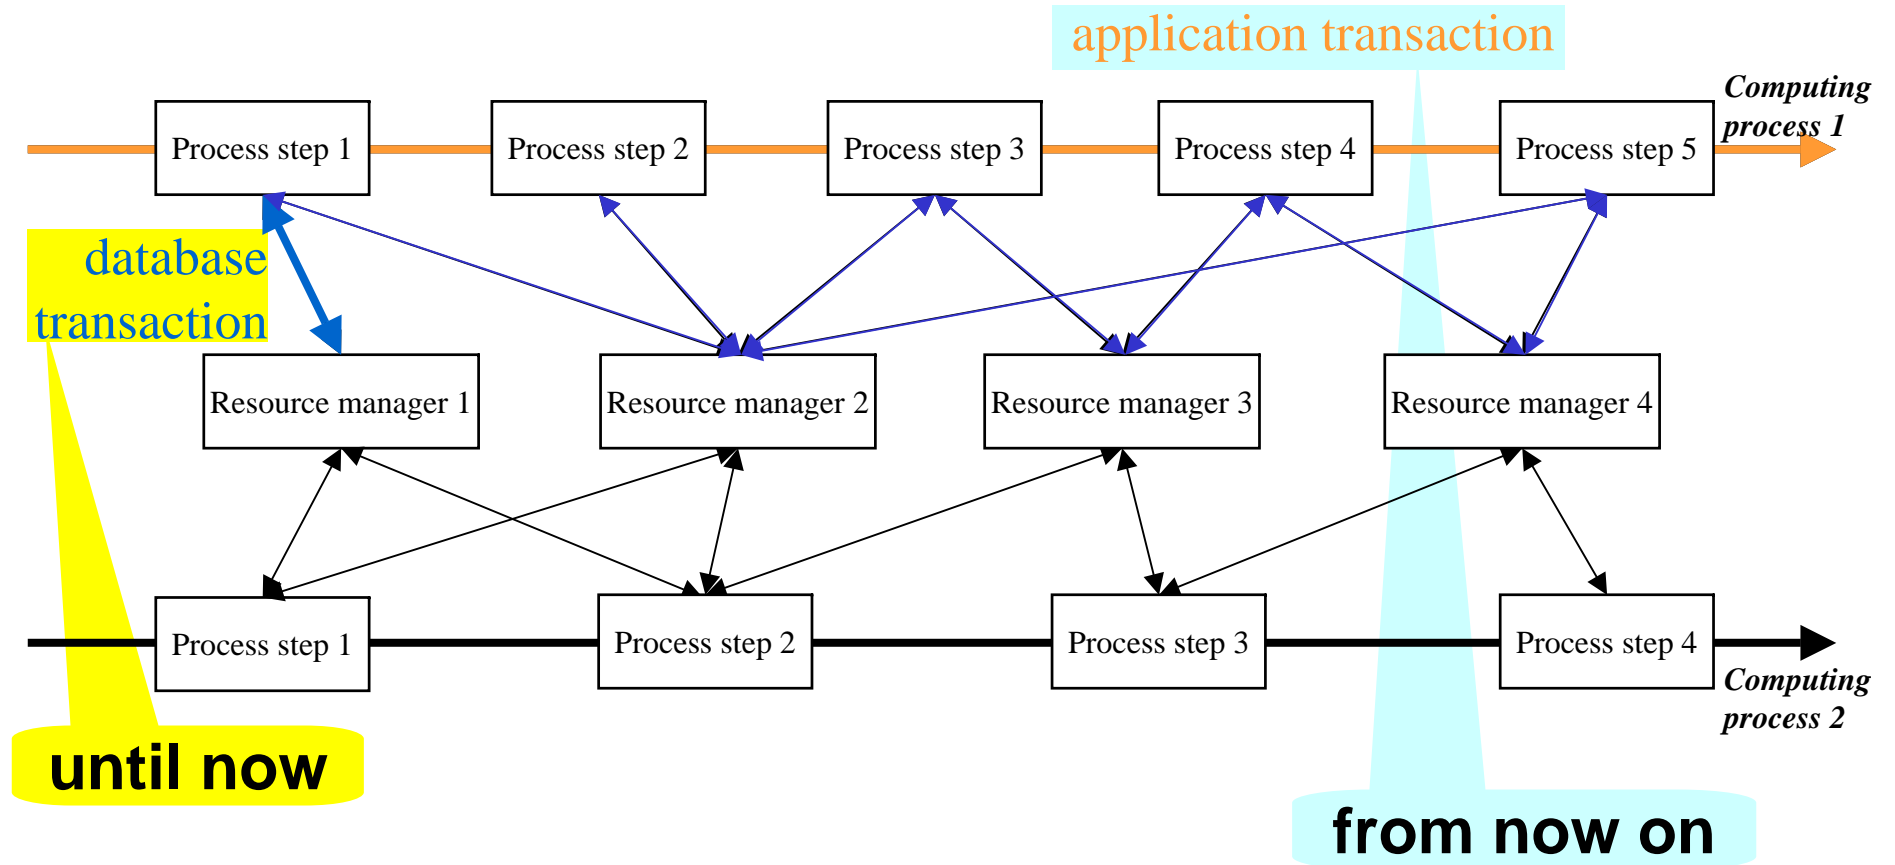


## Chapter 9


---

# Distributed Transactions: Characteristics

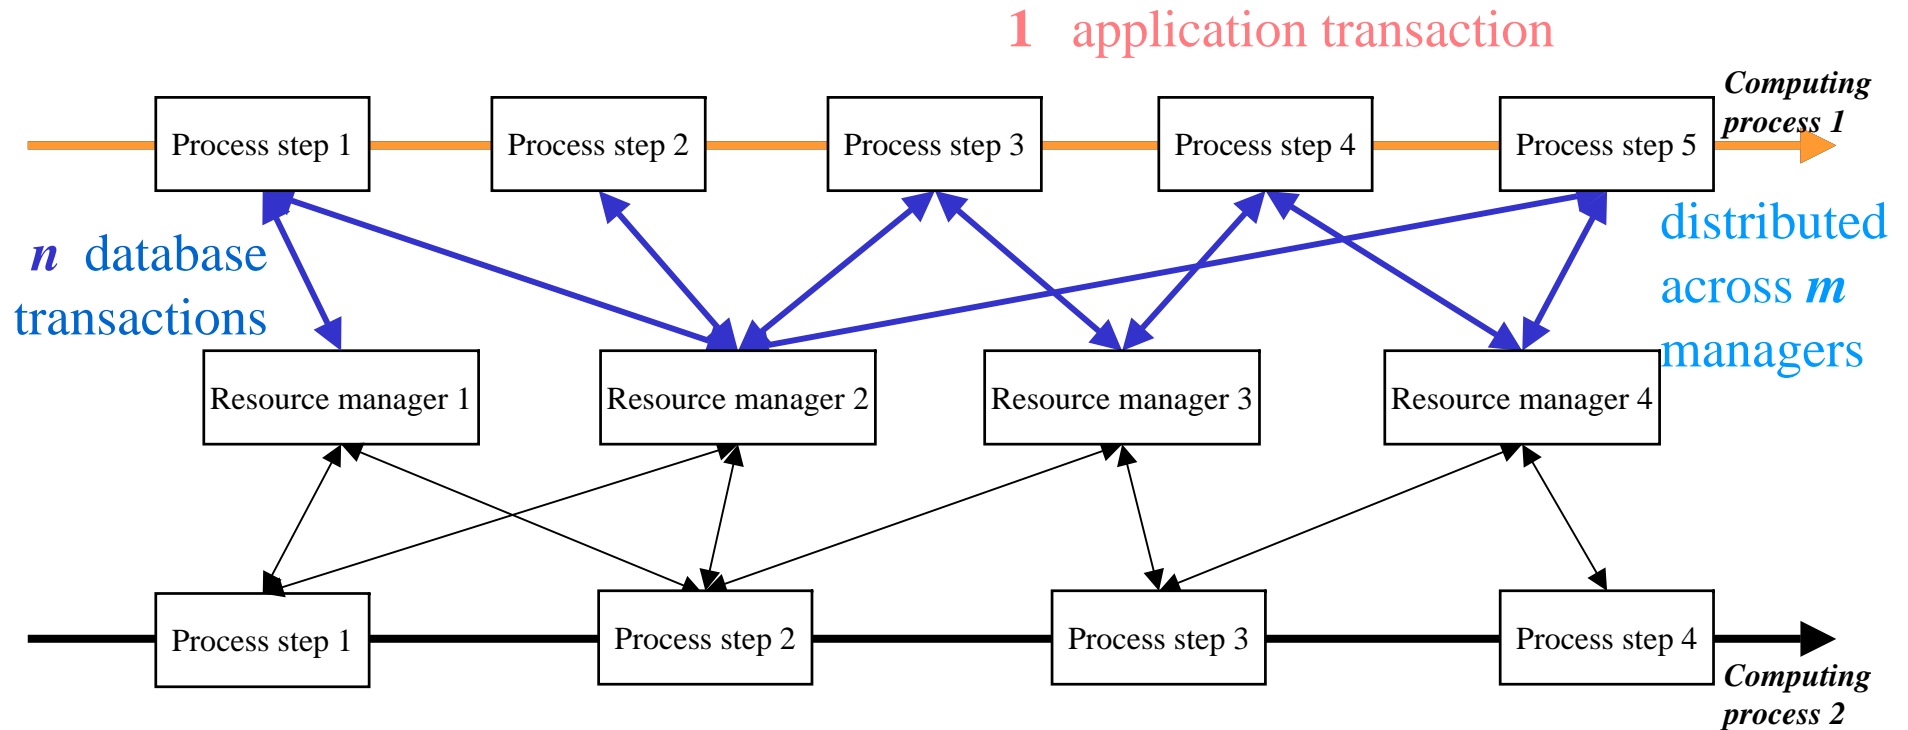
# Transactions



# Distributed transactions



# Complication: Autonomy



Modern complication: All servers – hence, all data servers – are autonomous!

# Autonomie

Jedes DBMS ist in *dreierlei* Hinsicht **autonom**:

- **Entwurfs-Autonomie**: Unabhängigkeit bzgl. verwendetem Datenmodell, logischem und physikalischem Entwurf, Anfragesprache, Protokollen für Nebenläufigkeitskontrolle und Recovery.
- **Ausführungs-Autonomie**: Lokale Transaktionen werden von den DBMS selbständig und ohne Kommunikation bearbeitet. Teile von globalen Transaktionen, die ein DBMS bearbeitet, werden wie lokale Transaktionen behandelt.
- **Kommunikations-Autonomie**: DBMS können mit anderen Systemen kommunizieren, können aber weder zur Kommunikation an sich, noch zu einem Zeitpunkt für die Kommunikation gezwungen werden.

