Weiterentwicklung von DB2 Server für VSE & VM

IBM Software Group

Torsten Roeber

GSE Herbsttagung, September 2004

DB2 Information Management Software

e business on demand software

Agenda

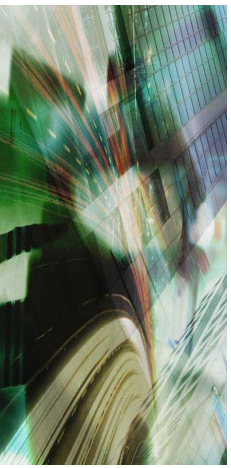

- DB2 Server for VSE & VM Version 7 Release 4
- Datenreplikation in der DB2 Familie
- Implementierung der ‚Q Replication' im VM & VSE
- Quo Vadis - DB2 Server for VSE & VM ?
- Weiter Informationen
- Haben Sie Fragen?

DB2 Server for VSE & VM Version 7 Release 4

- Verfügbar seit 17. September 2004
- Ein neues Kapitel der DB2 Datenreplikation
  - Transaktionsorientierte Replikation basierend auf MQSeries
  - Einschränkungen
    - Keine SQL-Replication mehr
    - VM nur MQ-Client
    - Kein Apply im VM & VSE

Replikation in der DB2 Familie

- Current SQL-Replication Architecture
- Why create another Replication Architecture?
- Q Replication
  - Q Subscription Process
  - Event Publication Process
  - Subscription Types
  - Q Apply Performance
  - Apply Load Options
  - Conflict Detection
  - Replication Administration

