



Assignment No. 02

Course Name: Cloud Computing

Submitted By:

saqibkhan (13158)

BS (SE) Section: A

Submitted To:

Sir M Omer Rauf

Dated: 23 May 2020

**Department of Computer Science,
IQRA National University, Peshawar Pakistan**

Question 01:

Explain in detail Service Oriented Architecture (SOA) in cloud computing.

Answer:

Definition:

Service-Oriented Architecture (SOA) is a style of software design where services are provided to the other components by application components, through a communication protocol over a network. Its principles are independent of vendors and other technologies. In [service oriented architecture](#), a number of services communicate with each other, in one of two ways: through passing data or through two or more services coordinating an activity.

Characteristics of Service-Oriented Architecture

While the defining concepts of Service-Oriented Architecture vary from company to company, there are six key tenets that overarch the broad concept of Service-Oriented Architecture. These core values include:

- Business value
- Strategic goals
- Intrinsic inter-operability
- Shared services
- Flexibility
- Evolutionary refinement

Each of these core values can be seen on a continuum from older format distributed computing to Service-Oriented Architecture to cloud computing (something that is often seen as an offshoot of Service-Oriented Architecture).

Service-Oriented Architecture Patterns

There are three roles in each of the Service-Oriented Architecture building blocks: service provider; service broker, service registry, service repository; and service requester/consumer.

The service provider works in conjunction with the service registry, debating the whys and how of the services being offered, such as security, availability, what to charge, and more. This role also determines the service category and if there need to be any trading agreements.

The service broker makes information regarding the service available to those requesting it. The scope of the broker is determined by whoever implements it.

The service requester locates entries in the broker registry and then binds them to the service provider. They may or may not be able to access multiple services; that depends on the capability of the service requester.

Why Service-Oriented Architecture Is Important

There are many benefits to service-oriented architecture, especially in a web service based business. We'll outline a few of those benefits here, in brief:

Use Service-Oriented Architecture to create reusable code: Not only does this cut down on time spent on the development process, but there's no reason to reinvent the coding wheel every time you need to create a new service or process. Service-Oriented Architecture also allows for using multiple coding languages because everything runs through a central interface.

Use Service-Oriented Architecture to promote interaction: With Service-Oriented Architecture, a standard form of communication is put in place, allowing the various systems and platforms to function independent of each other. With this interaction, Service-Oriented Architecture is also able to work around firewalls, allowing "companies to share services that are vital to operations."

Use Service-Oriented Architecture for scalability: It's important to be able to scale a business to meet the needs of the client, however certain dependencies can get in the way of that scalability. Using Service-Oriented Architecture cuts back on the client-service interaction, which allows for greater scalability.

Use Service-Oriented Architecture to reduce costs: With Service-Oriented Architecture, it's possible to reduce costs while still "maintaining a desired level of output." Using Service-Oriented Architecture allows businesses to limit the amount of analysis required when developing custom solutions.

How Service-Oriented Architecture and Cloud Computing Work Together

First, it's important to note that Service-Oriented Architecture can work with or without cloud computing, although more and more businesses are moving file

storage to the cloud so it makes sense to use cloud computing and Service-Oriented Architecture together.

In a nutshell, using cloud computing allows users to easily and immediately implement services tailored to the requirements of their clients, “without needing to consult an IT department.”

One downfall of using Service-Oriented Architecture and cloud computing together is that some aspects of it are not evaluated, such as security and availability. When using cloud computing, users are often at the mercy of the provider.

There is one fairly major challenge businesses face when merging cloud computing and Service-Oriented Architecture is the integration of existing data and systems into the cloud solution. There needs to be continuity from beginning to end in order for there to be a seamless transition. It’s also important to keep in mind that not every IT aspect can be outsourced to the cloud — there are some things that still need to be done manually.

You can read more about how service-oriented architecture and cloud computing work together right here.

Question 02:

Explain in detail prominent security threats to the cloud computing.

Answer:

The top security threats organizations face when using cloud services:

Threat No. 1: Data breaches

Cloud environments face many of the same threats as traditional corporate networks, but due to the vast amount of data stored on cloud servers, providers become an attractive target. The severity of potential damage tends to depend on the sensitivity of the data exposed. Exposed personal financial information tends to get the headlines, but breaches involving health information, trade secrets, and intellectual property can be more devastating.