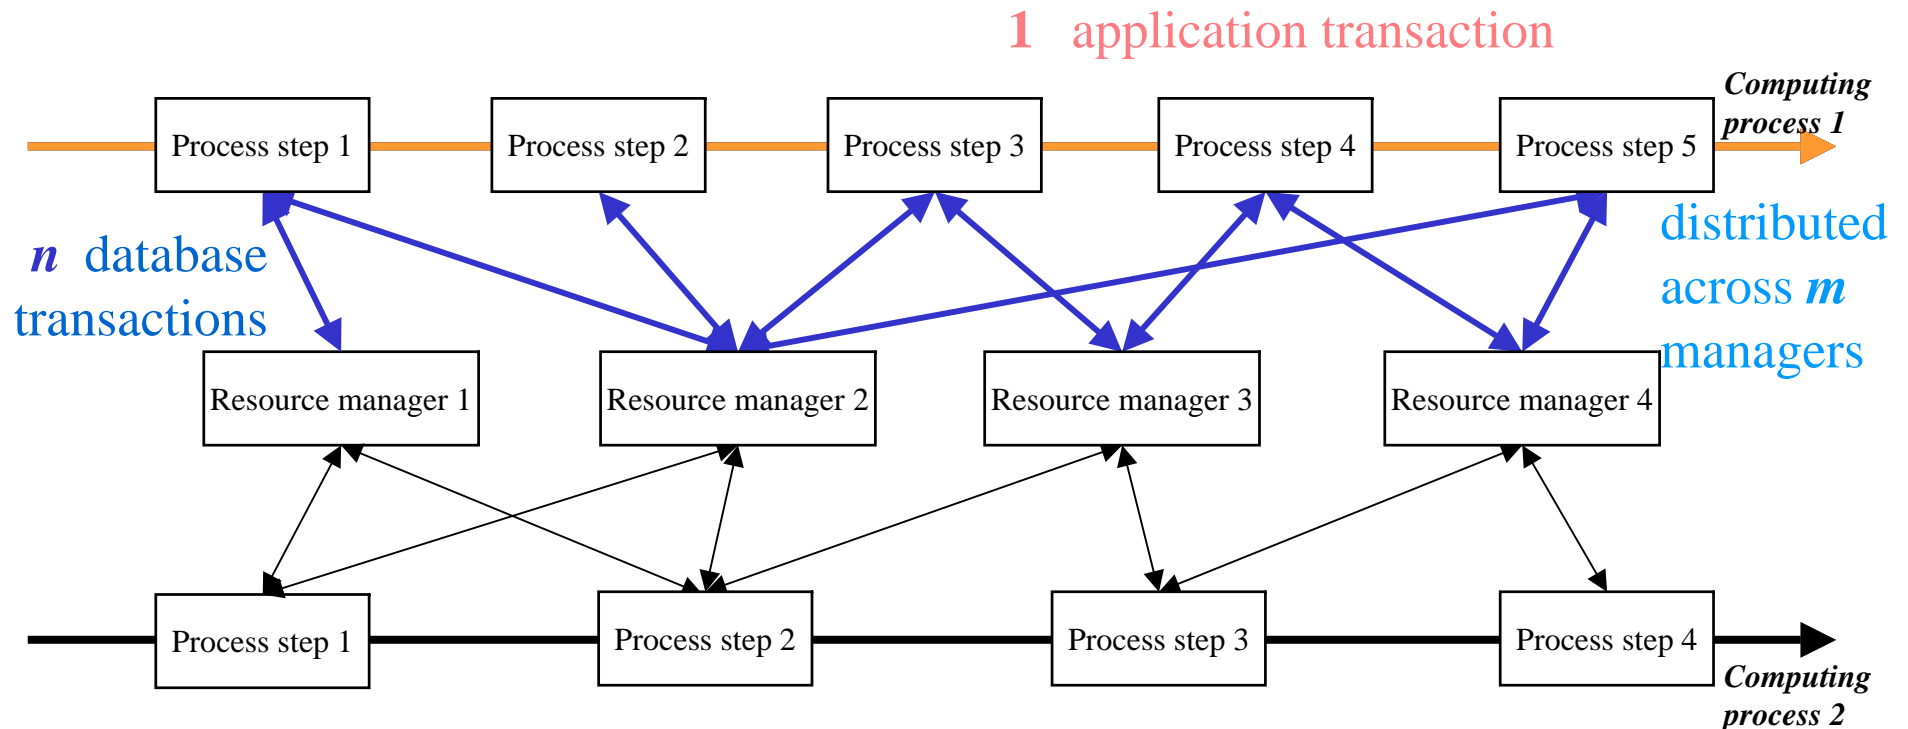
# DBMS Federations

**Modern focus:**

**Heterogeneous federation:** Participating servers

- are autonomous and independent;
- have no uniformity of protocols;
- their distribution is transparent to users.

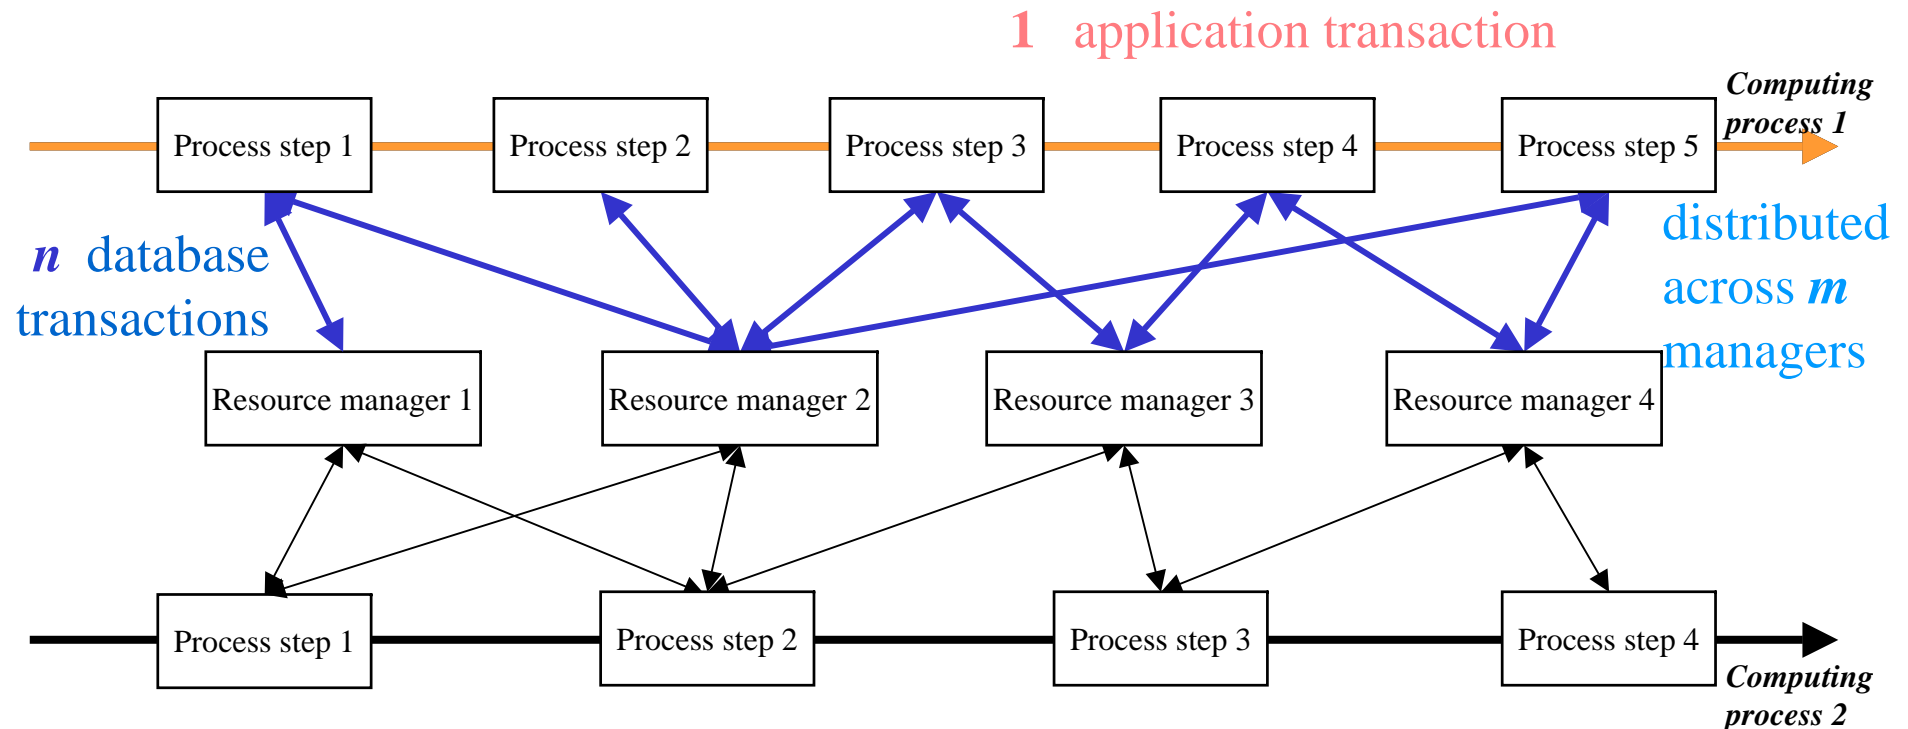
# Complication: More failure potential



Not only data servers may fail,  
but also communication links!

