allocated to virtual servers, and are typically broken down into individual processing cores

Resource Pooling Architecture

Dedicated pools can be created for each type of IT resource and individual pools can be grouped into a larger pool; in which case each individual pool becomes a sub-pool (Figure.2)




Figure .2. A sample resource pool that is composed of four sub-poo's of CPUs, memory, cloud storage devices, and virtual network devices.

Resource Pooling Architecture

Resource pools can become highly complex, with multiple pools created for specific cloud consumers or applications. A hierarchical structure can be established to form parent, sibling, and nested pools in order to facilitate the organization of diverse resource pooling requirements (Figure .3).

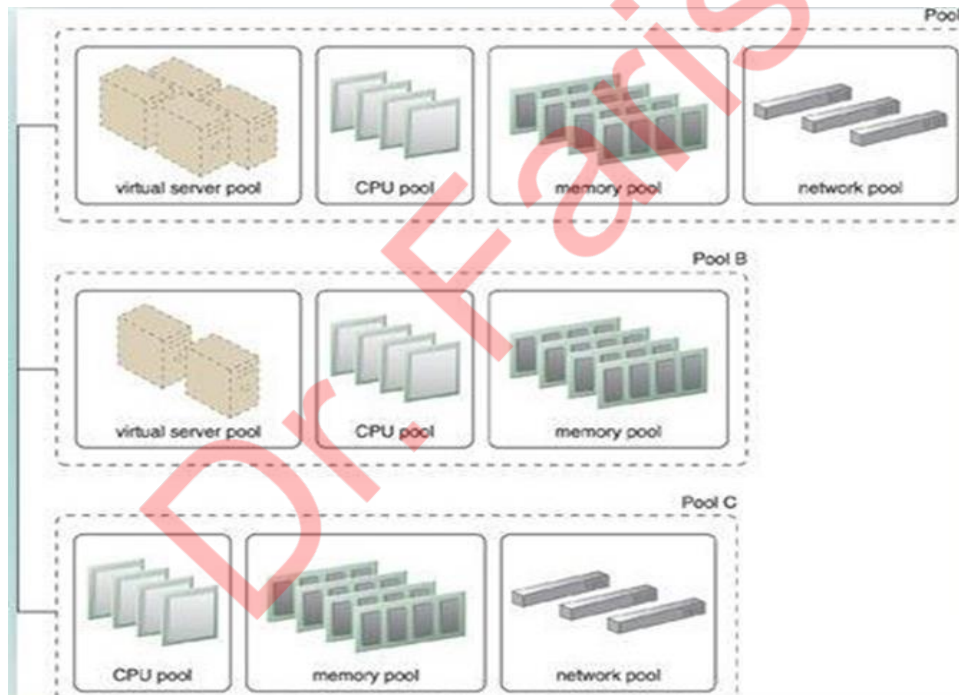


Figure .3. Pools B and C are sibling pools that are taken from the larger Pool A, which has been allocated to a cloud consumer. This is an alternative to taking the IT resources for Pool B and Pool C from a general reserve of IT resources that is shared throughout the cloud.

Resource Pooling Architecture

Sibling resource pools are usually drawn from physically grouped IT resources, as opposed to IT resources that are spread out over different data centers.

Sibling pools are isolated from one another so that each cloud consumer is only provided access to its respective pool.

In the nested pool model, larger pools are divided into smaller pools that individually group the same type of IT resources together (Figure .4).

Nested pools can be used to assign resource pools to different departments or groups in the same cloud consumer organization.