MORE ON NETWORK WORLD: How to build a private cloud

When a data breach occurs, companies may incur fines, or they may face lawsuits or criminal charges. Breach investigations and customer notifications can rack up significant costs. Indirect effects, such as brand damage and loss of business, can impact organizations for years.

Threat No. 2: Compromised credentials and broken authentication

Data breaches and other attacks frequently result from lax authentication, weak passwords, and poor key or certificate management. Organizations often struggle with identity management as they try to allocate permissions appropriate to the user's job role. More important, they sometimes forget to remove user access when a job function changes or a user leaves the organization.

Multifactor authentication systems such as one-time passwords, phone-based authentication, and smartcards protect cloud services because they make it harder for attackers to log in with stolen passwords. The Anthem breach, which exposed more than 80 million customer records, was the result of stolen user credentials. Anthem had failed to deploy multifactor authentication, so once the attackers obtained the credentials, it was game over.

Many developers make the mistake of embedding credentials and cryptographic keys in source code and leaving them in public-facing repositories such as GitHub. Keys need to be appropriately protected, and a well-secured public key infrastructure is necessary, the CSA said. They also need to be rotated periodically to make it harder for attackers to use keys they've obtained without authorization.

Organizations planning to federate identity with a cloud provider need to understand the security measures the provider uses to protect the identity platform. Centralizing identity into a single repository has its risks. Organizations need to weigh the trade-off of the convenience of centralizing identity against the risk of having that repository become an extremely high-value target for attackers.

Threat No. 3: Hacked interfaces and APIs

Practically every cloud service and application now offers APIs. IT teams use interfaces and APIs to manage and interact with cloud services, including those that offer cloud provisioning, management, orchestration, and monitoring.

The security and availability of cloud services -- from authentication and access control to encryption and activity monitoring -- depend on the security of the API. Risk increases with third parties that rely on APIs and build on these interfaces, as organizations may need to expose more services and credentials, the CSA warned. Weak interfaces and APIs expose organizations to security issues related to confidentiality, integrity, availability, and accountability.

APIs and interfaces tend to be the most exposed part of a system because they're usually accessible from the open Internet. The CSA recommends adequate controls as the "first line of defense and detection." Threat modeling applications and systems, including data flows and architecture/design, become important parts of the development lifecycle. The CSA also recommends security-focused code reviews and rigorous penetration testing.

Threat No. 4: Exploited system vulnerabilities

System vulnerabilities, or exploitable bugs in programs, are not new, but they've become a bigger problem with the advent of multitenancy in cloud computing. Organizations share memory, databases, and other resources in close proximity to one another, creating new attack surfaces.

Fortunately, attacks on system vulnerabilities can be mitigated with "basic IT processes," says the CSA. Best practices include regular vulnerability scanning, prompt patch management, and quick follow-up on reported system threats.

According to the CSA, the costs of mitigating system vulnerabilities "are relatively small compared to other IT expenditures." The expense of putting IT processes in place to discover and repair vulnerabilities is small compared to the potential damage. Regulated industries need to patch as quickly as possible, preferably as part of an automated and recurring process, recommends the CSA. Change control processes that address emergency patching ensure that remediation activities are properly documented and reviewed by technical teams.

Threat No. 5: Account hijacking

Phishing, fraud, and software exploits are still successful, and cloud services add a new dimension to the threat because attackers can eavesdrop on activities, manipulate transactions, and modify data. Attackers may also be able to use the cloud application to launch other attacks.

Common defense-in-depth protection strategies can contain the damage incurred by a breach. Organizations should prohibit the sharing of account credentials between users and services, as well as enable multifactor authentication schemes where available. Accounts, even service accounts, should be monitored so that every transaction can be traced to a human owner. The key is to protect account credentials from being stolen, the CSA says.

Threat No. 6: Malicious insiders

The insider threat has many faces: a current or former employee, a system administrator, a contractor, or a business partner. The malicious agenda ranges from data theft to revenge. In a cloud scenario, a hellbent insider can destroy whole infrastructures or manipulate data. Systems that depend solely on the cloud service provider for security, such as encryption, are at greatest risk.