Current SQL Replication Architecture

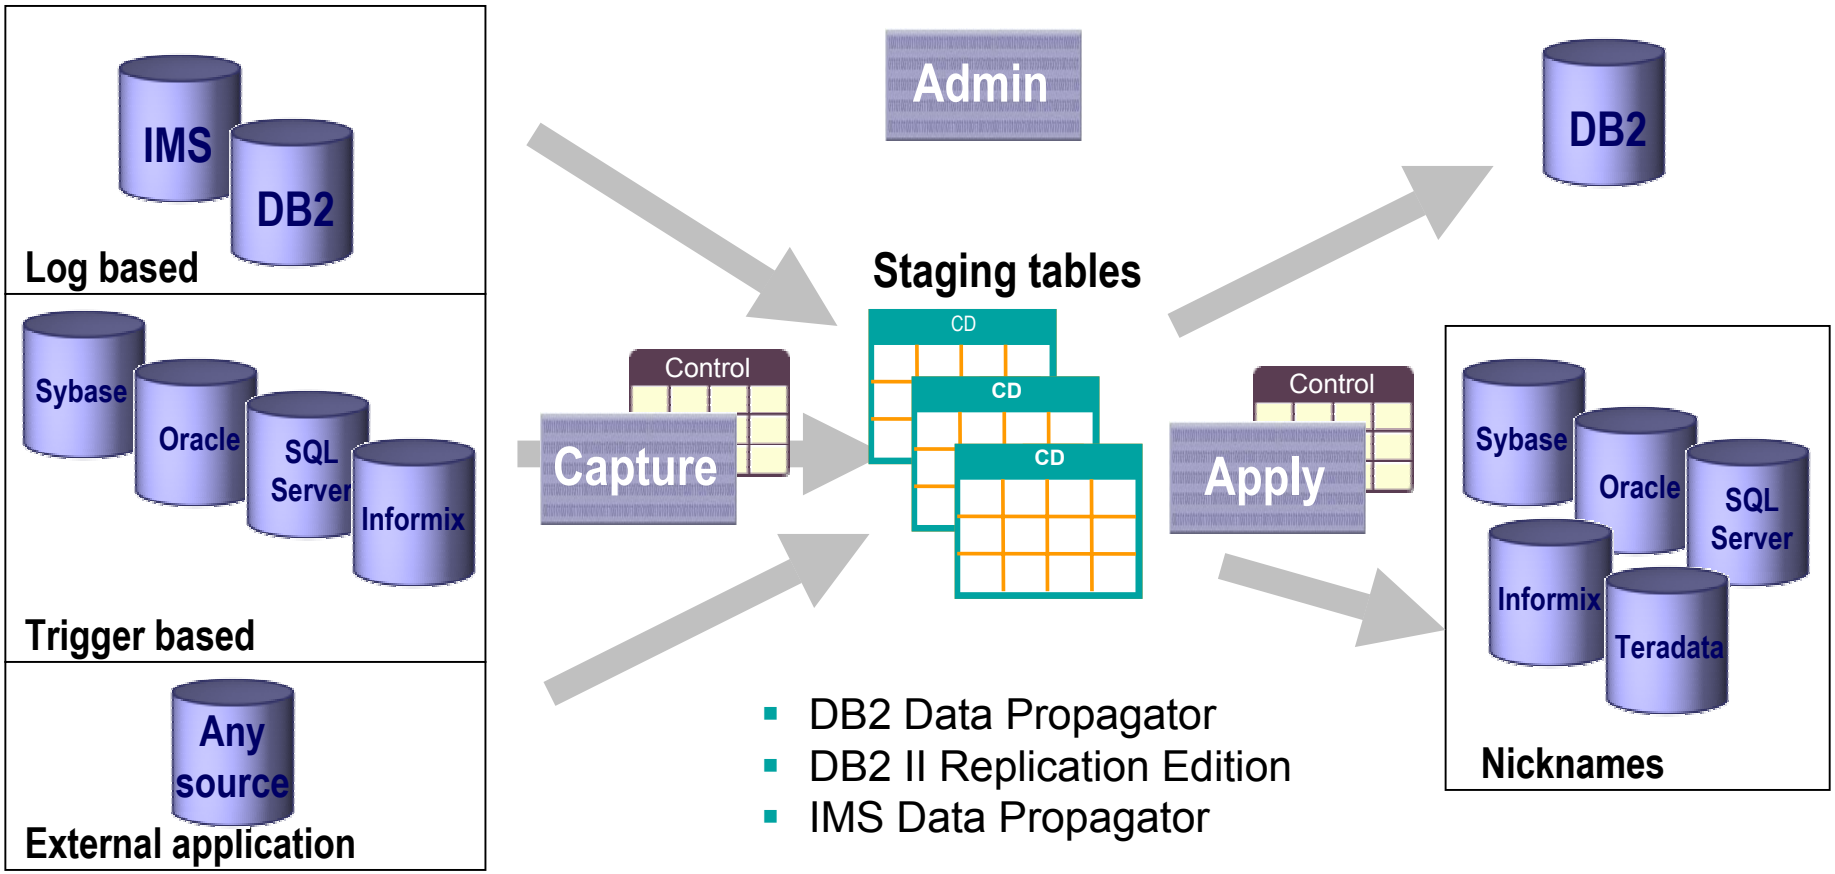

- DB2 Data Propagator
- DB2 II Replication Edition
- IMS Data Propagator

Why Create Another Replication Architecture?

- **Performance:** Combine high throughput with low latency

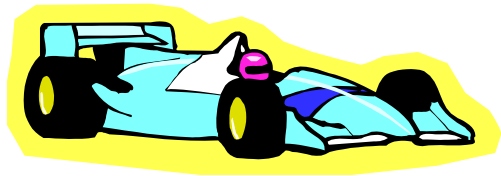

- **Capability:** Significantly improve multi-directional replication support

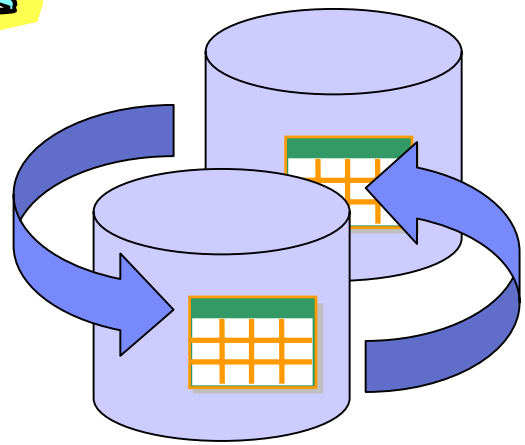

- **New function:** Event publishing, table difference utility

- **Manageability:** Reduce the number of replication objects to be defined and managed, ease the definition process with new Replication Center wizards