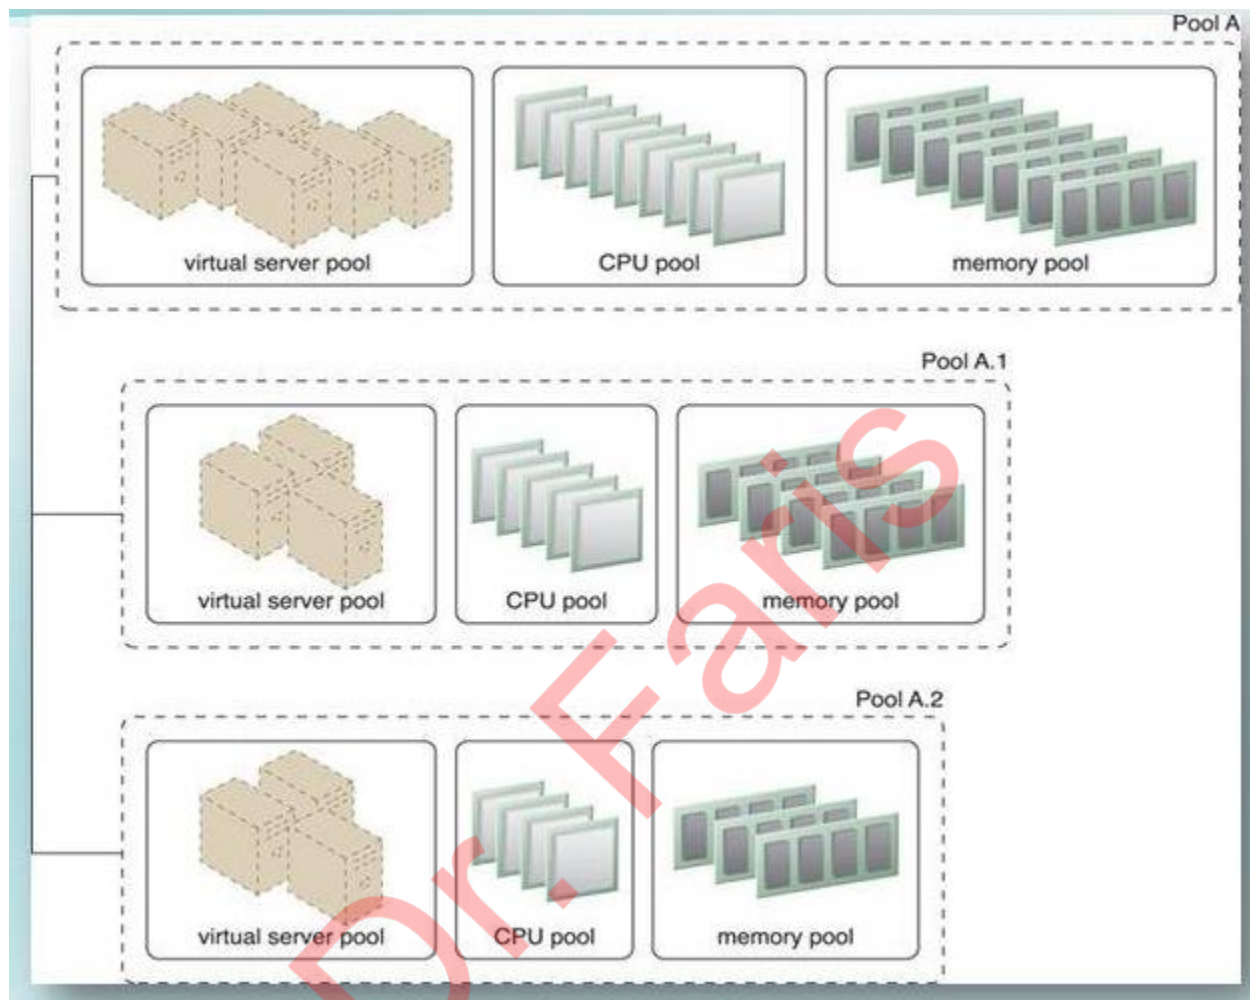


Figure .4. Nested Pools A1 and Pool A2 are comprised often same IT resources as Pool A, but jn different quantities. Nested pools are typically used to provision cloud services that need to be rapidly instantiated using the same type of IT resources with the same configuration settings.

Resource Pooling Architecture

The following mechanisms can also be part of this cloud architecture:

- **Audit Monitor** — This mechanism monitors resource pool usage to ensure compliance with privacy and regulation requirements, especially when pools contain cloud storage devices or data loaded into memory.
- **Cloud Usage Monitor**— Various cloud usage monitors are involved in the runtime tracking and synchronization that are required by the pooled IT resources and any underlying management systems.
- **Hypervisor** — The hypervisor mechanism is responsible for providing virtual servers with access to resource pools, in addition to hosting the virtual servers and sometimes the resource pools themselves.

Resource Pooling Architecture

- **Logical Network Perimeter** — The logical network perimeter is used to logically organize and isolate resource pools.
- **Pay-Per-Use Monitor** The pay-per-use monitor collects usage and billing information on how individual cloud consumers are allocated and use IT resources from various pools.
- **Remote Administration System** — This mechanism is commonly used to interface with backend systems and programs in order to provide resource pool administration features via a front-end portal.
- **Resource Management System** — The resource management system mechanism supplies cloud consumers with the tools and permission management options for administering resource pools.
- **Resource Replication** — This mechanism is used to generate new instances of IT resources for resource pools.

Dynamic Scalability Architecture

The dynamic scalability architecture is an architectural model based on a system of predefined scaling conditions that trigger the dynamic allocation of IT resources from resource pools.

Dynamic allocation enables variable utilization as dictated by usage demand fluctuations, since unnecessary IT resources are efficiently reclaimed without requiring manual interaction.

Dynamic Scalability Architecture

The following types of dynamic scaling are commonly used:

Dynamic Horizontal Scaling — IT resource instances are scaled out and in to handle fluctuating workloads. The automatic scaling listener monitors requests and signals resource replication to initiate IT resource duplication, as per requirements and permissions.

Dynamic Vertical Scaling — IT resource instances are scaled up and down when there is a need to adjust the processing capacity of a single IT resource. For example, a virtual server that is being overloaded can have its memory dynamically increased or it may have a processing core added.

Dynamic Relocation — The IT resource is relocated to a host with more capacity. For example, a database may need to be moved from a tape-based SAN storage device with 4 GB per second I/O capacity to another disk-based SAN storage device with 8 GB per second I/O capacity.

Figures .5 to .7 illustrate the process of dynamic horizontal scaling.

