

Q.1. (a) [4 marks] List and briefly explain four reasons why resource sharing is beneficial.

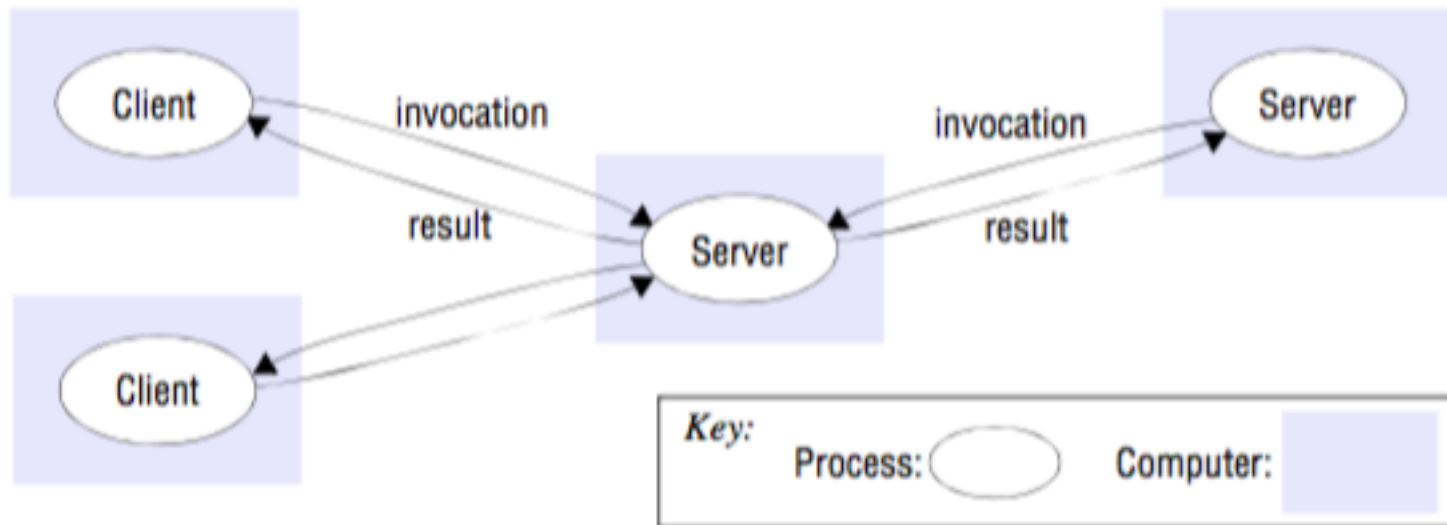
- Reduces cost by allowing a single resource for a number of users, rather than a identical resource for each user.
- Facilitates interactions among users, e.g. through a shared file system.
- Increases capacity, e.g. by allowing otherwise unused local disk space to be used remotely.
- Increases availability, e.g. through redundancy of resources.

(b) [1 marks] In the context of distributed systems, what is meant by the term independent failure?

- Components can fail separately, leaving others still running.
 - Faults in the network result in the isolation of the computers that are connected to it, but that doesn't mean that they stop running.
 - The crash of a program is not immediately made known to the other components with which it communicates.