The CSA recommends that organizations control the encryption process and keys, segregating duties and minimizing access given to users. Effective logging, monitoring, and auditing administrator activities are also critical.

As the CSA notes, it's easy to misconstrue a bungling attempt to perform a routine job as "malicious" insider activity. An example would be an administrator who accidentally copies a sensitive customer database to a publicly accessible server. Proper training and management to prevent such mistakes becomes more critical in the cloud, due to greater potential exposure.

Threat No. 7: The APT parasite

The CSA aptly calls advanced persistent threats (APTs) "parasitical" forms of attack. APTs infiltrate systems to establish a foothold, then stealthily exfiltrate data and intellectual property over an extended period of time.

APTs typically move laterally through the network and blend in with normal traffic, so they're difficult to detect. The major cloud providers apply advanced techniques to prevent APTs from infiltrating their infrastructure, but customers

need to be as diligent in detecting APT compromises in cloud accounts as they would in on-premises systems.

Common points of entry include spear phishing, direct attacks, USB drives preloaded with malware, and compromised third-party networks. In particular, the CSA recommends training users to recognize phishing techniques.

Regularly reinforced awareness programs keep users alert and less likely to be tricked into letting an APT into the network -- and IT departments need to stay informed of the latest advanced attacks. Advanced security controls, process management, incident response plans, and IT staff training all lead to increased security budgets. Organizations should weigh these costs against the potential economic damage inflicted by successful APT attacks.

Threat No. 8: Permanent data loss

As the cloud has matured, reports of permanent data loss due to provider error have become extremely rare. But malicious hackers have been known to permanently delete cloud data to harm businesses, and cloud data centers are as vulnerable to natural disasters as any facility.

Cloud providers recommend distributing data and applications across multiple zones for added protection. Adequate data backup measures are essential, as well as adhering to best practices in business continuity and disaster recovery. Daily data backup and off-site storage remain important with cloud environments.

The burden of preventing data loss is not all on the cloud service provider. If a customer encrypts data before uploading it to the cloud, then that customer must be careful to protect the encryption key. Once the key is lost, so is the data.

Compliance policies often stipulate how long organizations must retain audit records and other documents. Losing such data may have serious regulatory consequences. The new EU data protection rules also treat data destruction and corruption of personal data as data breaches requiring appropriate notification. Know the rules to avoid getting in trouble.

Threat No. 9: Inadequate diligence

Organizations that embrace the cloud without fully understanding the environment and its associated risks may encounter a “myriad of commercial, financial, technical, legal, and compliance risks,” the CSA warned. Due diligence applies whether the organization is trying to migrate to the cloud or merging (or working) with another company in the cloud. For example, organizations that fail to scrutinize a contract may not be aware of the provider’s liability in case of data loss or breach.

Operational and architectural issues arise if a company's development team lacks familiarity with cloud technologies as apps are deployed to a particular cloud. The CSA reminds organizations they must perform extensive due diligence to understand the risks they assume when they subscribe to each cloud service.

Threat No. 10: Cloud service abuses

Cloud services can be commandeered to support nefarious activities, such as using cloud computing resources to break an encryption key in order to launch an attack. Other examples including launching DDoS attacks, sending spam and phishing emails, and hosting malicious content.

Providers need to recognize types of abuse -- such as scrutinizing traffic to recognize DDoS attacks -- and offer tools for customers to monitor the health of their cloud environments. Customers should make sure providers offer a mechanism for reporting abuse. Although customers may not be direct prey for malicious actions, cloud service abuse can still result in service availability issues and data loss.

Threat No. 11: DoS attacks

DoS attacks have been around for years, but they've gained prominence again thanks to cloud computing because they often affect availability. Systems may slow to a crawl or simply time out. “Experiencing a denial-of-service attack is like being caught in rush-hour traffic gridlock; there is one way to get to your destination and there is nothing you can do about it except sit and wait,” the report said.

DoS attacks consume large amounts of processing power, a bill the customer may ultimately have to pay. While high-volume DDoS attacks are very common,

organizations should be aware of asymmetric, application-level DoS attacks, which target Web server and database vulnerabilities.

Cloud providers tend to be better poised to handle DoS attacks than their customers, the CSA said. The key is to have a plan to mitigate the attack before it occurs, so administrators have access to those resources when they need them.