# Objective and problems (1)

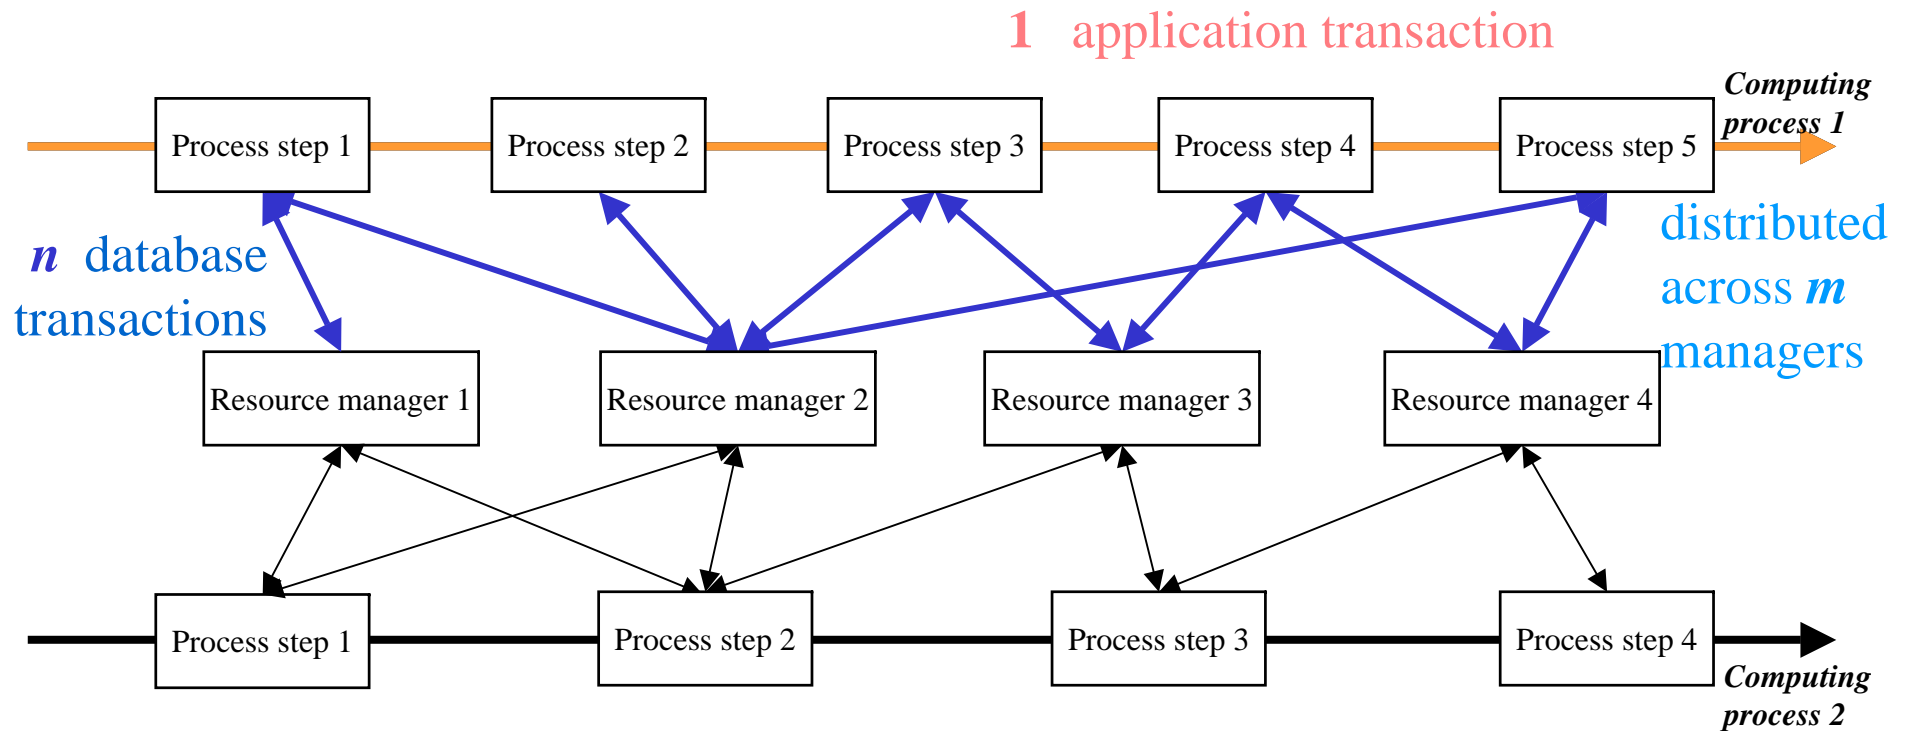
8



**Ideally:** Application transactions are ACID.

**Consistency problem 1:** Isolation of application transactions  $\Rightarrow$  Can we still enforce serializability?

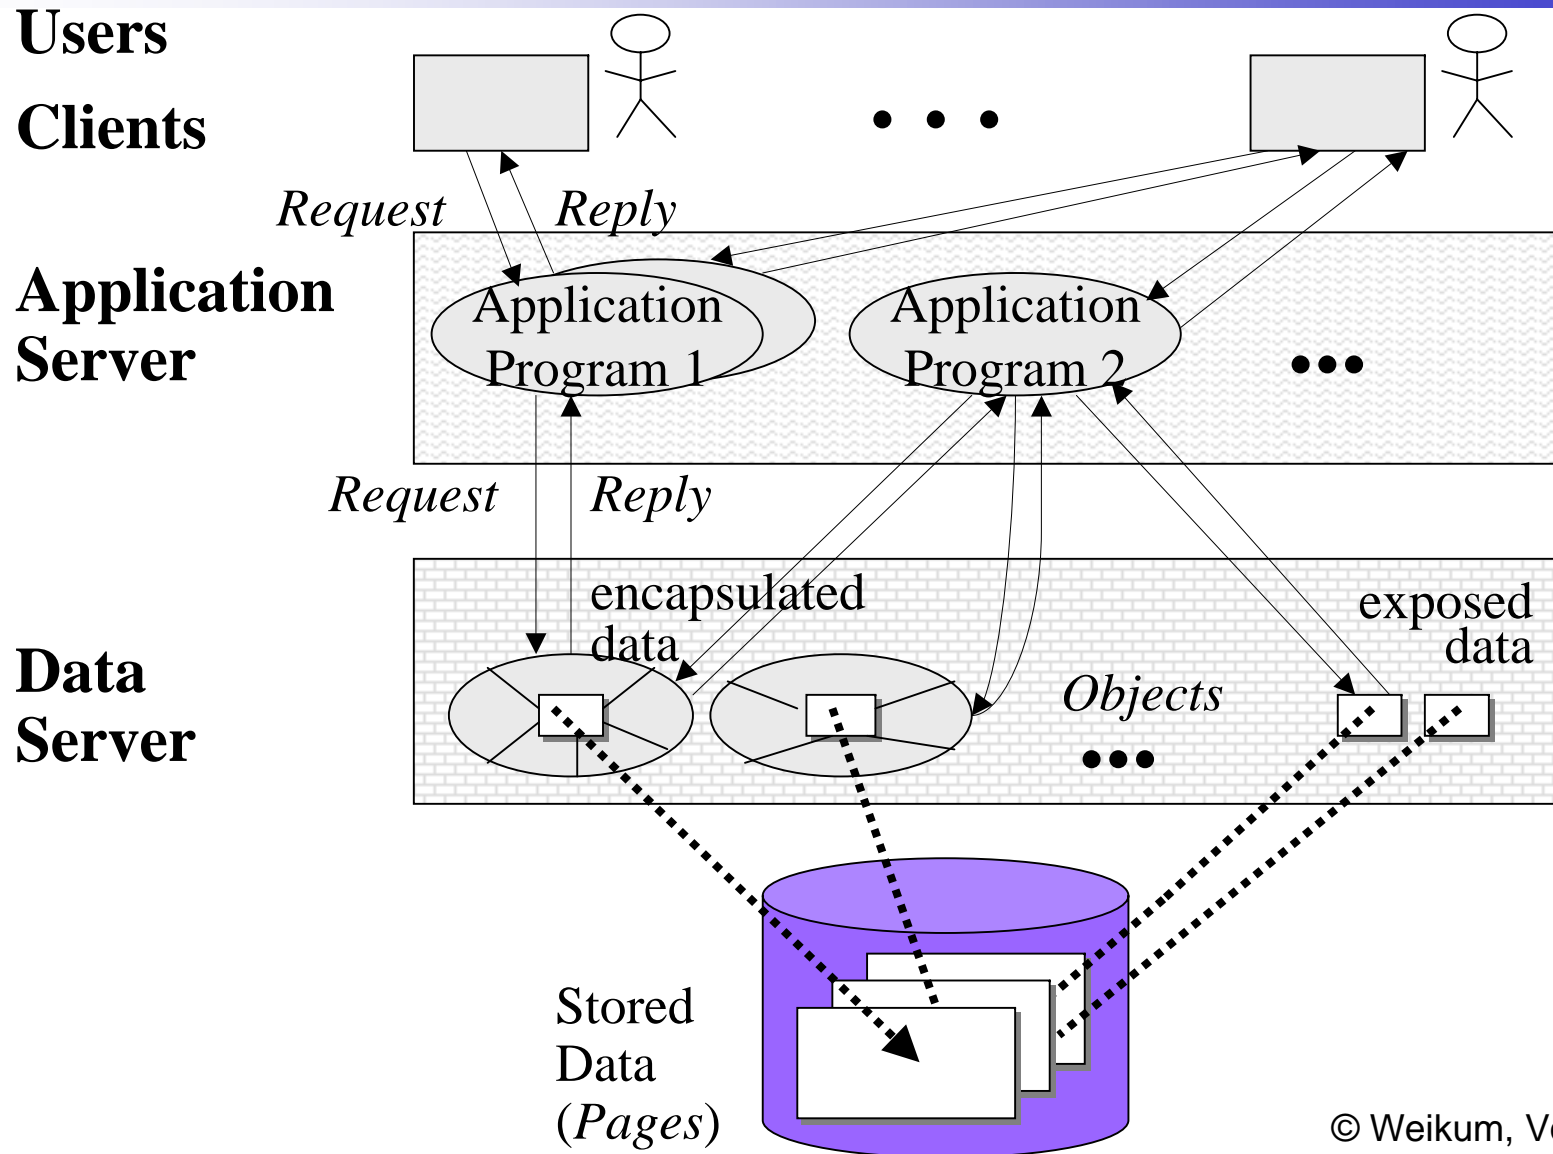
# Objective and problems (2)