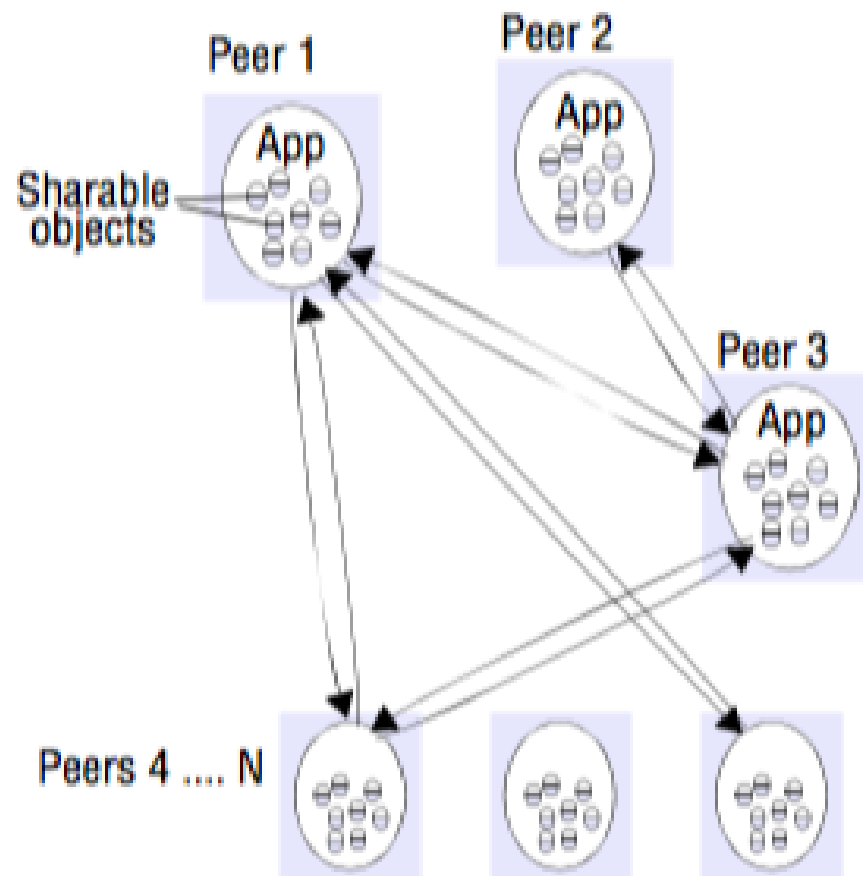
Q.3. (a) [4 marks] Draw a high level architecture diagram for each of the following architectural models and briefly explain each diagram:

i. Client/Server




Q.3. (a) [4 marks] Draw a high level architecture diagram for each of the following architectural models and briefly explain each diagram:

ii. Peer-to-Peer




Q.3. (a) [4 marks] Draw a high level architecture diagram for each of the following architectural models and briefly explain each diagram:

iii. A service provided by multiple servers

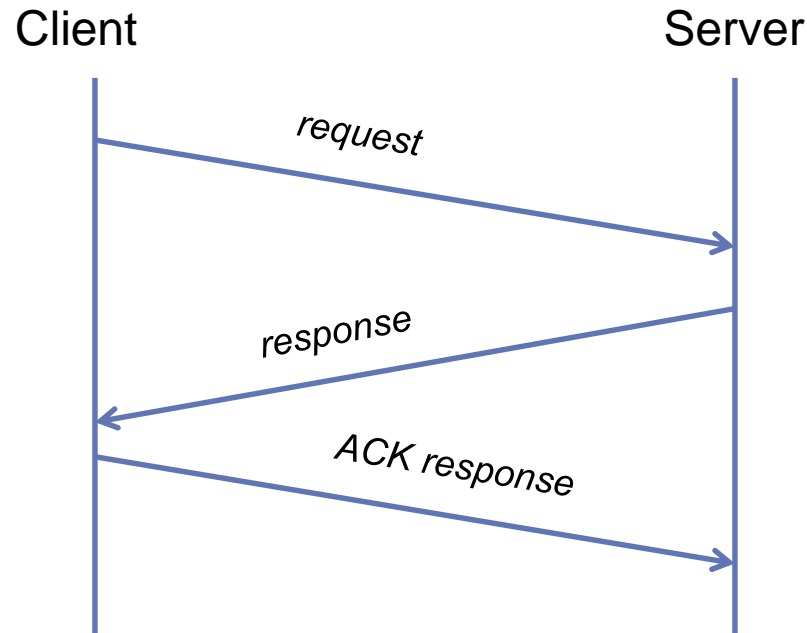


Q.3. (a) [4 marks] Draw a high level architecture diagram for each of the following architectural models and briefly explain each diagram:

ii. Proxy server



(b) [4 marks] Draw an interaction diagram to explain the RRA protocol. Explain the diagram. Explain what problem(s) the RRA protocol solves that the RR protocol does not.



- The RRA protocol acknowledges the response message
 - Allows the server to discard cached/stored results when a response acknowledgement is received

(c) [2 marks] Consider a client/server system where the server is sending a stream of temperature readings to the client. Each temperature reading is a 32bit floating point number. Design a format for the stream, using either a well-known data representation, or your own representation. Explain your design exactly, so that a third person could implement it without ambiguity.

- Binary stream: First byte indicates number of temperature readings to follow (0 to 255), next 4 bytes correspond to a reading until the number of readings has been reached, last byte indicates whether the stream should be closed (0) or whether more readings will follow.
- Other options:
 - JSON
 - XML

Q.4. (a) [4 marks] Answer the following questions about Java RMI:

i. Can objects be created remotely? Explain your answer.

