Prof. Dr. Carsten Vogt
Exercises „Operating Systems and Distributed Systems 2“
2019/20

Technology
Arts Sciences
TH Köln

Distributed File Systems – Solution

Exercise 1: File System Structures

- Given:
  - Four groups of files:
    - "database":
      db.exe (1.4 MB), data.dat (300 KB)
    - "prog":
      cc.exe (500 KB), p1.c (30 KB), p2.c (60 KB)
    - "inet":
      mail.exe (400 KB), ftp.exe (250 KB), www.exe (150 KB)
    - "texts":
      word.exe (750 KB), letter.doc (220 KB)
  - A file tree:
    - Main directory "users"
      with subdirectories "user1" and "user2" for the files of two users
  - Three (very small) file server nodes in the network:
    - node_A: disk size 1 MB
    - node_B: disk size 1.5 MB
    - node_C: disk size 2 MB
  - A client node (i.e. a fourth node in the network) from where users access the files
- Do the following:
  - Place the files on the nodes
    → Specify which file shall be stored on which server node.
    → Try to place all files of a group on the same node.

  Solution:
  - on node_A (1 MB): texts (970 KB total)
  - on node_B (1.5 MB): prog (590 KB total), inet (900 KB total)
  - on node_C (2 MB): database (1700 KB total)

- Sketch the file trees
  → One diagram for each of the three techniques, as seen in the lecture.
- Let user1 (who is working on the client node) access the file "p1.c"
  → For each of the three techniques: State the file path and name
  and (if necessary) the additional operations.

Solution:

Local File Trees:

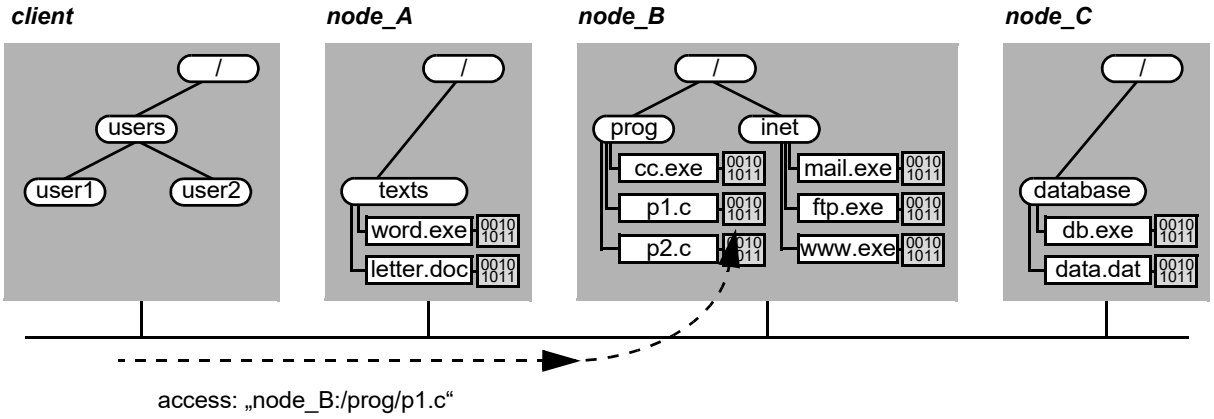

Mounting:

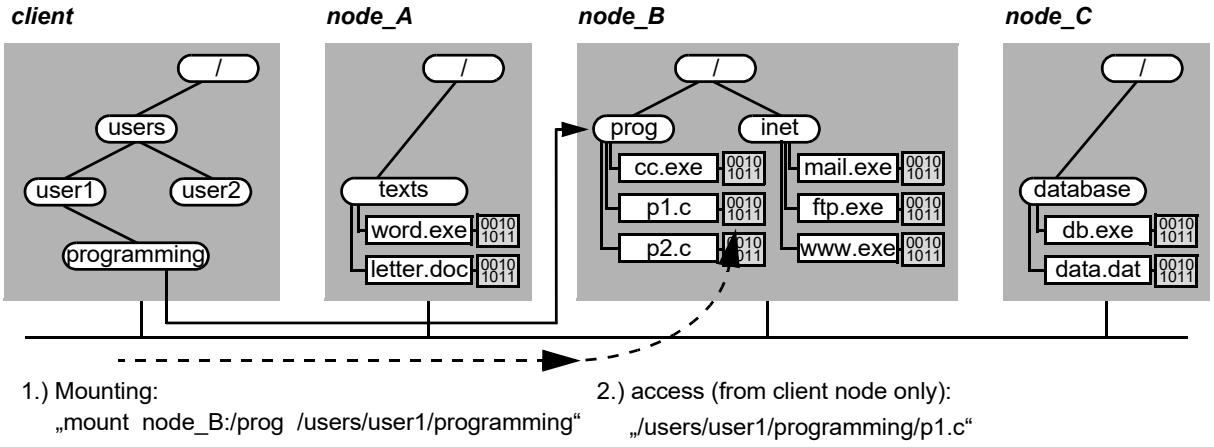

Solution (cont.):