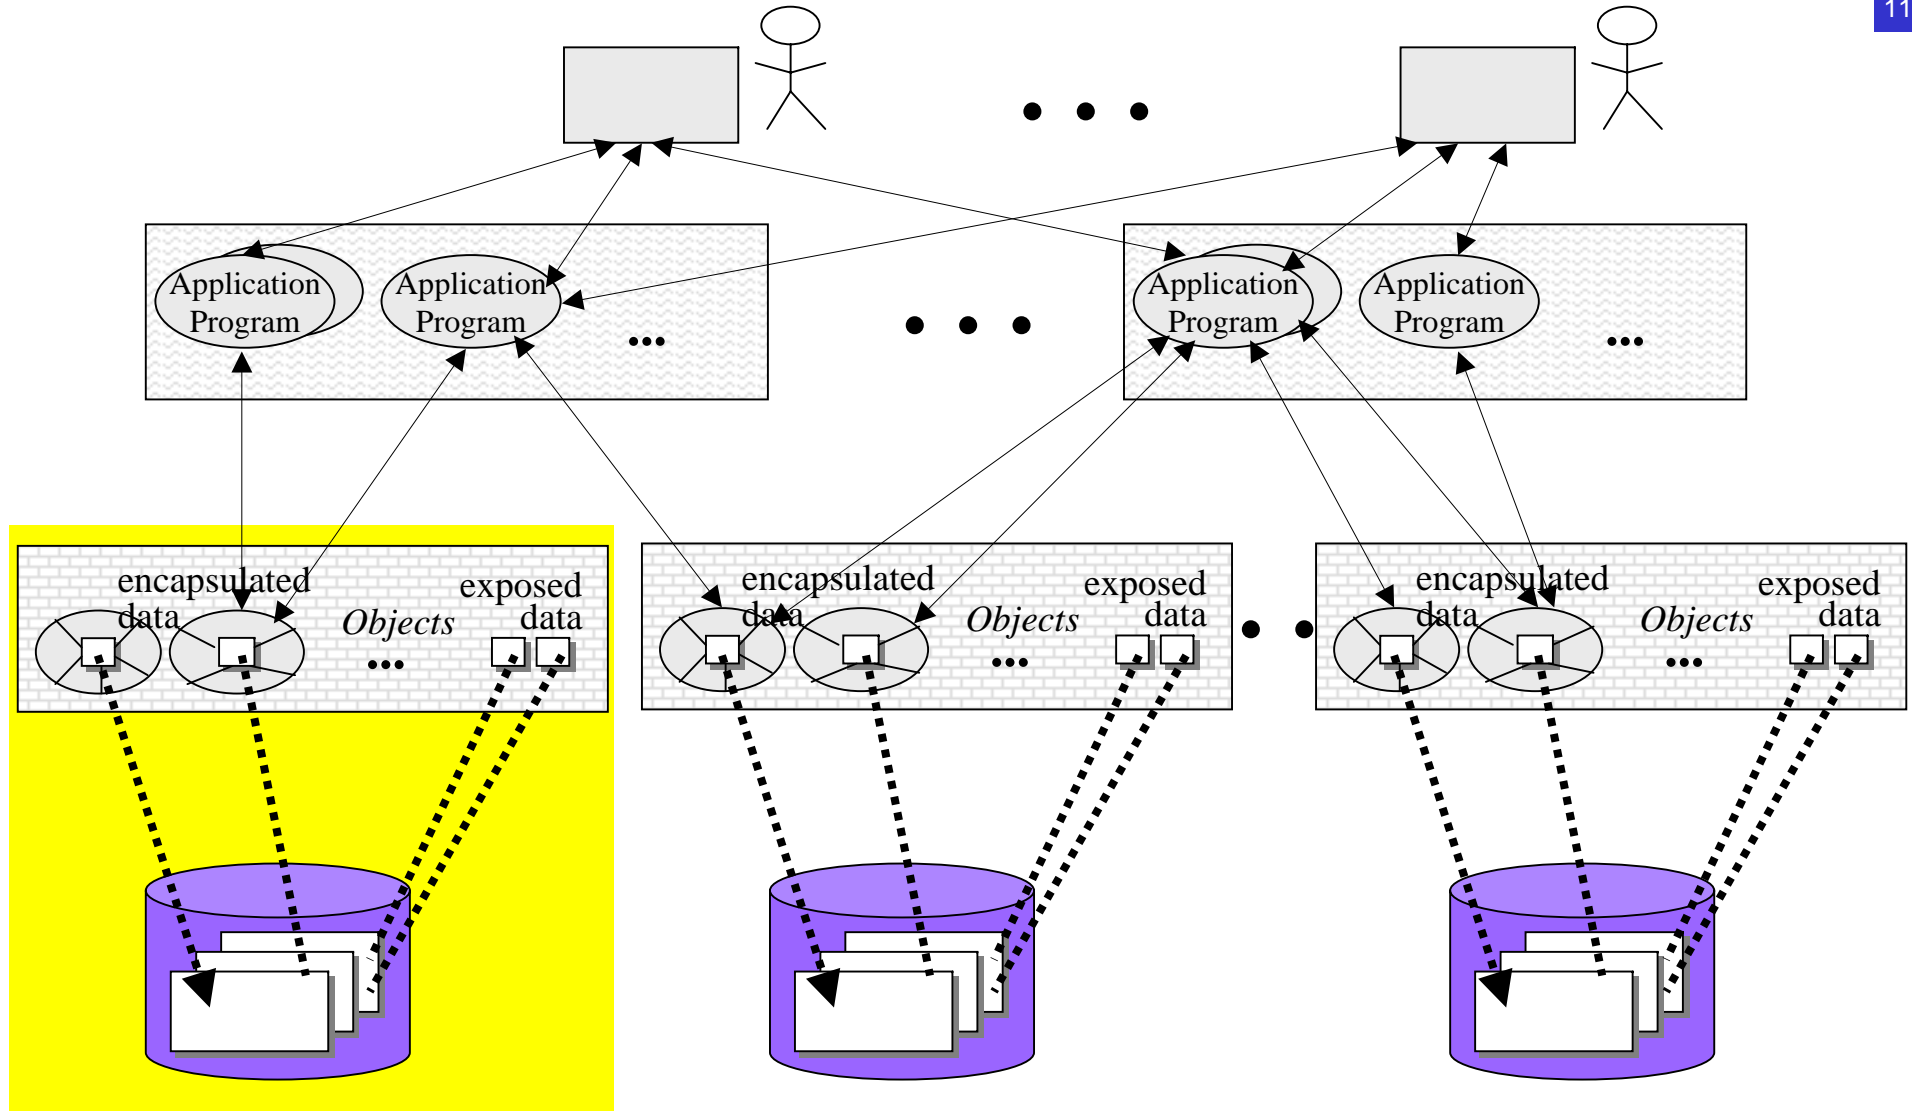
**Ideally:** Application transactions are ACID.