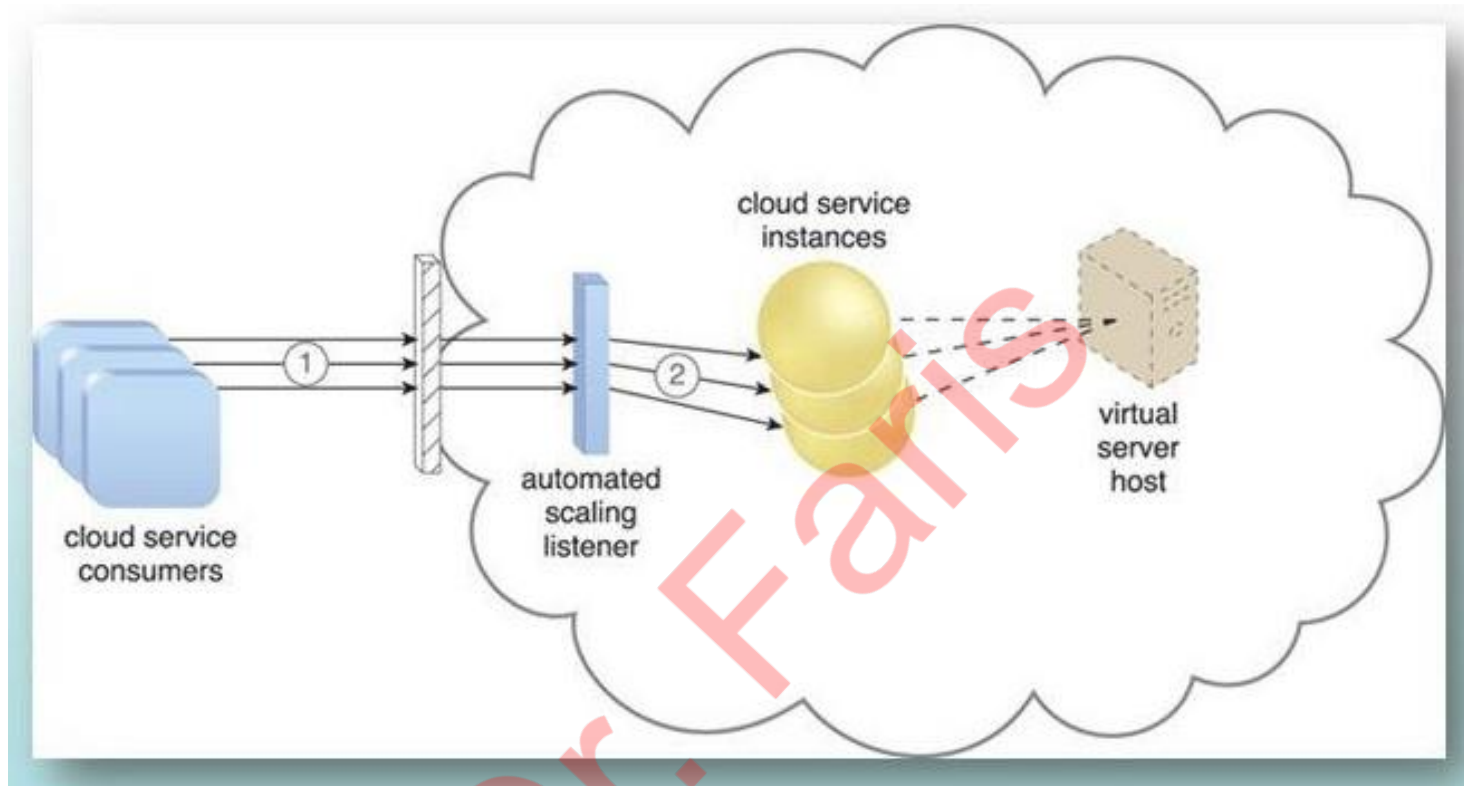


Figure 5. Cloud service consumers are sending requests to a cloud service. The automated scaling listener monitors the cloud service to determine if predefined capacity thresholds are being exceeded (2).