Threat No. 12: Shared technology, shared dangers

Vulnerabilities in shared technology pose a significant threat to cloud computing. Cloud service providers share infrastructure, platforms, and applications, and if a vulnerability arises in any of these layers, it affects everyone. "A single vulnerability or misconfiguration can lead to a compromise across an entire provider's cloud," the report said.

If an integral component gets compromised -- say, a hypervisor, a shared platform component, or an application -- it exposes the entire environment to potential compromise and breach. The CSA recommended a defense-in-depth strategy, including multifactor authentication on all hosts, host-based and network-based intrusion detection systems, applying the concept of least privilege, network segmentation, and patching shared resources.

Question 03:

Explain in detail Cloud Infrastructure Mechanisms.

Answer:

Cloud Infrastructure Mechanisms

- 1.** The following cloud infrastructure mechanisms are described in this chapter: • Logical Network Perimeter • Virtual Server • Cloud Storage Device • Cloud Usage Monitor • Resource Replication • Ready-Made Environment Not all of these mechanisms are necessarily broad-reaching, nor does each establish an individual architectural layer. Instead, they should be viewed as core components that are common to cloud platforms.
- 2.** Logical Network Perimeter Cloud Infrastructure Mechanisms Defined as the isolation of a network environment from the rest of a communications network,

the logical network perimeter establishes a virtual network boundary that can encompass and isolate a group of related cloud-based IT resources that may be physically distributed. (Figure 7.1). Figure 7.1. The dashed line notation used to indicate the boundary of a logical network perimeter. This mechanism can be implemented to:

- Isolate IT resources in a cloud from non-authorized users
- Isolate IT resources in a cloud from non-users
- Isolate IT resources in a cloud from cloud consumers
- Control the bandwidth that is available to isolated IT resources

Logical network perimeters are typically established via network devices that supply and control the connectivity of a data center and are commonly deployed as virtualized IT environments that include:

- Virtual Firewall – An IT resource that actively filters network traffic to and from the isolated network while controlling its interactions with the Internet.

3. Cloud Infrastructure Mechanisms • Virtual Network – Usually acquired through VLANs, this IT resource isolates the network environment within the data center infrastructure. Figure 7.2 introduces the notation used to denote these two IT resources. Figure 7.3 depicts a scenario in which one logical network perimeter contains a cloud consumer's on-premise environment, while another contains a cloud provider's cloud-based environment. These perimeters are connected through a VPN that protects communications, since the VPN is typically implemented by point-to-point encryption of the data packets sent between the communicating endpoints. Figure 7.2. The symbols used to represent a virtual firewall (top) and a virtual network (bottom).

4. Cloud Infrastructure Mechanisms Figure 7.3. Two logical network perimeters surround the cloud consumer and cloud provider environments.

5. Cloud Infrastructure Mechanisms Figure 7.4. A logical network layout is established through a set of logical network perimeters using various firewalls and virtual networks.

6. Cloud Infrastructure Mechanisms A virtual server is a form of virtualization software that emulates a physical server. Virtual servers are used by cloud providers to share the same physical server with multiple cloud consumers by providing cloud consumers with individual virtual server instances. Figure 7.5 shows three virtual servers being hosted by two physical servers. The number of instances a given physical server can share is limited by its capacity. The virtual

server represents the most foundational building block of cloud environments. Each virtual server can host numerous IT resources, cloud-based solutions, and various other cloud computing mechanisms. Cloud consumers that install or lease virtual servers can customize their environments independently from other cloud consumers that may be using virtual servers hosted by the same underlying physical server. Figure 7.6 depicts a virtual server that hosts a cloud service being accessed by Cloud Service Consumer B, while Cloud Service Consumer A accesses the virtual server directly to perform an administration task. 7.2. Virtual Server Figure 7.5. The first physical server hosts two virtual servers, while the second physical server hosts one virtual server.

7. Cloud Infrastructure Mechanisms Figure 7.7. Virtual servers are created via the physical servers' hypervisors and a central VIM.

8. Cloud Infrastructure Mechanisms Figure 7.6. A virtual server hosts an active cloud service and is further accessed by a cloud consumer for administrative purposes.