**Consistency problem 2:** Atomicity of application transactions  $\Rightarrow$  Can we still enforce all-or-nothing?

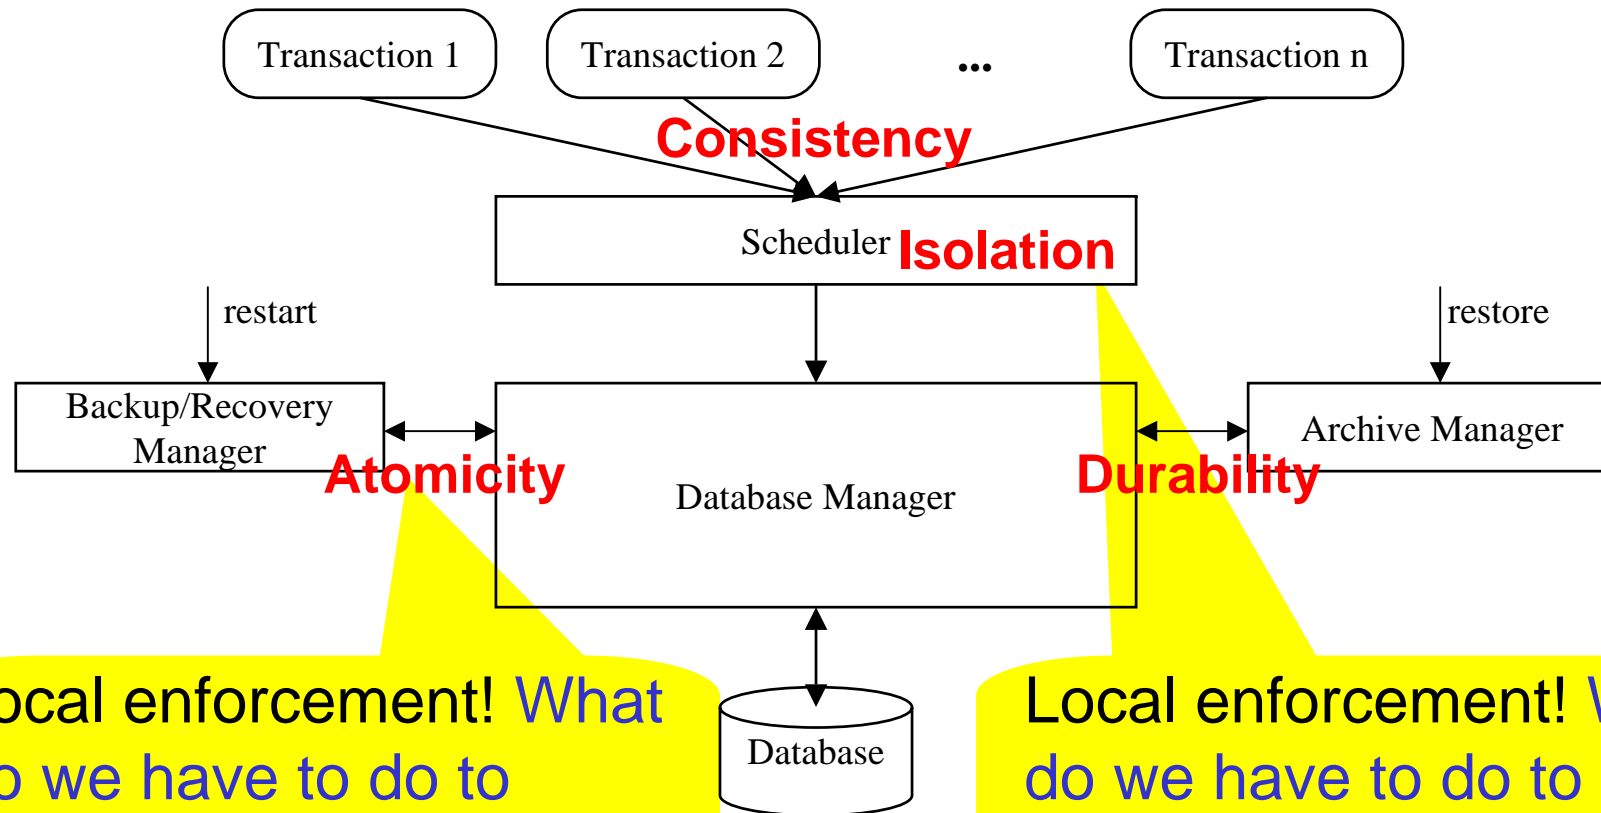
# 3-Tier Reference Architecture (1)