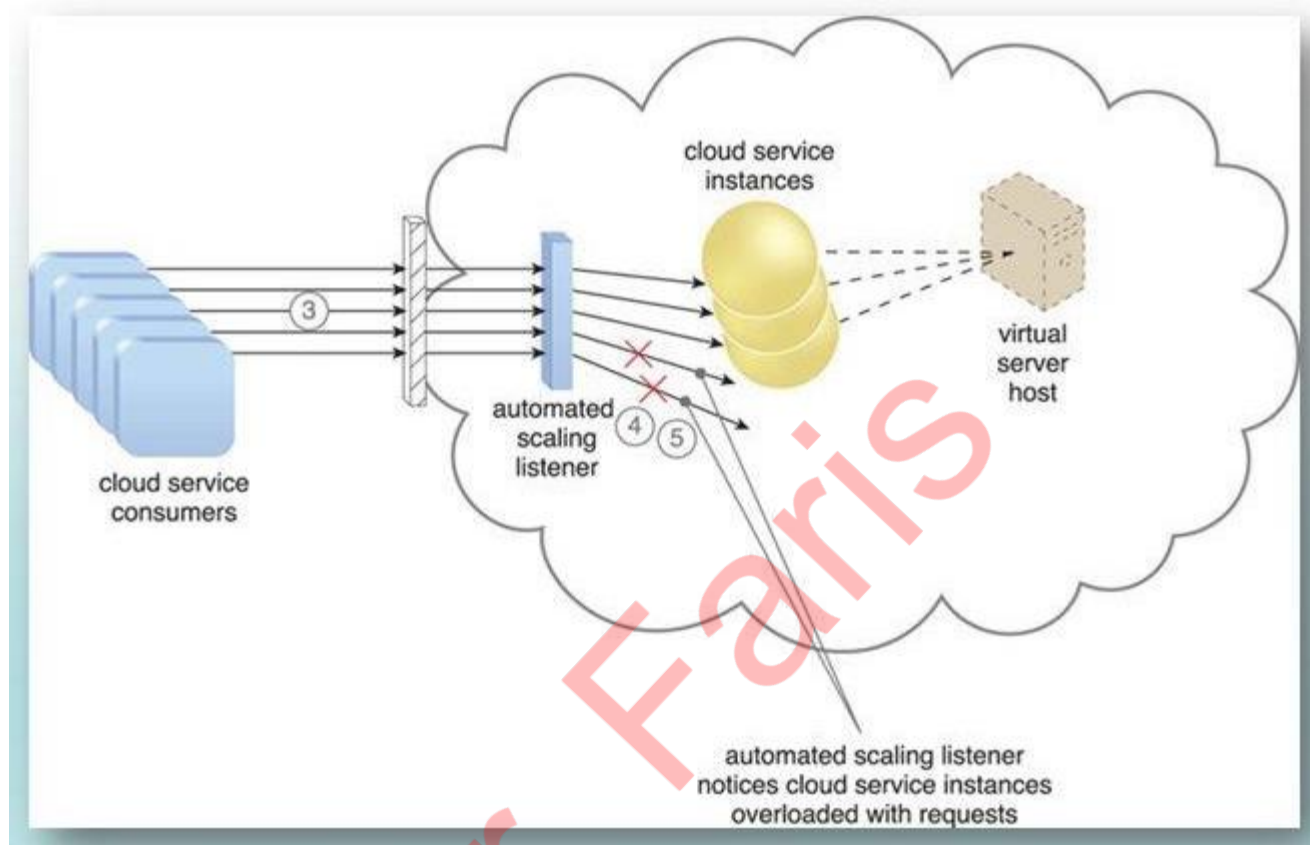


Figure 6. The number of requests coming from cloud service consumers increases (3). The workload exceeds the performance thresholds. The automated scaling listener determines the next course of action based on a predefined scaling policy (4). If the cloud service implementation is deemed eligible for additional scaling, the automated scaling listener initiates the scaling process

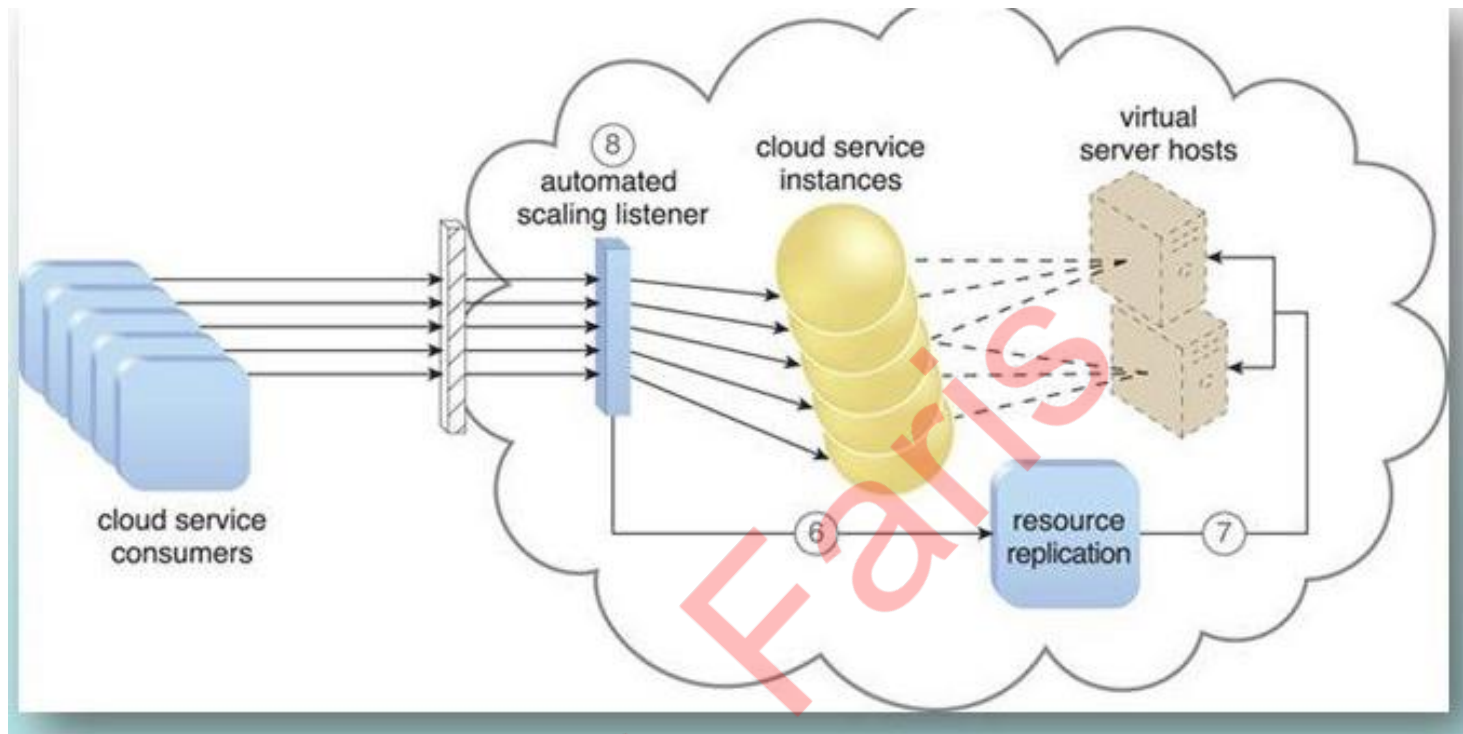


Figure 7. The automated scaling listener sends a signal to the resource replication mechanism which creates more instances of the cloud service (7). Now that the increased workload has been accommodated, the automated scaling listener resumes monitoring and detracting and adding iT resources, as required (8).