9. Cloud Infrastructure Mechanisms Figure 7.8. The cloud consumer uses the self-service portal to select a template virtual server for creation (1). A copy of the corresponding VM image is created in a cloud consumer-controlled cloud storage device (2). The cloud consumer initiates the virtual server using the usage and administration portal. (3), which interacts with the VIM to create the virtual server instance via the underlying hardware. (4). The cloud consumer is able to use and customize the virtual server via other features on the usage and administration portal. (5). (Note that the self-service portal and usage and administration portal are explained in chapter – 9.

10. Cloud Infrastructure Mechanisms 7.3. Cloud Storage Device The cloud storage device mechanism represents storage devices that are designed specifically for cloud-based provisioning. Instances of these devices can be virtualized, similar to how physical servers can spawn virtual server images. Cloud storage devices are commonly able to provide fixed-increment capacity allocation in support of the pay-per-use mechanism. Cloud storage devices can be exposed for remote access via cloud storage services. A primary concern related to cloud storage is the security, integrity, and confidentiality of data, which becomes more prone to being compromised when entrusted to external cloud providers and other third

parties. Cloud Storage Levels Cloud storage device mechanisms provide common logical units of data storage, such as:

- Files – Collections of data are grouped into files that are located in folders.
- Blocks – The lowest level of storage and the closest to the hardware, a block is the smallest unit of data that is still individually accessible.
- Datasets – Sets of data are organized into a table-based, delimited, or record format.
- Objects – Data and its associated metadata are organized as Web-based resources.

11. Cloud Infrastructure Mechanisms Each of these data storage levels is commonly associated with a certain type of technical interface which corresponds to a particular type of cloud storage device and cloud storage service used to expose its API (Figure 7.9). Figure 7.9. Different cloud service consumers utilize different technologies to interface with virtualized cloud storage devices.

(Adapted from the CDMI Cloud Storage Reference Model.)

12. Cloud Infrastructure Mechanisms Network Storage Interfaces Legacy network storage most commonly falls under the category of network storage interfaces. It includes storage devices in compliance with industry standard protocols, such as SCSI for storage blocks and the server message block (SMB), common Internet file system (CIFS), and network file system (NFS) for file and network storage. File storage entails storing individual data in separate files that can be different sizes and formats and organized into folders and subfolders. Original files are often replaced by the new files that are created when data has been modified. Storage processing levels and thresholds for file allocation are usually determined by the file system itself. Block storage requires data to be in a fixed format (known as a data block), which is the smallest unit that can be stored and accessed and the storage format closest to hardware Object Storage Interfaces Various types of data can be referenced and stored as Web resources. This is referred to as object storage, which is based on technologies that can support a range of data and media types. The Storage Networking Industry Association's Cloud Data Management Interface (SNIA's CDMI) supports the use of object storage interfaces.

13. Cloud Infrastructure Mechanisms Database Storage Interfaces Cloud storage device mechanisms based on database storage interfaces typically support a query language in addition to basic storage operations. Storage management is

carried out using a standard API or an administrative user-interface. This classification of storage interface is divided into two main categories according to storage structure, as follows : I. Relational Data Storage Relational databases (or relational storage devices) rely on tables to organize similar data into rows and columns. Tables can have relationships with each other to give the data increased structure, to protect data integrity, and to avoid data redundancy (which is referred to as data normalization). II. Non-Relational Data Storage Non-relational storage (also commonly referred to as NoSQL storage) moves away from the traditional relational database model in that it establishes a “looser” structure for stored data with less emphasis on defining relationships and realizing data normalization.

14. Cloud Infrastructure Mechanisms Figure 7.10. The cloud consumer interacts with the usage and administration portal to create a cloud storage device and define access control policies (1). The usage and administration portal interact with the cloud storage software to create the cloud storage device instance and apply the required access policy to its data objects (2). Each data object is assigned to a cloud storage device and all of the data objects are stored in the same virtual storage volume. The cloud consumer uses the proprietary cloud storage device UI to interact directly with the data objects (3). (Note that the usage and administration portal is explained in Chapter 9.)