# 3-Tier Reference Architecture (2)



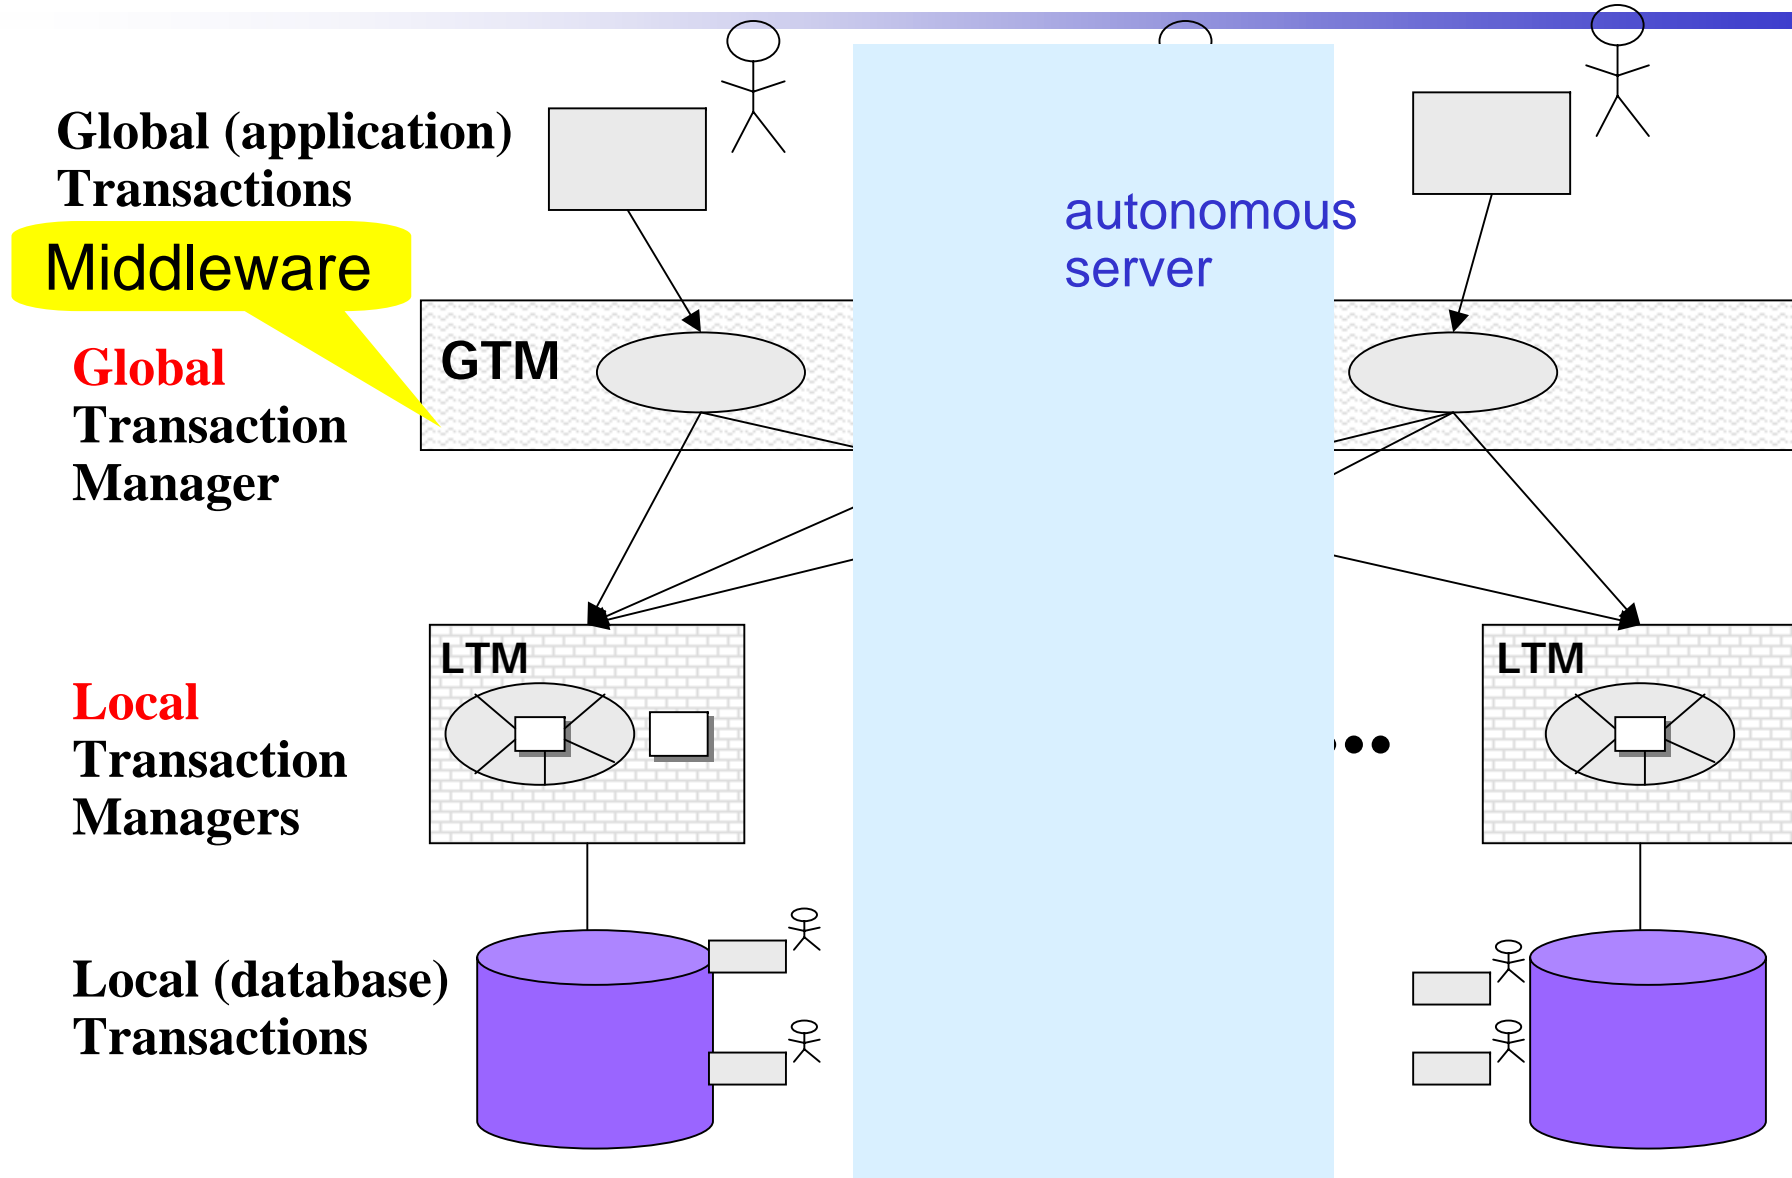
# Local server architecture



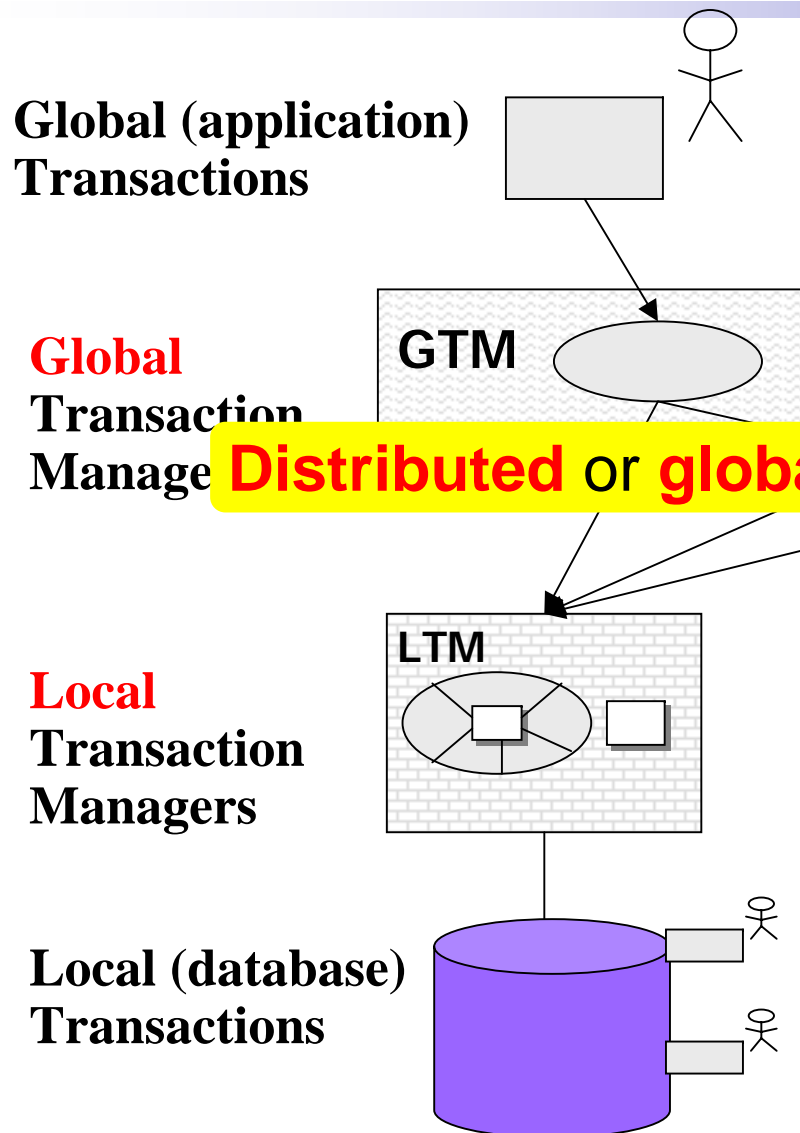
Local enforcement! What do we have to do to obtain global atomicity?

Local enforcement! What do we have to do to obtain global isolation?

# Federation Model



# Distributed Transaction Model (1)



**home/coordinator node:**

The node that starts the transaction, i.e., executes BOT. Issues all other DB operations including commit.

**Global Transaction Manager**

**Distributed or global transaction** across several nodes.

**Local Transaction Managers**

**Local transaction:** executed exclusively at one node (all data available there).

# Distributed Transaction Model (2)

- Local and global transactions are sequential (total order of DB operations).
- Each global transaction starts in each local DBMS at most one local (sub-)transaction.
- Global transactions invoke commit in each local DBMS where they are active.
- Replications of data are not recognized as such.

# Distributed Transaction Model (3)

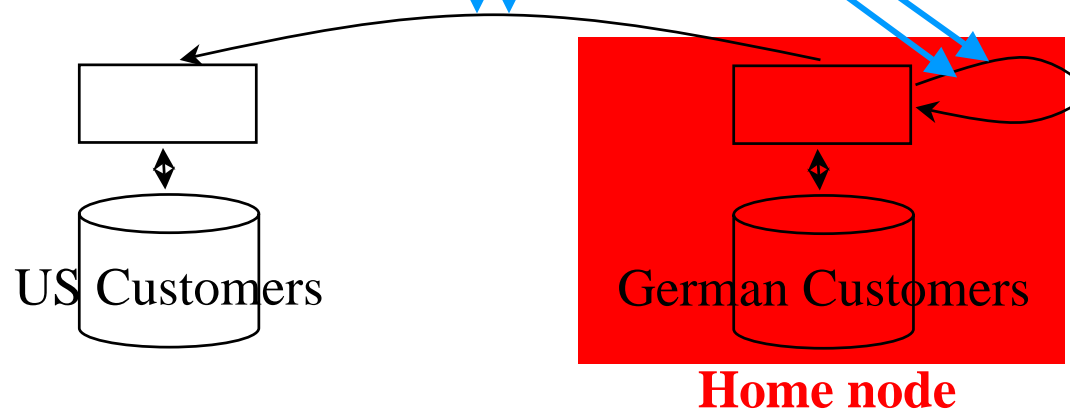
16

## Global transaction:

Transfer USD 500,--  
from Klemens' account to Jim's account.

- Increment Jim's balance.
- Record receipt of money from Klemens.
- Decrement Klemens' balance.