Dynamic Scalability Architecture

Besides the core automated scaling listener and resource replication mechanisms, the following mechanisms can also be used in this form of cloud architecture:

- **Cloud Usage Monitor** — Specialized cloud usage monitors can track runtime usage in response to dynamic fluctuations caused by this architecture.
- **Hypervisor** — The hypervisor is invoked by a dynamic scalability system to create or remove virtual server instances, or to be scaled itself.
- **Pay-Per-Use Monitor** — The pay-per-use monitor is engaged to collect usage cost information in response to the scaling of IT resources.

Elastic Resource Capacity Architecture

The elastic resource capacity architecture is primarily related to the dynamic provisioning of virtual servers, using a system that allocates and reclaims CPUs and RAM in immediate response to the fluctuating processing requirements of hosted IT resources.

(Figures.8 and .9)

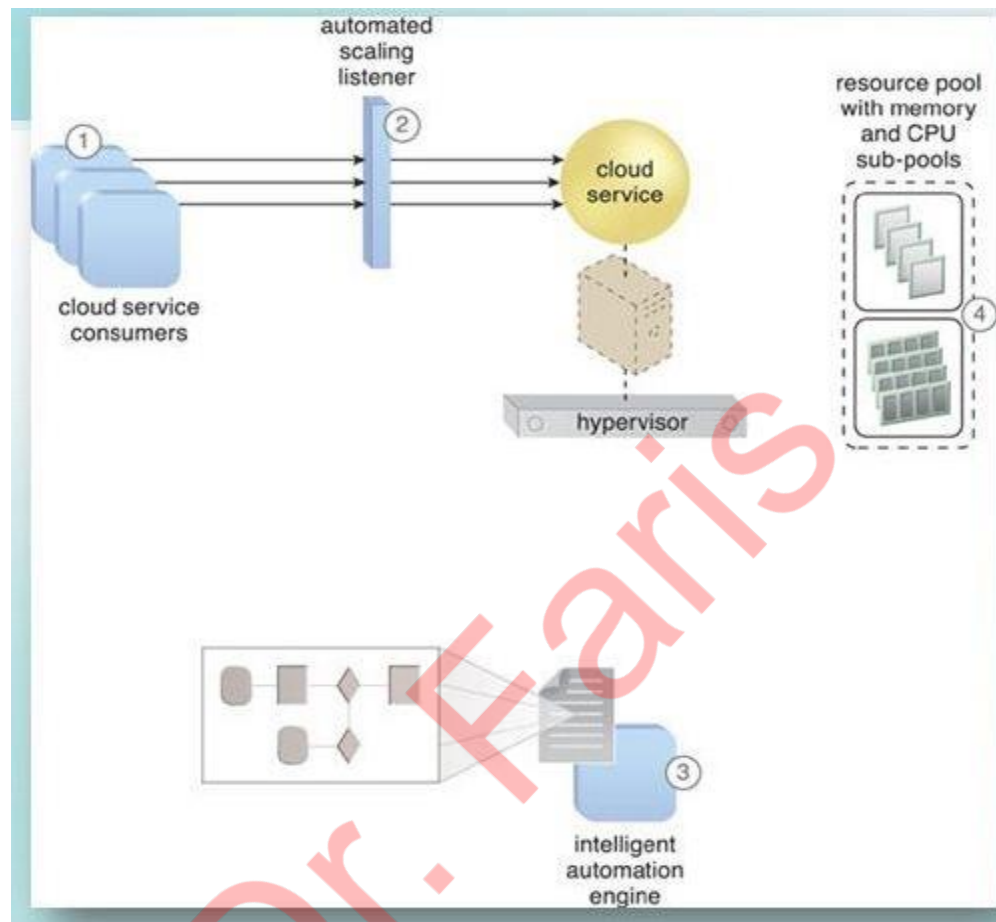


Figure 8. Cloud service consumers are actively sending requests to a cloud service (1), which are monitored by an automated scaling “listener” (2). An intelligent automation engine script is deployed with workflow logic (3) that is capable of notifying the resource pool using allocation requests