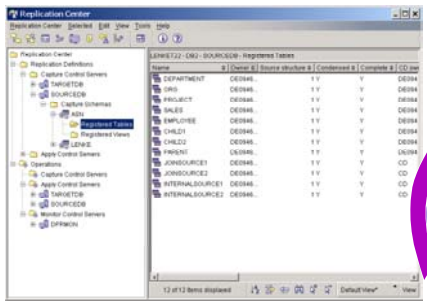

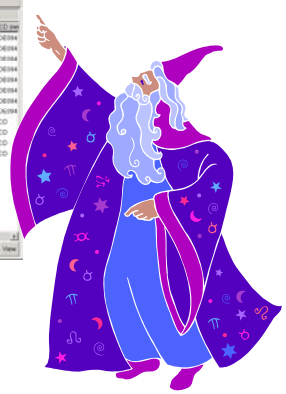

Q Replication – Q Subscription Process

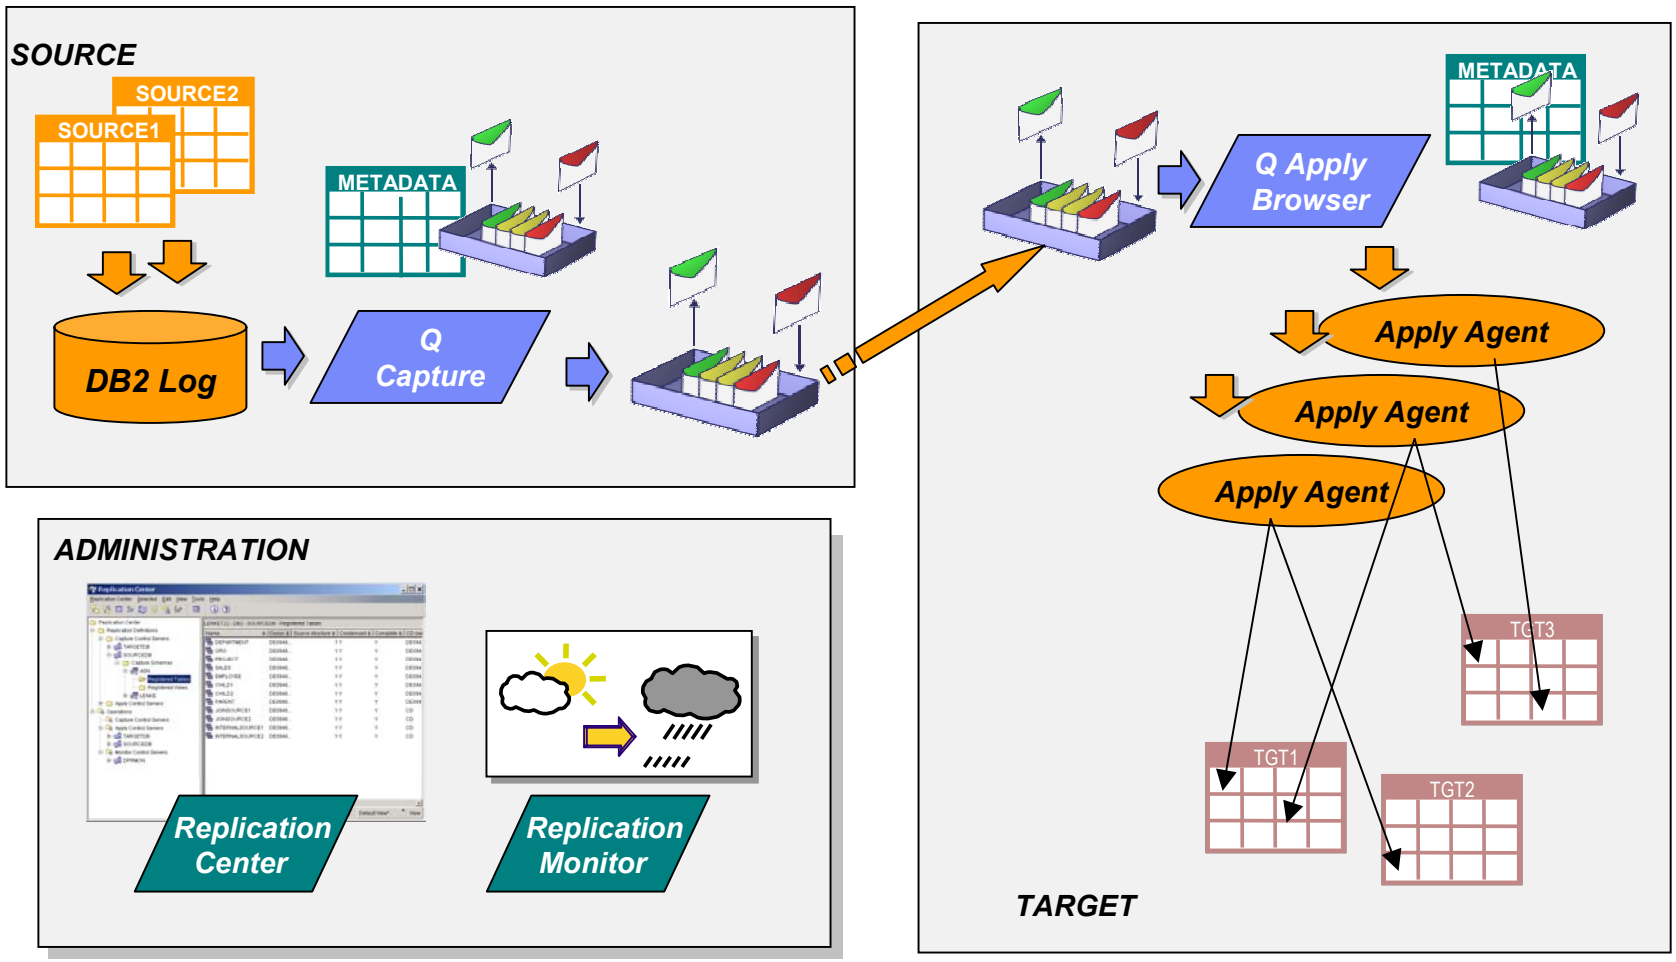

Q Replication – Event Publication Process

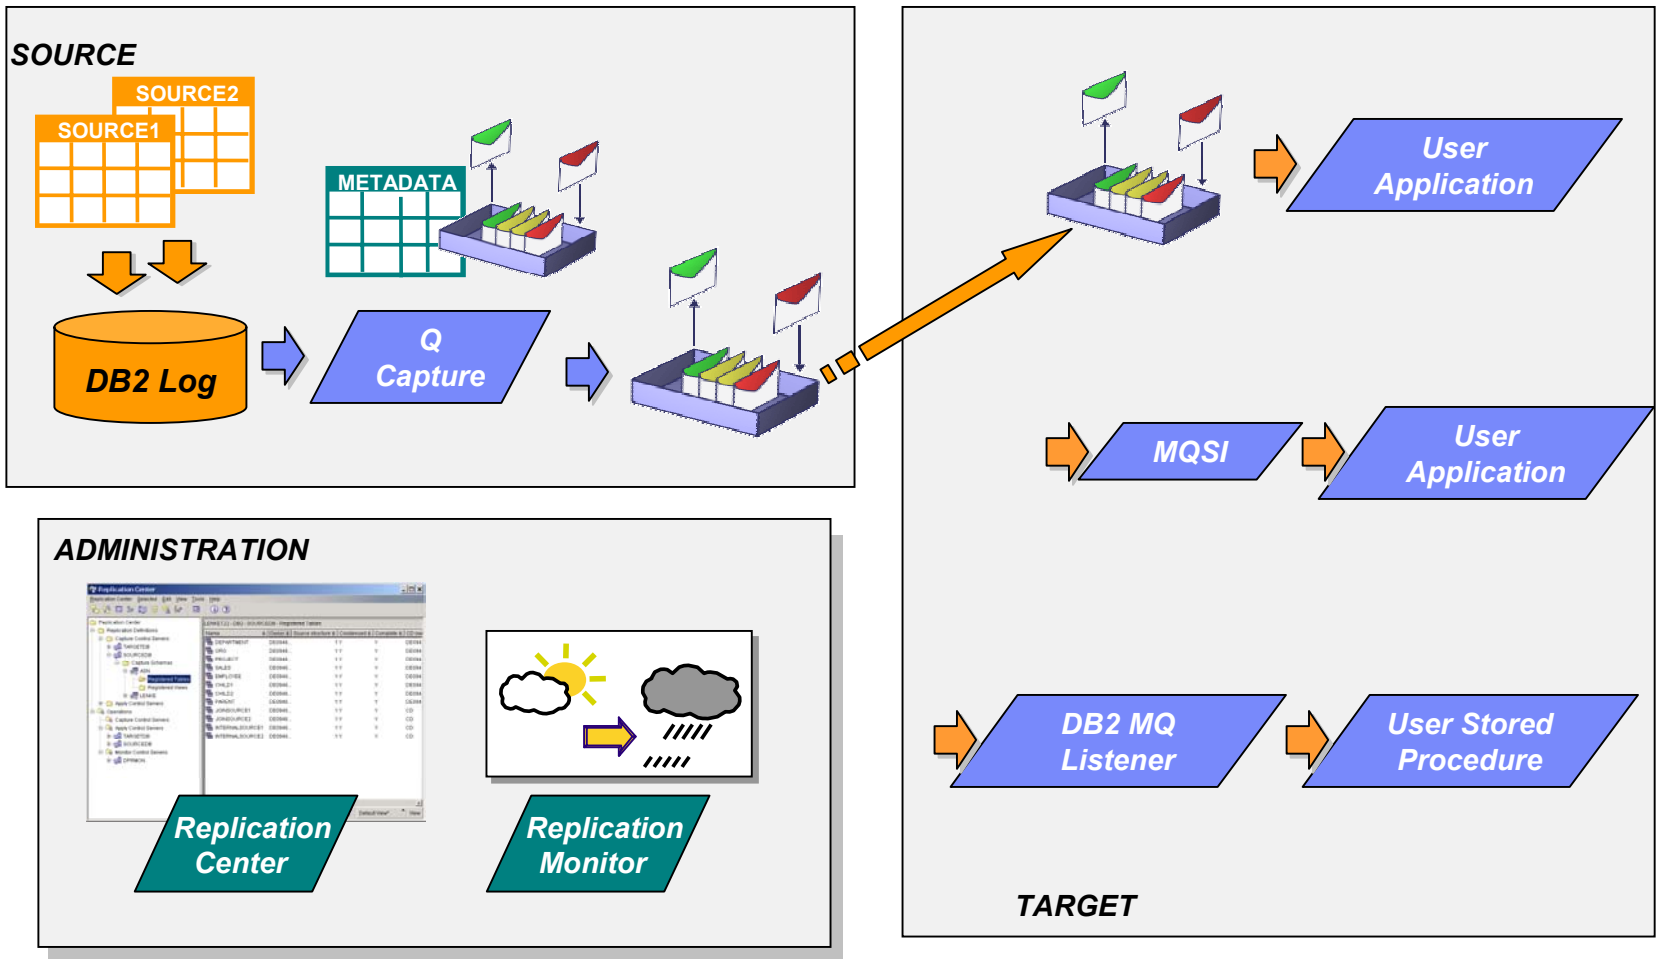

Subscription Types

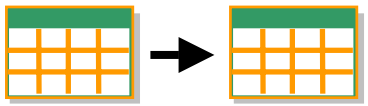

- **Unidirectional**
  - Changes are replicated in one direction between two servers (i.e. from source to target)
  - Changes can be filtered and transformed

- **Bidirectional**
  - Changes are replicated in two directions between two servers
  - Utilizes **VALUE** based conflict detection

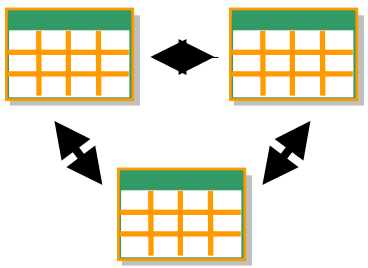

- **Peer to peer**
  - Changes are replicated between 2 or more servers
  - Utilizes **VERSION** based conflict detection

Q Apply Performance

- A Q Apply program will start one Apply browser for each receive queue defined
- Each Apply browser will read from its receive queue and apply transactions in parallel
  - Using as many agents as the user has defined
  - Performing dependency analysis and serializing transactions as necessary to preserve data integrity
  - Also possible to observe strict transactional order without parallelism
- It is the user's choice regarding how many queues to define
  - A high degree of parallelism can be reached using one queue
  - All related tables should be subscribed to on the same queue (using the same replication queue map)
  - In some ways, a receive queue is similar to the SQL Replication subscription set

Apply Load Options

- A subscription is defined as either: automatic load, manual load, no load required
- Automatic load:
  - Load is performed by Apply, with automatic coordination of the simultaneous capture of changes, loading of the new table, and apply of changes to other tables.
- Manual load:
  - Load is performed by user, coordination is required, and will be handled by user (with some help from our administration).
- No load:
  - No loading required, no coordination required, can immediately capture and apply changes
  - Example: target system is built through backup/restore, with replication started from an inactive source