- Record transfer of money to Jim.

(local)  
subtransactions

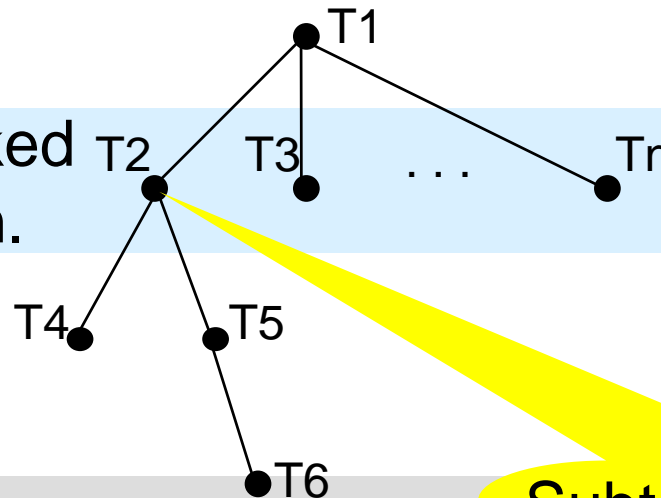


# Distributed Transaction Model (4)

17

**Primary transaction** (root transaction) – subtransaction executed at home node.

Subtransactions invoked by primary transaction.



Subtransaction invokes further subtransactions at other nodes.

## Transaction tree

- not balanced, height only limited by number of nodes,
- subtransactions may execute in parallel.

# Transaction execution

18

- Transaction starts from the home site, results are returned from there.
- If the LTM at the home site cannot execute an operation, GTM passes them on to the LTM of some other suitable node.
- Subtransaction may concurrently execute at the various nodes.

# Einfluss der Autonomie

19

- Es kann lediglich unterstellt werden, dass jedes lokale System für lokale Isolation und Atomizität sorgt, aber nicht wie.
- Der GTM hat keinen Einfluss auf die Ausführung lokaler Transaktionen. Er kann Kollisionen zwischen globalen und lokalen Transaktionen nicht verhindern.
- Subtransaktionen können in einem lokalen DBMS abstürzen und in anderen nicht.
- Der Abbruch lokaler Transaktionen wird außen nicht bemerkt.

- How to adapt scheduling protocols to maintain serializability?
- How to adapt recovery protocols to maintain atomicity?