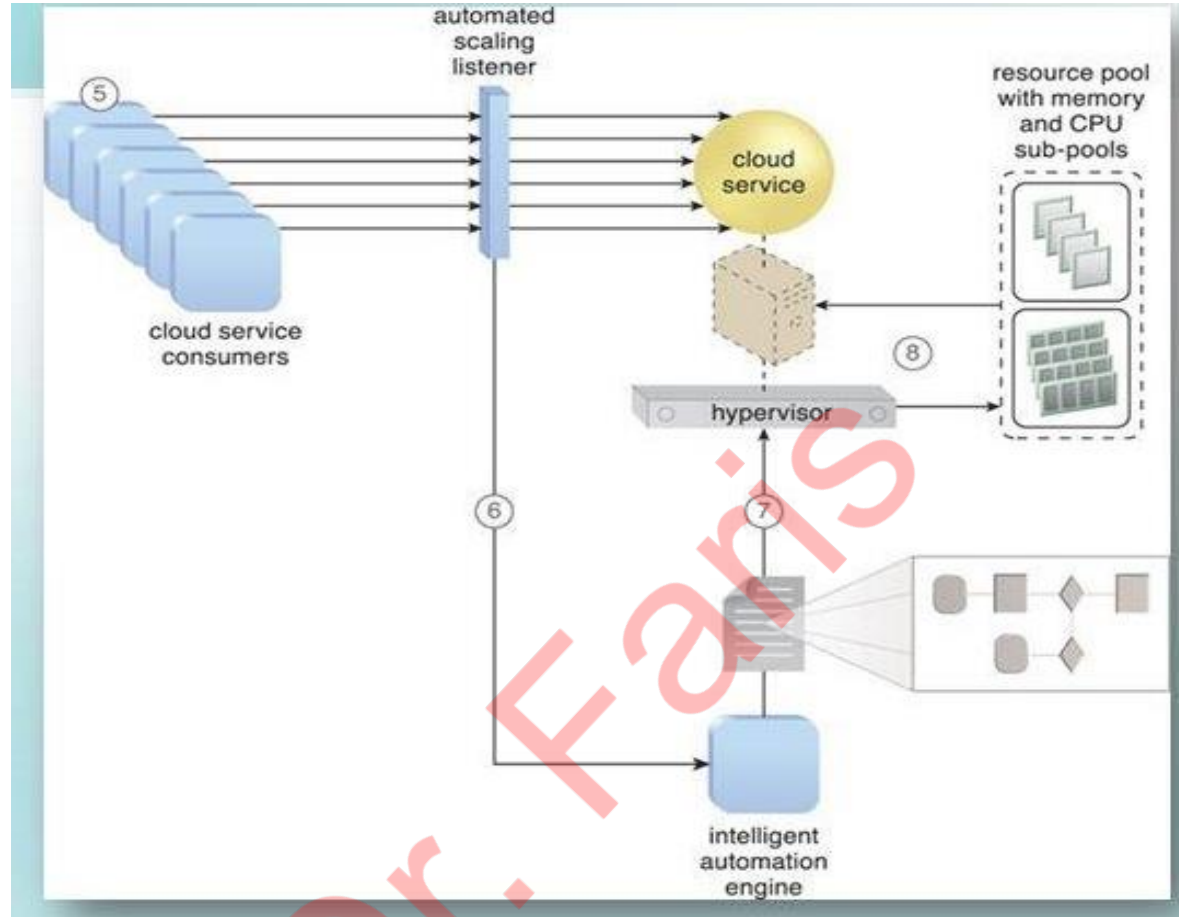



Figure 9. Cloud service consumer requests increase (5) causing the automated scaling listener to signal the intelligent automation engine to execute the script. The script runs the workflow Yogic that signals the hypervisor to allocate more IT resources from the resource pools (7). The hypervisor allocates additional CPU and RAM to the virtual server, enabling the increased workload to be handled (8).

Elastic Resource Capacity Architecture



Virtual servers that participate in elastic resource allocation systems may require rebooting in order for the dynamic resource allocation to take.

Elastic Resource Capacity Architecture

Some additional mechanisms that can be included in this cloud architecture are the following:

- **Cloud Usage Monitor** — Specialized cloud usage monitors collect resource usage information on IT resources before, during, and after scaling, to help define the future processing capacity thresholds of the virtual servers.
- **Pay-Per-Use Monitor** — The pay-per-use monitor is responsible for collecting resource usage cost information as it fluctuates with the elastic provisioning.
- **Resource Replication** — Resource replication is used by this architectural model to generate new instances of the scaled IT resources.

Service Load Balancing Architecture

The service load balancing architecture can be considered a specialized variation of the workload distribution architecture that is geared specifically for scaling cloud service implementations.

Redundant deployments of cloud services are created, with a load balancing system added to dynamically distribute workloads.

The duplicate cloud service implementations are organized into a resource pool, while the load balancer is positioned as either an external or built-in component to allow the host servers to balance the workloads themselves.