Global File Tree:

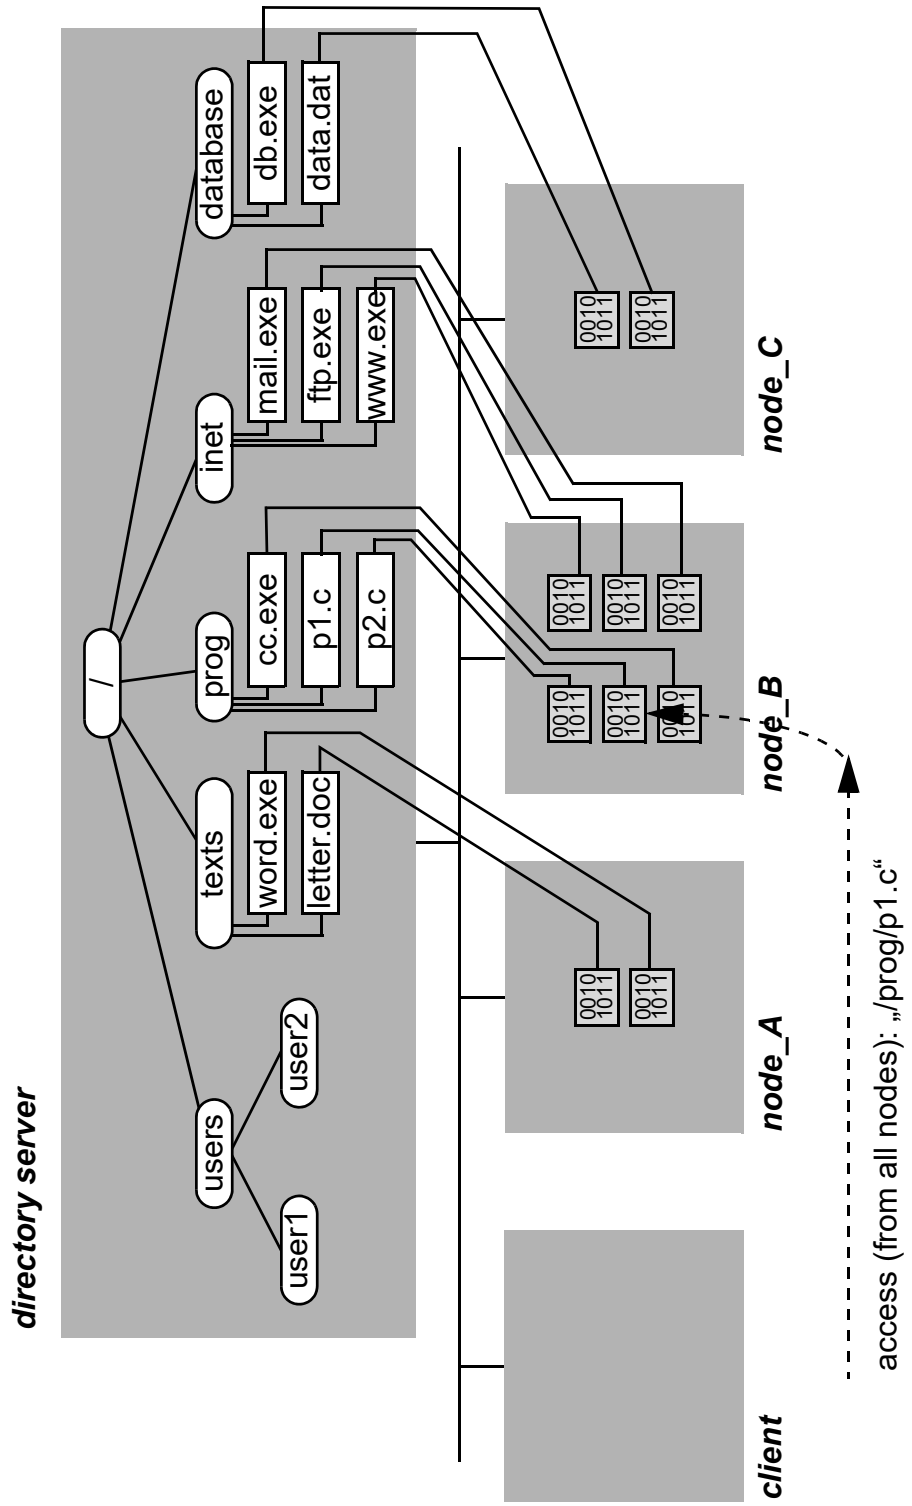

Exercise 2: Replication

- Given: A configuration of network nodes and disks with files

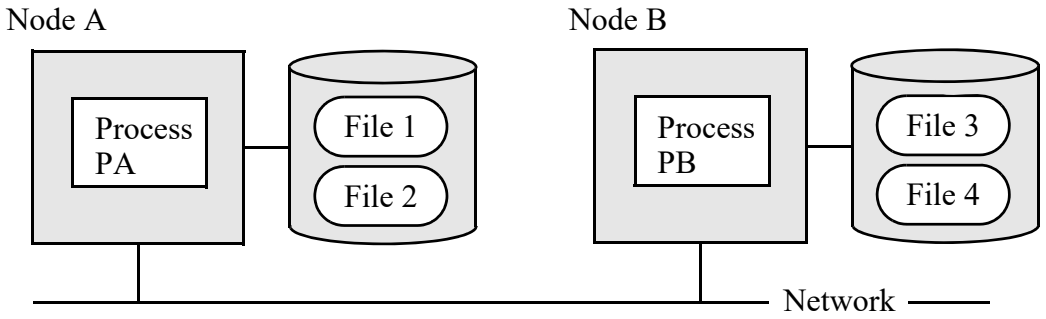

  - The number of file accesses per minute is given by the following table:

File accesses per minute	... to File 1	... to File 2	... to File 3	... to File 4
Reads from PA ...	20	5	4	3
Reads from PB ...	8	30	1	4
Writes from PA ...	2	3	2	18
Writes from PB ...	2	1	15	2

  - Assume that a local file access needs *L* milliseconds, a remote access *R* milliseconds. *L* and *R* are constants.

- Do the following:

  - Calculate the "total costs" *K* per minute, *K* being the sum of the durations of all file accesses in a minute:
    1.) Set up a formula that is based on a.) the symbols *L* and *R* and b.) the concrete number of accesses, as given by the table.
    2.) Use the formula to calculate the numerical costs *K* for *L*=15 ms and *R*=40 ms.

Solution:
K_{norep} = (20+5+1+4+2+3+15+2)*L + (4+3+8+30+2+18+2+1)*R = 52*L + 68*R
(Explanation: The reads/writes of PA to File 1 and 2 and of PB to File 3 and 4 are local. Therefore the corresponding numbers from the table have to be multiplied by L.
The other operations are remote. Therefore their numbers have to be multiplied by R.
The index "norep" stands for "no replication" to distinguish this formula from the following formulae.)
→ K_{norep} = 3500 ms für L=15 ms und R=40 ms