Apply - Load Processing Example

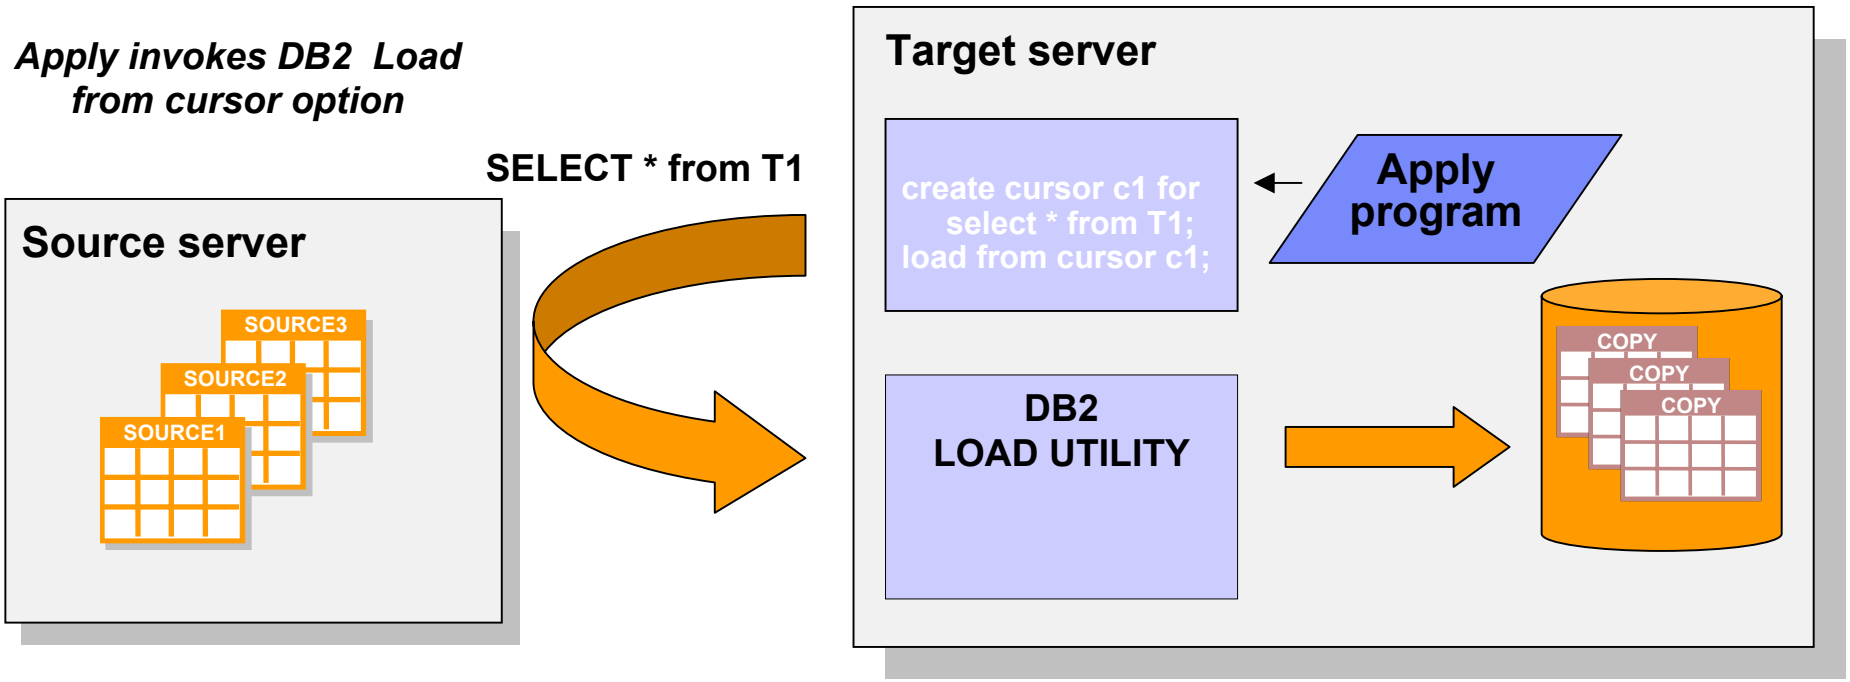

Selected data blocks feed directly into load utility

Subscription Start with Load Processing

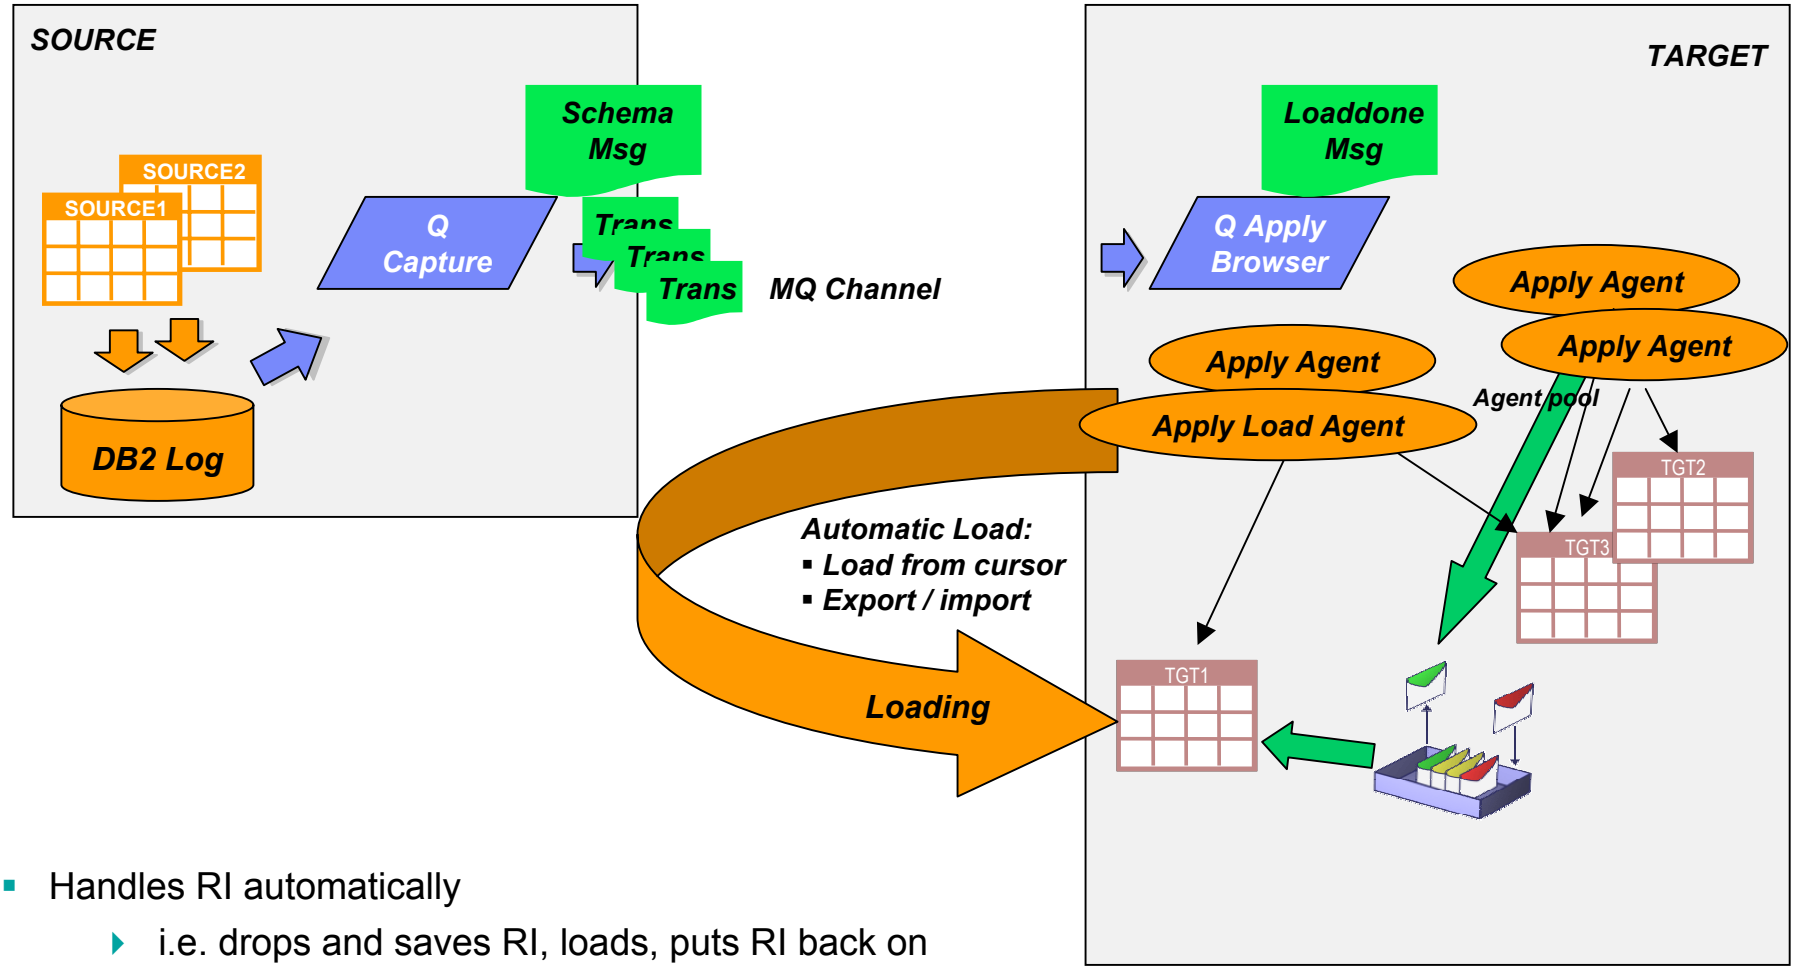

- Handles RI automatically
  - i.e. drops and saves RI, loads, puts RI back on
- Can handle Peer to Peer loading of n peers
  - Includes starting a new peer into an existing active configuration

Conflict Detection and Resolution

- Enables multi-directional replication that may result in conflicts
- Important for
  - “Active” standby systems
  - Workload balancing
- Value based conflict resolution
  - 2 participating nodes
  - Minimal overhead
- Version based conflict resolution
  - 2 or more participating nodes (practical limit around 6)
  - Requires extra columns and triggers
  - Most robust conflict detection and resolution (with some restrictions)

Value Based Conflict Detection

- Do the current row values at the apply target match the old row (before values) carried over from the source update?
- Designated site wins.

