0907741 Distributed Systems (Spring 2014)

Midterm Exam

الاسم: الرقم التسلسلي: رقم الشعبة: 1

==========

Instructions: Time **70** minutes. Open book and notes exam. No electronics. Please answer all problems in the space provided and limit your answer to the space provided. **No questions are allowed**.

<Good Luck>

Q1. Assume that you want to build a global system for providing weather forecasts to millions of users all over the world. Describe a scalable solution for this service and discuss the scalability and reliability of your solution.

[6 marks]

The system should comprise multiple content providing servers distributed in data centers in multiple locations in the world. These servers will receive user requests and provide weather information and forecasts. Each data center could have multiple servers with fast load distribution mechanism. Faulty servers must be detected and isolated. The weather forecasts could be developed in a central massively-parallel system with redundant components. This system collects current weather information and runs detailed global weather models to develop the required weather forecasts. These weather information and forecasts are distributed to the content providing servers by the parallel system.

Scalability is achieved through using multiple content providing servers distributed all over the world. The massively-parallel server should have enough computational resources to develop timely global weather forecasts.

Reliability is achieved through the presence of multiple redundant servers and components.

Q2. Which one of the following alternative client-server organizations would you choose for the global weather forecasts system described in Problem 1 assuming that the clients access the system through smart phones? And explain your choice?

[6 marks]

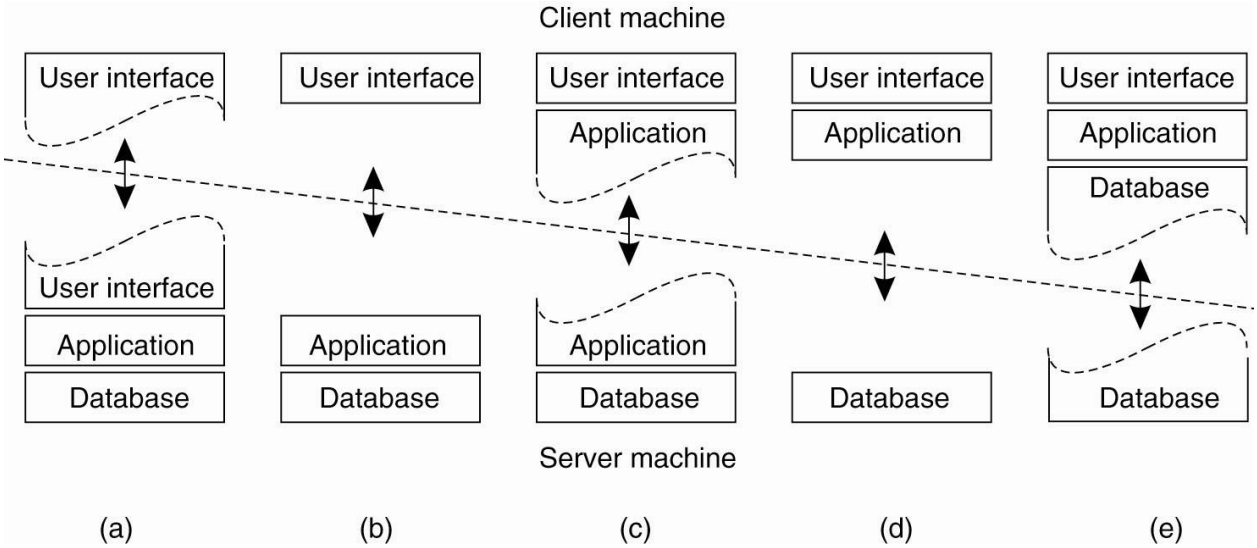

Organziation (e) is suitable for this application. The smart phone application would handle the user interface, communicate with the content providing service, and perform some processing like accessing the GPS location to retreive the weather information of the current location.

It should cache the retrieved information to reduce communication and power and to provide some content even with working offline.

Q3. Recall the communication types and mechanisms studied in the class. What communication type and mechanism would you choose (and why) for the global weather forecasts system described in Problem 1 for the following:

[6 marks]

a) Among the clients from one side and servers from the other side?

Transient synchronous communication from the clients to the content providing servers is sufficient as both parties should be active during the communication. This would save the costs of setting up the persistent communication infrastructure.

Sockets communication is suitable in this case where a server receives and responds to requests from many servers.

b) Among the servers?

Persistent asynchronous communication (possibly with multicast features) from the parallel central server to the content providing servers insures that updates will reach the content providing servers without much overhead on the central server side.

RPC and RMI mechanisms are usually used to deliver the message-oriented communication to the destination content providing servers.

Q4. Recall the process/thread options studied in the class. What process/thread design would you choose (and why) for the global weather forecasts system described in Problem 1 for the following:

[6 marks]

a) Client program?

The client program should use kernel threads to hide the communication overheads with the content providing servers. When one thread is handling the communication and another one is handling the user interface, the program keeps responding to the user even when the program is still waiting for some communication request.

b) Server program?

The content providing server program should use hybrid user and kernel threads to achieve highest performance when dealing with many client requests. A new thread is spawned for every new client request.

Q5. Given the following directory graph, describe what the operating system needs to do in order to read the first block of the file /home/max.

[6 marks]

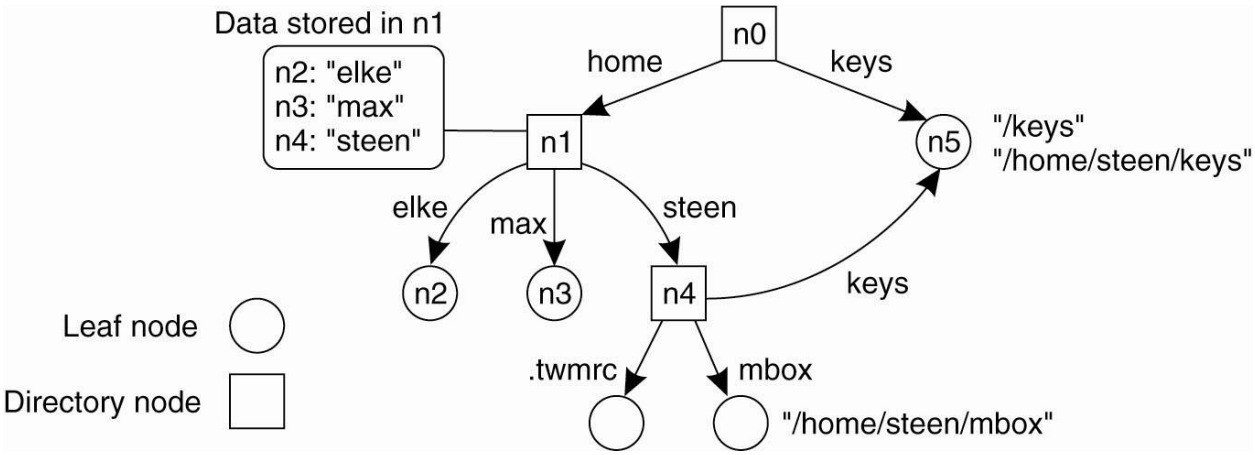

The OS first parses the string "/home/max" into the tokens "/", "home", and "max". It accesses the known home node n0 and searches for the entry "home". The hit on "home" returns the address of directory node n1. n1 is read and its entries are searched for "max". The hit on max returns the address of the file n3 which is accessed to read the contents of the required file.