Service Load Balancing Architecture

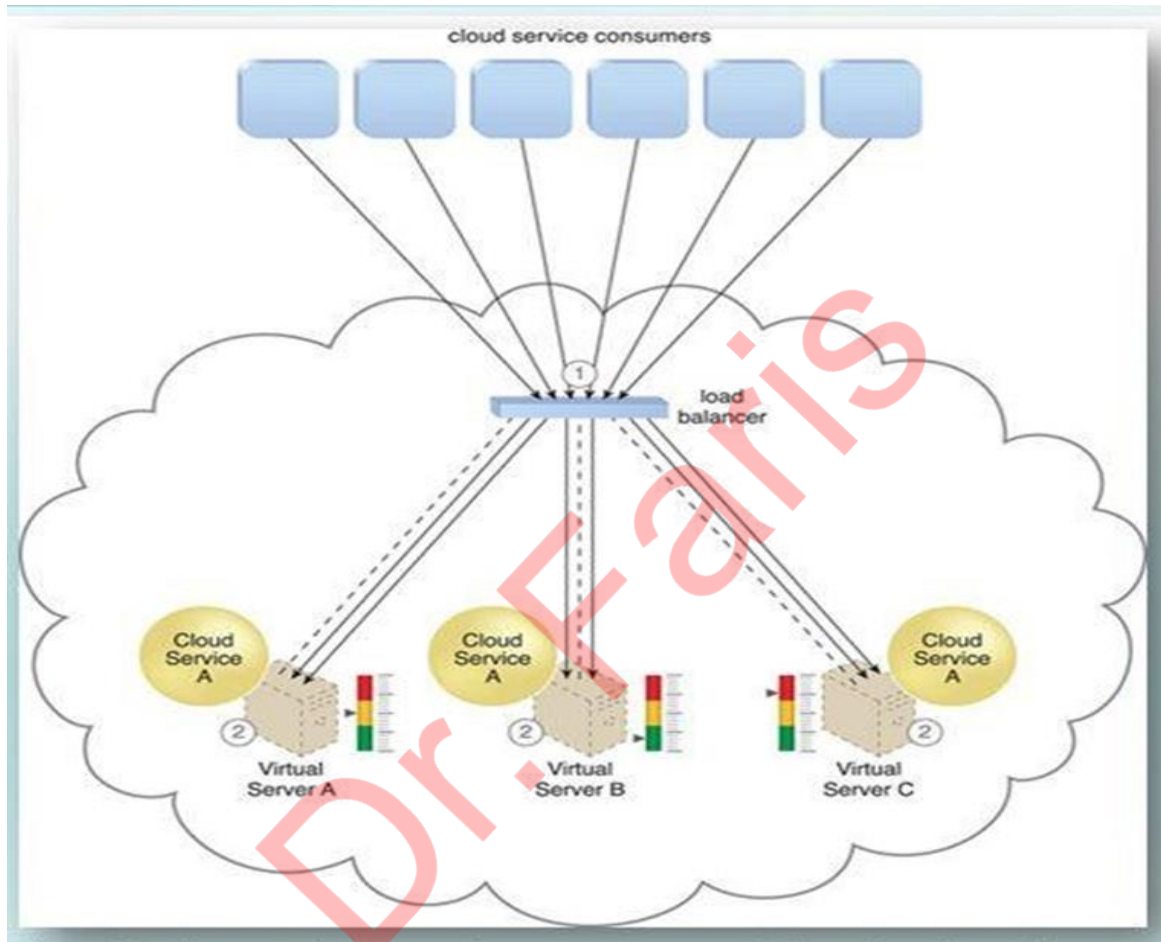


Figure 10. The load balancer intercepts messages sent by cloud service consumers (1) and forwards them to the virtual servers so that the workload processing is horizontally scaled

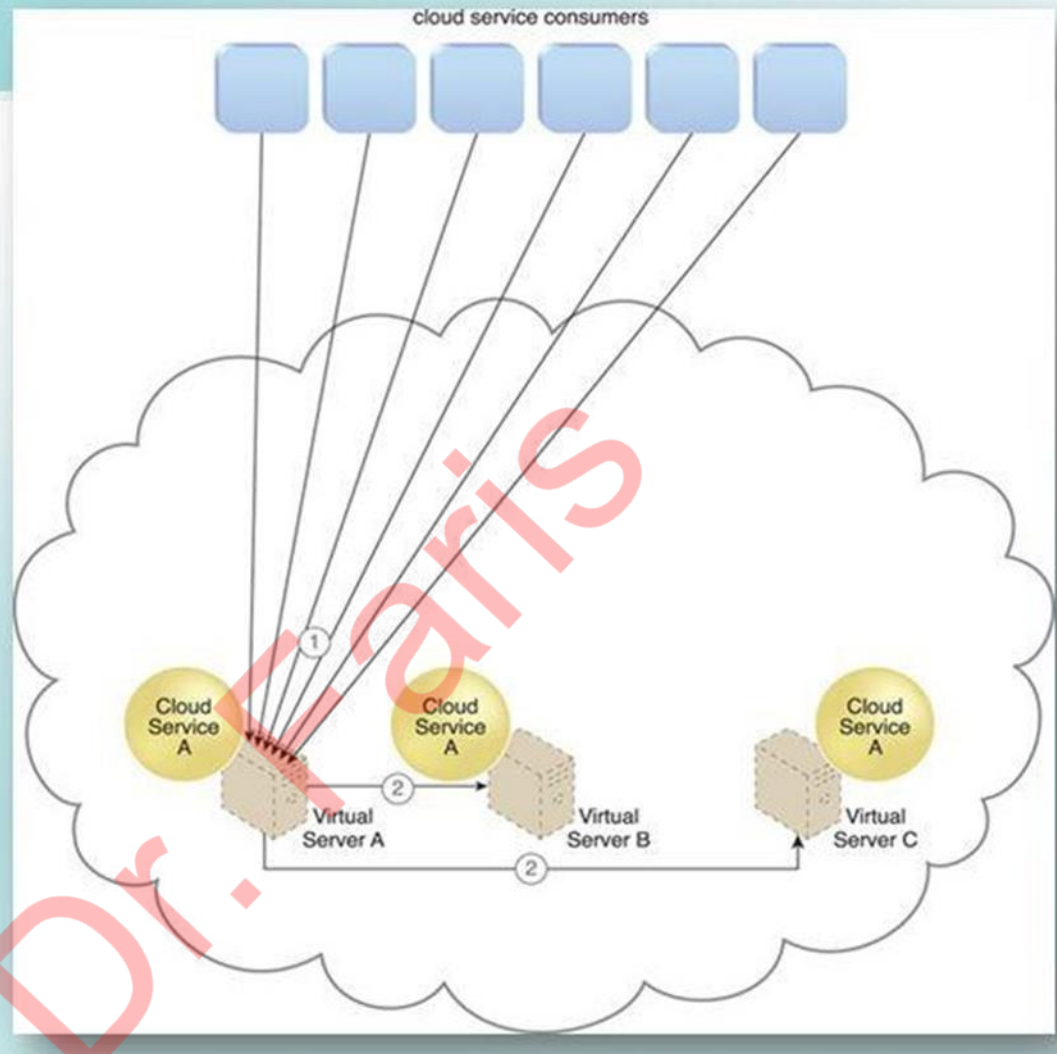


Figure 11. Cloud service consumer requests are sent to Cloud Service A on Virtual Server A (1). The cloud service implementation includes built-in load balancing logic that is capable of distributing requests to the neighboring Cloud Service A implementations on Virtual Servers B and C (2).