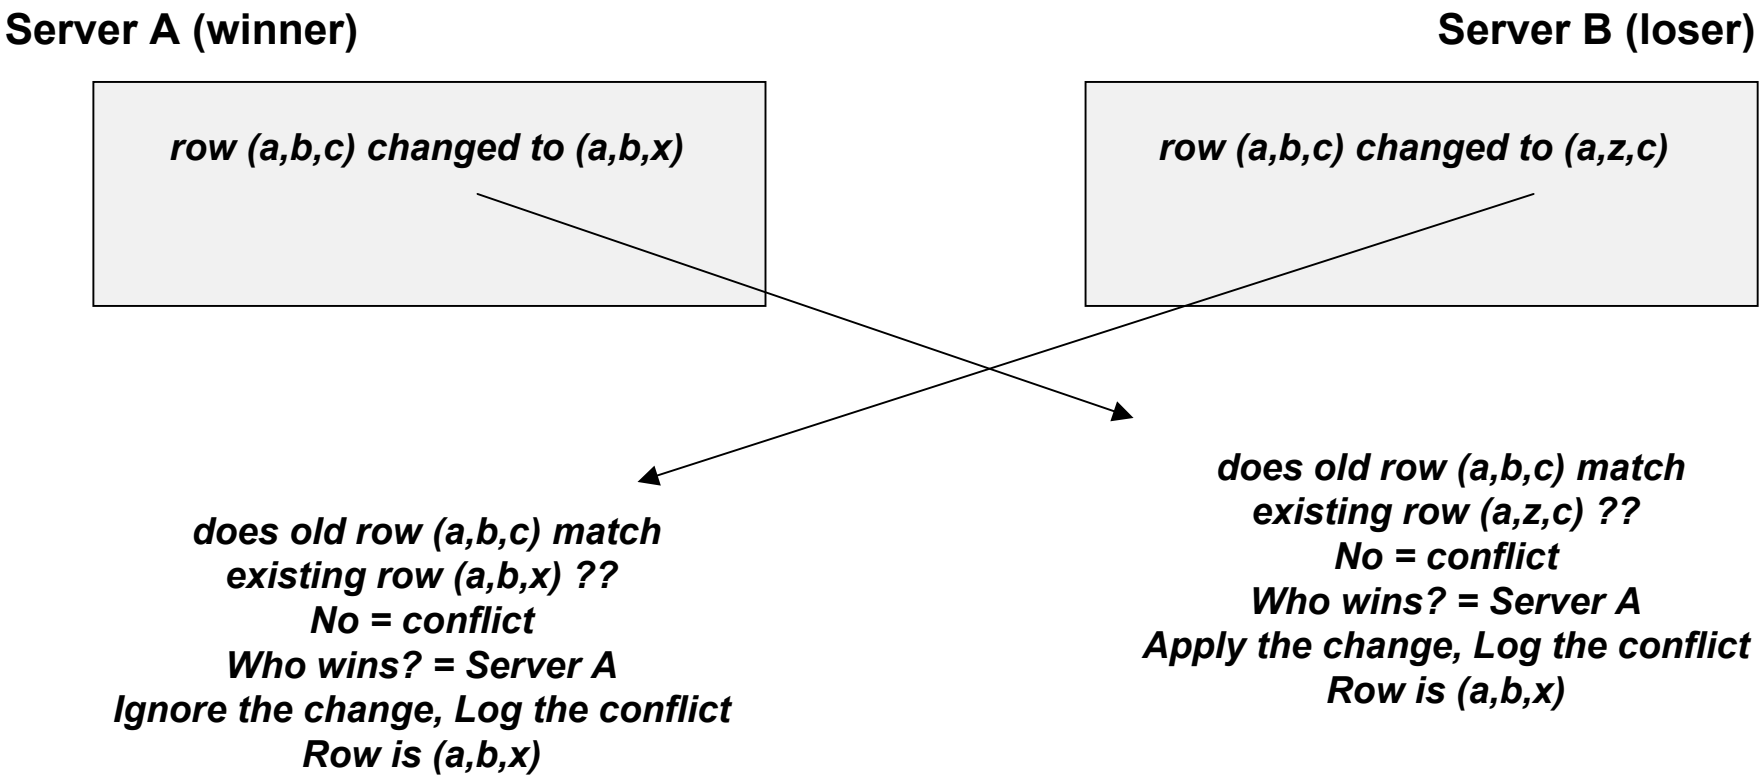

Value Based Conflict Detection and Resolution

- Conflict level options offered:
  - Check all columns on update — requires capture/transmit/apply of all old values
    - ✓ update target table where key values = < > and all current target values = all old source values
  - Check only changed columns on update - allows for merge
    - ✓ update target table where key values = < > and current target values = old source values for columns changed at source
  - No conflict checking

- Resolution choices offered: Force or Ignore
  - Force Action — requires capture/transmit/apply of all new values
    - ✓ force convergence on conflicts
    - ✓ log the conflict
  - Ignore Action
    - ✓ log the conflict
  - Force/Ignore used together in a pair provides "very good" convergence

Value Based Conflict Detection and Resolution

- Disadvantage: Does not detect all possible forms of conflict
  - Does not detect insert/insert+delete conflicts
  - Does not detect other fairly academic conflict cases
  - Is not offered for more than 2 participating database nodes
- Advantage: Requires less overhead
  - No extra columns or triggers
  - No effect to source updating applications
  - Problematic conflict cases may not be applicable to user applications
  - Can supplement with reconciliation utility (Tdiff/Trepair)
  - Might be appropriate for planned outage/failover/DR

Version Based Conflict Detection

- All rows are augmented with a “Version” = timestamp Tx and smallint Nx, indicating when and by which server the row was last updated
- Do the current values of Tx and Nx at the apply target match the old values of Tx and Nx carried over from the source update?
- Most current timestamp Tx wins.