- Assume now that all files are replicated, i.e. each of the two disks stores copies of all four files. Again, set up a formula for the total costs and calculate its value for L=15 ms and R=40 ms. Have in mind that a read operation needs only to access the local disk but that a write operation must access both local and remote disk in order to keep the copies consistent. Local and remote accesses can be done in parallel, i.e. their total duration is the maximum of L and R.

Solution:
K_{allrep} =(20+5+4+3+8+30+1+4)*L + (2+3+2+18+2+1+15+2)*R = 75*L + 45*R
(Explanation: As stated above, all read operations are executed locally. Therefore their numbers are multiplied by L. All write operations require a local and a remote access that are executed in parallel. Assuming that R>L, their numbers are multiplied by R.)
$\rightarrow K_{allrep}$ = 2925 ms für L=15 ms und R=40 ms

- Does it pay off (for L=15 ms and R=40 ms) to replicate all four files? If so, why? If no, which file should better not be replicated?

Solution:
File 3 should better not be replicated because PB has many write accesses to File 3. If File 3 is not replicated to the other disk these accesses can be done locally. Otherwise, they must also be executed remotely in order to keep both copies consistent.
If File 3 is not replicated the total costs are reduced to 2650 ms (which is lower than the value K_{allrep} = 2925 ms from above).

- Assume now that you only have the alternative to replicate *all four* files or *none* of them. From which ratio R/L on does it pay off to replicate all files (based on the given table and the formulae from above)?

Solution:
Replication pays off if $K_{allrep} < K_{norep}$ or (equivalently for the numbers of the table) if 75*L + 45*R < 52*L + 68*R or (after simplification) R/L > 1. This will always be the case, as we can assume R > L.