Service Load Balancing Architecture

The service load balancing architecture can involve the following mechanisms in addition to the load balancer:

- **Cloud Usage Monitor** — Cloud usage monitors may be involved with monitoring cloud service instances and their respective IT resource consumption levels, as well as various runtime monitoring and usage data collection tasks.
- **Resource Cluster** — Active-active cluster groups are incorporated in this architecture to help balance workloads across different members of the cluster.
- **Resource Replication** — The resource replication mechanism is utilized to generate cloud service implementations in support of load balancing requirements.

Cloud Bursting Architecture

The cloud bursting architecture establishes a form of dynamic scaling that scales or "bursts out" on-premise IT resources into a cloud whenever predefined capacity thresholds have been reached.

The corresponding cloud-based IT resources are redundantly pre-deployed but remain inactive until cloud bursting occurs. After they are no longer required, the cloud-based IT resources are released and the architecture "bursts in" back to the on-premise environment.

Cloud bursting is a flexible scaling architecture that provides cloud consumers with the option of using cloud-based IT resources only to meet higher usage demands.

The foundation of this architectural model is based on the automated scaling listener and resource replication mechanisms.

Cloud Bursting Architecture

The automated scaling listener determines when to redirect requests to cloud-based IT resources, and resource replication is used to maintain synchronicity between on premise and cloud-based IT resources in relation to state information (Figure 12).

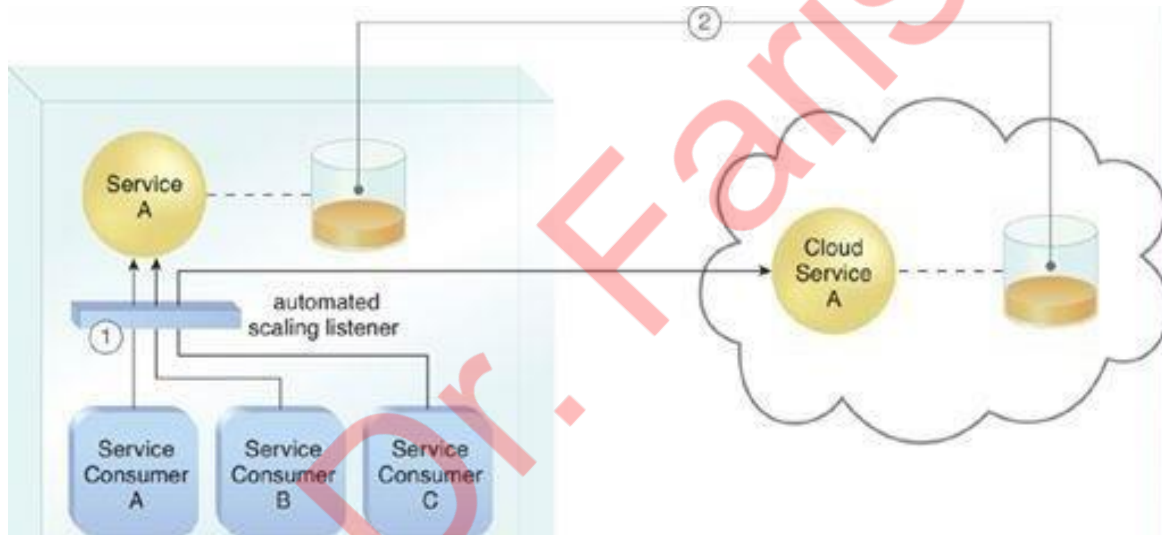


Figure 12 An automated scaling listener monitors the usage of on-premise Service A, and redirects Service Consumer C's request to Service A's redundant implementation in the cloud (Cloud Service A) once Service A's usage threshold has been exceeded. (1) A resource replication system is used to keep state management databases synchronized

Elastic Disk Provisioning Architecture

Cloud consumers are commonly charged for cloud-based storage space based on fixed-disk storage allocation, meaning the charges are predetermined by disk capacity and not aligned with actual data storage consumption.

Figure 13 demonstrates this by illustrating a scenario in which a cloud consumer provisions a virtual server with the Windows Server operating system and three 150 GB hard drives.

The cloud consumer is billed for using 450 GB of storage space after installing the operating system, even though the operating system only requires 15 GB of storage space.




Figure 13. The cloud consumer requests a virtual server with three hard disks, each with a capacity of 150 GB.

- (1). The virtual server is provisioned according to the elastic disk provisioning architecture, with a total of 450 GB of disk space.
- (2). The 450 GB is allocated to the virtual server by the cloud provider.
- (3). The cloud consumer has not installed any software yet, meaning the actual used space is currently 0 GB.
- (4). Because the 450 GB are already allocated and reserved for the cloud consumer, it will be charged for 450 GB of disk usage as of the point of allocation (5).

Elastic Disk Provisioning Architecture

The elastic disk provisioning architecture establishes a dynamic storage provisioning system that ensures that the cloud consumer is granularly billed for the exact amount of storage that it actually uses.

This system uses thin-provisioning technology for the dynamic allocation of storage space, and is further supported by runtime usage monitoring to collect accurate usage data for billing purposes (Figure 14).




Figure 14. The cloud consumer requests a virtual server with three hard disks, each with a capacity of 150 GB

- (1). The virtual server is provisioned by this architecture with a total of 450 GB of disk space.
- (2). The 450 GB are set as the maximum disk usage that is allowed for this virtual server, although no physical disk space has been reserved or allocated yet.
- (3). The cloud consumer has not installed any software, meaning the actual used space is currently at 0 GB.
- (4). Because the allocated disk space is equal to the actual used space (which is currently at zero), the cloud consumer is not charged for any disk space usage