15. Cloud Infrastructure Mechanisms Figure 7.11. The cloud consumer uses the usage and administration portal to create and assign a cloud storage device to an existing virtual server. (1). The usage and administration portal interacts with the VIM software. (2a), which creates and configures the appropriate LUN (2b). Each cloud storage device uses a separate LUN controlled by the virtualization platform. The cloud consumer remotely logs into the virtual server directly . (3a) to access the cloud storage device 3b).

16. Cloud Infrastructure Mechanisms 7.4. Cloud Usage Monitor The cloud usage monitor mechanism is a lightweight and autonomous software program responsible for collecting and processing IT resource usage data. Cloud usage monitors can exist in different formats. The upcoming sections describe three common agent-based implementation formats. Each can be designed to forward collect usage data to a log database for post-processing and reporting purposes.

Monitoring Agent A monitoring agent is an intermediary, event-driven program that exists as a service agent and resides along existing communication paths to transparently monitor and analyze data flows (Figure 7.12). This type of cloud usage monitor is commonly used to measure network traffic and message metrics.

17. Cloud Infrastructure Mechanisms Figure 7.12. A cloud service consumer sends a request message to a cloud service (1). The monitoring agent intercepts the message to collect relevant usage data (2) before allowing it to continue to the cloud service (3a). The monitoring agent stores the collected usage data in a log database (3b). The cloud service replies with a response message (4) that is sent back to the cloud service consumer without being intercepted by the monitoring agent (5).

18. Cloud Infrastructure Mechanisms Resource Agent A resource agent is a processing module that collects usage data by having event-driven interactions with specialized resource software (Figure 7.13). This module is used to monitor usage metrics based on pre-defined, observable events at the resource software level, such as initiating, suspending, resuming, and vertical scaling. Figure 7.13. The resource agent is actively monitoring a virtual server and detects an increase in usage (1). The resource agent receives a notification from the underlying resource management program that the virtual server is being scaled up and stores the collected usage data in a log database, as per its monitoring metrics (2).

19. Cloud Infrastructure Mechanisms Polling Agent A polling agent is a processing module that collects cloud service usage data by polling IT resources. This type of cloud service monitor is commonly used to periodically monitor IT resource status, such as uptime and downtime (Figure 7.14). Figure 7.14. A polling agent monitors the status of a cloud service hosted by a virtual server by sending periodic polling request messages and receiving polling response messages that report usage status "A" after a number of polling cycles, until it receives a usage status of "B" (1), upon which the polling agent records the new usage status in the log database (2).

20. Resource Replication Defined as the creation of multiple instances of the same IT resource, replication is typically performed when an IT resource's availability and performance need to be enhanced. Virtualization technology is

used to implement the resource replication mechanism to replicate cloud-based IT resources (Figure 7.16). Figure 7.16. The hypervisor replicates several instances of a virtual server, using a stored virtual server image.

21. Cloud Infrastructure Mechanisms Figure 7.17. A high-availability virtual server is running in Data Center A. VIM instances in Data Centers A and B are executing a coordination function that allows detection of failure conditions. Stored VM images are replicated between data centers as a result of the high-availability architecture.

22. Cloud Infrastructure Mechanisms 7.6. Ready-Made Environment The ready-made environment mechanism (Figure 7.20) is a defining component of the PaaS cloud delivery model that represents a pre-defined, cloud-based platform comprised of a set of already installed IT resources, ready to be used and customized by a cloud consumer. These environments are utilized by cloud consumers to remotely develop and deploy their own services and applications within a cloud. Typical ready-made environments include pre-installed IT resources, such as databases, middleware, development tools, and governance tools. Figure 7.20. A cloud consumer accesses a ready-made environment hosted on a virtual server. A ready-made environment is generally equipped with a complete software development kit (SDK) that provides cloud consumers with programmatic access to the development technologies that comprise their preferred programming stacks.

23. Cloud Infrastructure Mechanisms Figure 7.21. The developer uses the provided SDK to develop the Part Number Catalog Web application (1). The application software is deployed on a Web platform that was established by two ready-made environments called the front-end instance (2a) and the back-end instance (2b). The application is made available for usage and one end-user accesses its front-end instance (3). The software running in the front-end instance invokes a long-running task at the back-end instance that corresponds to the processing required by the end-user (4). The application software deployed at both the front-end and back-end instances is backed by a cloud storage device that provides persistent storage of the application data (5).