Objects cannot be directly created remotely, i.e. with the new command, but a servant factory object can be used to indirectly create an object on a remote machine.

ii. Explain the difference between a local invocation and a remote invocation. If two Java Virtual Machines are on the same physical machine, and invocation is made between them, is this local or remote?

A local invocation takes place between objects within the same Java virtual machine; whereas a remote invocation takes place between objects on different Java virtual machines. If the two Java virtual machines are on the same physical machine, it is still considered a remote invocation.

iii. What is a remote reference? Explain how an object obtains a remote reference.

A remote reference is a reference to an object that is not within the local Java virtual machine. An object can obtain a remote reference by accessing the registry process, knowing the objects name.

(b) [2 marks] Explain the following aspects of a publish subscribe system:

i. **Event**

- A change of state that is of interest.

ii. **Notification**

- Information regarding an event that is sent to a subscriber.

iii. **Subscriber**

- Express interest in particular events and get notified when they occur.

iv. **Publisher**

- Publish/send structured events to an event service.

Q.5. (a) [2 marks] Explain what is a persistent asynchronous invocation and describe an example application that would benefit from this technique.

- A persistent asynchronous invocation tries indefinitely to perform the invocation, until it is known to have succeeded or failed, or until the application cancels the invocation. Generally, it will continue to be attempted over a long period of time if errors occurs.
- E.g. sending an SMS from a mobile phone in the presence of connection problems.

(b) [2 marks] Explain what is meant by process migration. Explain two major complications with process migration.

- The transfer of an executing process from one node to another.
- It involves suspending the execution of the running process, capturing its state within the kernel, and transferring it to another machine to resume its execution.
- Complications:
 - The local and remote machine may not share the same instruction set.
 - The process may have references to resources that are bound to the local machine.

(c) [4 marks] Apart from actual network delay, list and briefly explain four factors that contribute to the delay incurred when making an RMI call.

- **Marshalling and unmarshalling:** Involve copying and converting data and create a significant overhead as the amount of data grows.
- **Data copying:** Message data is copied several times in the course of an RMI (across the user–kernel boundary, across each protocol layer, between the network interface and kernel buffers).
- **Packet initialization:** Initializing protocol headers and trailers, including checksums. The cost is therefore proportional, in part, to the amount of data sent.
- **Thread scheduling and context switching:** Several context switches are made during an RMI, as stubs invoke the kernel's communication operations or one or more server threads are scheduled.

Q.6. (a) [3 marks] List and briefly explain three worst-case assumptions when designing a secure system.

- All communications between processes can be copied, modified and retransmitted. Attackers can obtain information that they should not and can pretend to be a legitimate party.
- All of the source code is known to the attacker. Knowing the source code can help the attacker discover vulnerabilities.
- The attacker has unlimited computing resources.

(b) [4 marks] Explain what is a digital certificate, including what is the basic technique used to create a digital certificate, and what is a certificate chain.

- A digital certificate is a document containing a statement signed by a principal.
- It binds information together, like a principal's id with its public key, and is digitally signed by a certificate authority.
- Sara creates a digital signature by encrypting a digest of the data to be signed with her private key.
- Certificate chain:
 - A series of certificates where each certificate's signature is authenticated by the subsequent certificate

Q.8. [3 marks] Before DNS, a single file was used to store all name information for computers on the Internet. This file was downloaded by everyone on the Internet on a daily basis, from the well known host that provided it. List and briefly explain three problems with this approach, that prompted people to develop DNS.

- Scalability
- Difficulty to maintain
- Availability

Q.9. [5 marks] Consider an image-resizing service. Clients submit an image to the service along with a resize factor and the service returns the image scaled by the resize factor. The service is expected to act in real time { e.g. for images of less than 1 megabyte in size the service should respond in less than 500ms for 99% of the requests. Also, the system should provide 99.9% availability in a month, meaning that the system is rarely unavailable. Using your knowledge of distributed systems, propose an overall approach for this system, draw an architecture diagram, and explain the protocol that you propose. Your approach should address the common challenges of a distributed system and specially the challenges suggested above. Explain how it does so. Your answer should be less than one page.

- Emphasis is on the architecture of the system that is best suited to provide availability and low response times.
- E.g. a multi-server architecture
 - I/O servers
 - Worker servers
- TCP or UDP? Why?
- Message format?