



Last name:

First name:

Student number:

---

## Question 1)

Points: .....

Maximum points: 4

How long does it take to transfer 7.5 TB via a 1 Gbps (= 1,000 Mbps) Ethernet?





Last name:

First name:

Student number:

---

## Question 4)

Points: .....

Maximum points:  $1+2+2+1+1=7$

- a) For exercise sheet 7, you implemented a remote desktop solution for a Linux instance and a Microsoft Windows instance both. Name a protocol, you used to implement a graphical remote desktop solution.
  
- b) If you create a cluster of virtual server instances in EC2, you can distribute the instances over multiple regions. Give an advantage and a drawback of this method.
  
  
  
  
  
  
  
  
  
  
- c) If you create a cluster of virtual server instances in EC2, you can distribute the instances over multiple availability zones. Give an advantage and a drawback of this method.
  
  
  
  
  
  
  
  
  
  
- d) For exercise sheet 8, you implemented with the infrastructure services of the Amazon Web Services, a highly available High Throughput Cluster of virtual web servers. Which web server software did you use?
  
  
  
  
  
  
  
  
  
  
- e) For exercise sheet 8, you implemented with the infrastructure services of the Amazon Web Services, a highly available High Throughput Cluster of virtual web servers. The web server data was stored in EBS volumes. Which Linux file system did you deploy on the EBS volumes?



Last name:

First name:

Student number:

---

## Question 6)

Points: .....

Maximum points:  $3+7+2=12$

Your local time in Frankfurt am Main is Monday 09:00 (UTC+1). You need to copy 3 TB of data into the storage service S3. You have two options:

- **Scenario 1:** You immediately start at 09:00 (UTC+1) to upload the 3 TB of data to S3 via the internet. Consider the data rate between your computer and S3 is 100 Mbit/s.

- **Scenario 2:** You use the AWS Import/Export service. Therefore you copy the data to a HDD, which is connected via USB 3.0. The transfer rate (for write) is 125 MB/s.

After you copied the data, you pack the HDD into a parcel and send it via a package delivery company to Amazon. DHL, UPS and FedEx can deliver a parcel from Frankfurt am Main in less than 24 hours to most places in Europe.

You need 15 Minutes to put the HDD into a parcel and another 15 Minutes to bring the parcel to the branch office of a package delivery company.

The parcel must arrive at the branch office of the package delivery company no later than 16:30 (UTC+1) to arrive at Amazon in Ireland at 9:00 (UTC) the next working day.

An Amazon employee needs to copy the data from the HDD to the S3 service. The transfer rate of the HDD (for read) is 150 MB/s.

Consider 3 hours additional overhead for the in-house mail at Amazon to ship the HDD to the correct employee.

Calculate...

- a) for the first scenario, how long it takes until the data is copied to S3.
- b) for the second scenario, how long it takes until the data is copied to S3.
- c) the data rate of the second scenario.

*(For all subtasks, the calculation steps must be visible.)*

Last name:

First name:

Student number:

---

## Question 6 – Additional Page)

Maximum points:  $3+7+2=12$



Last name:

First name:

Student number:

---

## Question 7)

Points: .....

Maximum points: 2+4=6

Company X runs 8,000 computer workplaces.

- Scenario 1: Fat clients (PC)
  - Electrical power rating per desktop: 350 watts
  - Electrical power rating per screen: 80 watts
- Scenario 2: Thin clients
  - Electrical power rating per thin client: 40 watts
  - Electrical power rating per screen: 80 watts
  - Electrical power rating per server blade: 400 watts
  - Each server blade has enough resources to run 50 virtual desktops

Calculate for both scenarios the electricity costs per year for 24/7 operation when the electricity price is 0,28 € per kWh?

Last name:

First name:

Student number:

---

## Question 8)

Points: .....

Maximum points: 4

- a) Google Cloud Print implements...  
 IaaS       PaaS       SaaS
- b) Amazon S3 implements...  
 IaaS       PaaS       SaaS
- c) Google App Engine implements...  
 IaaS       PaaS       SaaS
- d) Amazon EC2 implements...  
 IaaS       PaaS       SaaS
- e) AppScale implements...  
 IaaS       PaaS       SaaS
- f) Google Cloud Storage implements...  
 IaaS       PaaS       SaaS
- g) Google Compute Engine implements...  
 IaaS       PaaS       SaaS
- h) Microsoft Office 365 implements...  
 IaaS       PaaS       SaaS

Last name:

First name:

Student number:

# Question 9)


Points: .....

Maximum points: 9

- $PR_p$  = PageRank of a web page  $p$
- $L_{IN}(p)$  = Set of documents, which refer to  $p \implies$  incoming links
- $L_{OUT}(p)$  = Set of documents, to which  $p$  refers  $\implies$  outgoing links
- $d$  = damping factor between 0 and 1

$$PR(p) = (1 - d) + d * \sum_{p_i \in L_{IN}(p)} \frac{PR(p_i)}{\text{amount } L_{OUT}(p_i)}$$

Calculate the missing iterations of the PageRank algorithm for the given example scenario with  $d = 0.75$ .



	0	1	2	3	4	5	PR
A	1		1,28125		1,1494140625		1,127166748
B	1		1,09375		1,19921875		1,1918029785
C	1		0,625		0,6513671875		0,6810302734



Last name:

First name:


Student number:

# Question 11)

Points: .....

Maximum points: 1+1+1+10+1=14

- a) What is the drawback of linear search in the Chord ring?
- b) What way of searching in the Chord ring is preferred?
- c) To which node  $n$  gets a key  $k$  assigned to?
  - Direct predecessor
  - First node (starting from ID 1) without any keys assigned yet
  - The node with the same ID as the key
  - Direct successor
- d) Calculate the Finger Table values of node  $n = 22$  and insert the correct values into the provided Finger Table.



Finger Table of node  $n = 22$

Entry	Start	Node
1		
2		
3		
4		
5		

The table has 5 entries, because  $m$  contains the length of the ID in bits and  $m = 5$

The Start value of entry  $i$  of the table on node  $n$  is  $(n + 2^{i-1}) \bmod 2^m$

The Node value of entry  $i$  points to the first node, which follows to  $n$  at a distance of at least  $2^{i-1}$

- e) Which node is responsible for the key (resource) with ID 11 ?

Last name:

First name:

Student number:

---

## Question 12)

Points: .....

Maximum points: 5

*Only a single answer is correct in each subquestion.*

- a) Centralized services exist in . . .  
 Centralized P2P       Pure P2P       Hybrid P2P
- b) No centralized services exist with . . .  
 Centralized P2P       Pure P2P       Hybrid P2P
- c) A central point of attack exists with . . .  
 Centralized P2P       Pure P2P       Hybrid P2P
- d) Which architecture causes the biggest network overhead?  
 Centralized P2P       Pure P2P       Hybrid P2P
- e) Which architecture causes the lowest network overhead?  
 Centralized P2P       Pure P2P       Hybrid P2P
- f) Which architecture implements a kind of dynamic, centralized service?  
 Centralized P2P       Pure P2P       Hybrid P2P
- g) Napster (1999 - 2001) implemented . . .  
 Centralized P2P       Pure P2P       Hybrid P2P
- h) Which architecture implements Ultrappeers (= Supernodes)?  
 Centralized P2P       Pure P2P       Hybrid P2P
- i) Gnutella v0.4 implements . . .  
 Centralized P2P       Pure P2P       Hybrid P2P
- j) Gnutella v0.6 implements . . .  
 Centralized P2P       Pure P2P       Hybrid P2P