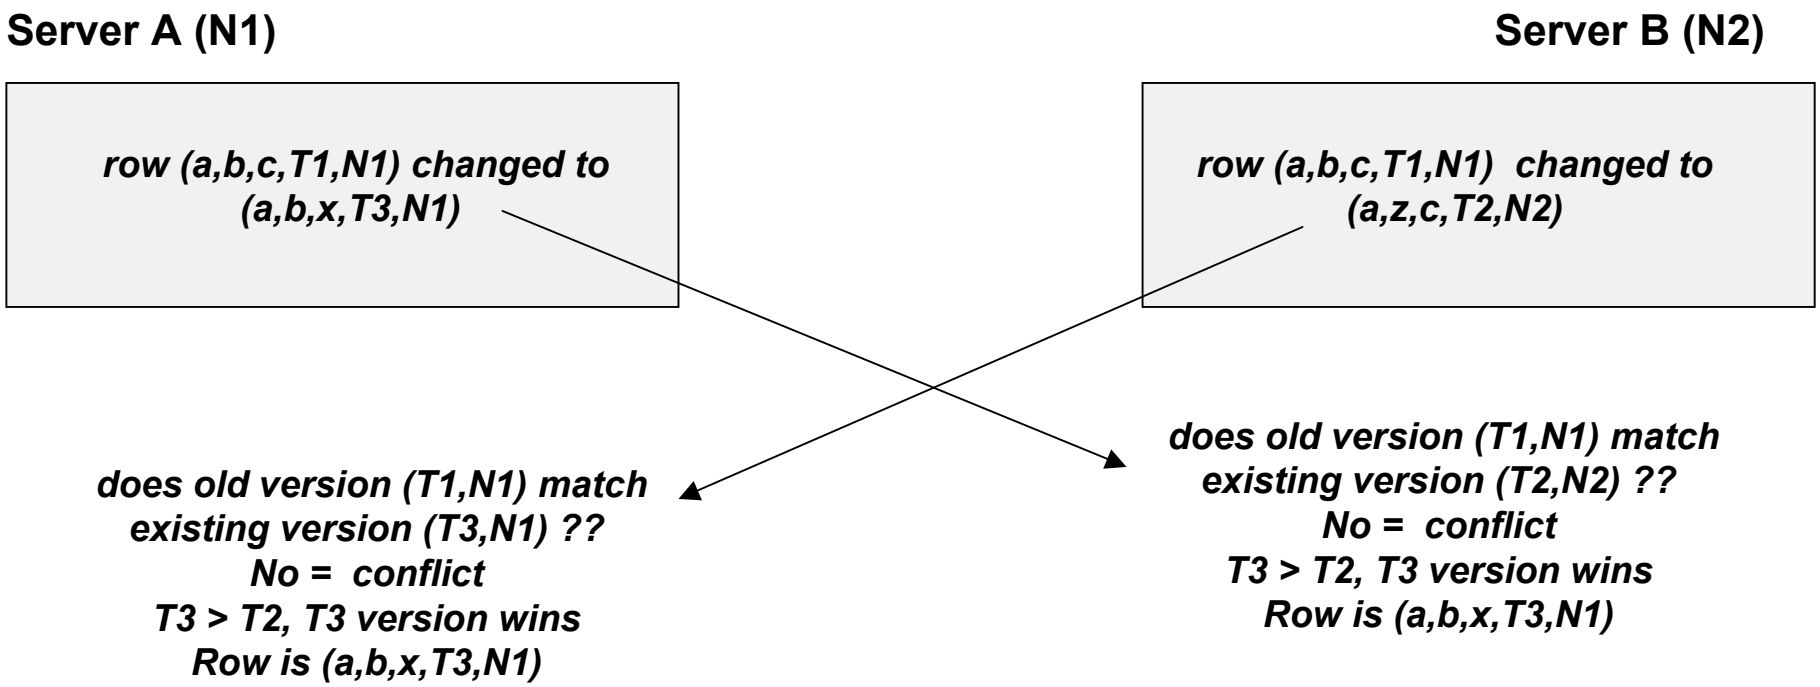

Replication Administration

- Replication Center GUI
  - Launchpads, Wizards, Online Help
  - Definitions, Operations, Monitoring

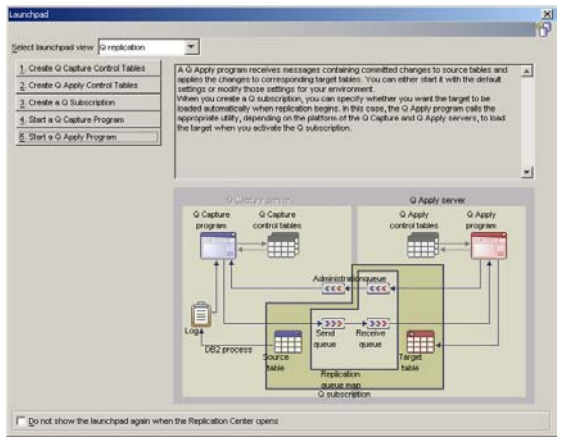

- Command Line Interface
  - Scripts or interactive mode
  - Example:

```
C:\asnclp
REPL > CREATE QSUB USING REPLQMAP ...
REPL > CREATE SUBSCRIPTION SET SETNAME ...
REPL > CREATE MEMBER IN SETNAME ...
```

- Java API's
  - Typically used when replication is embedded

Q Create Subscription Wizard

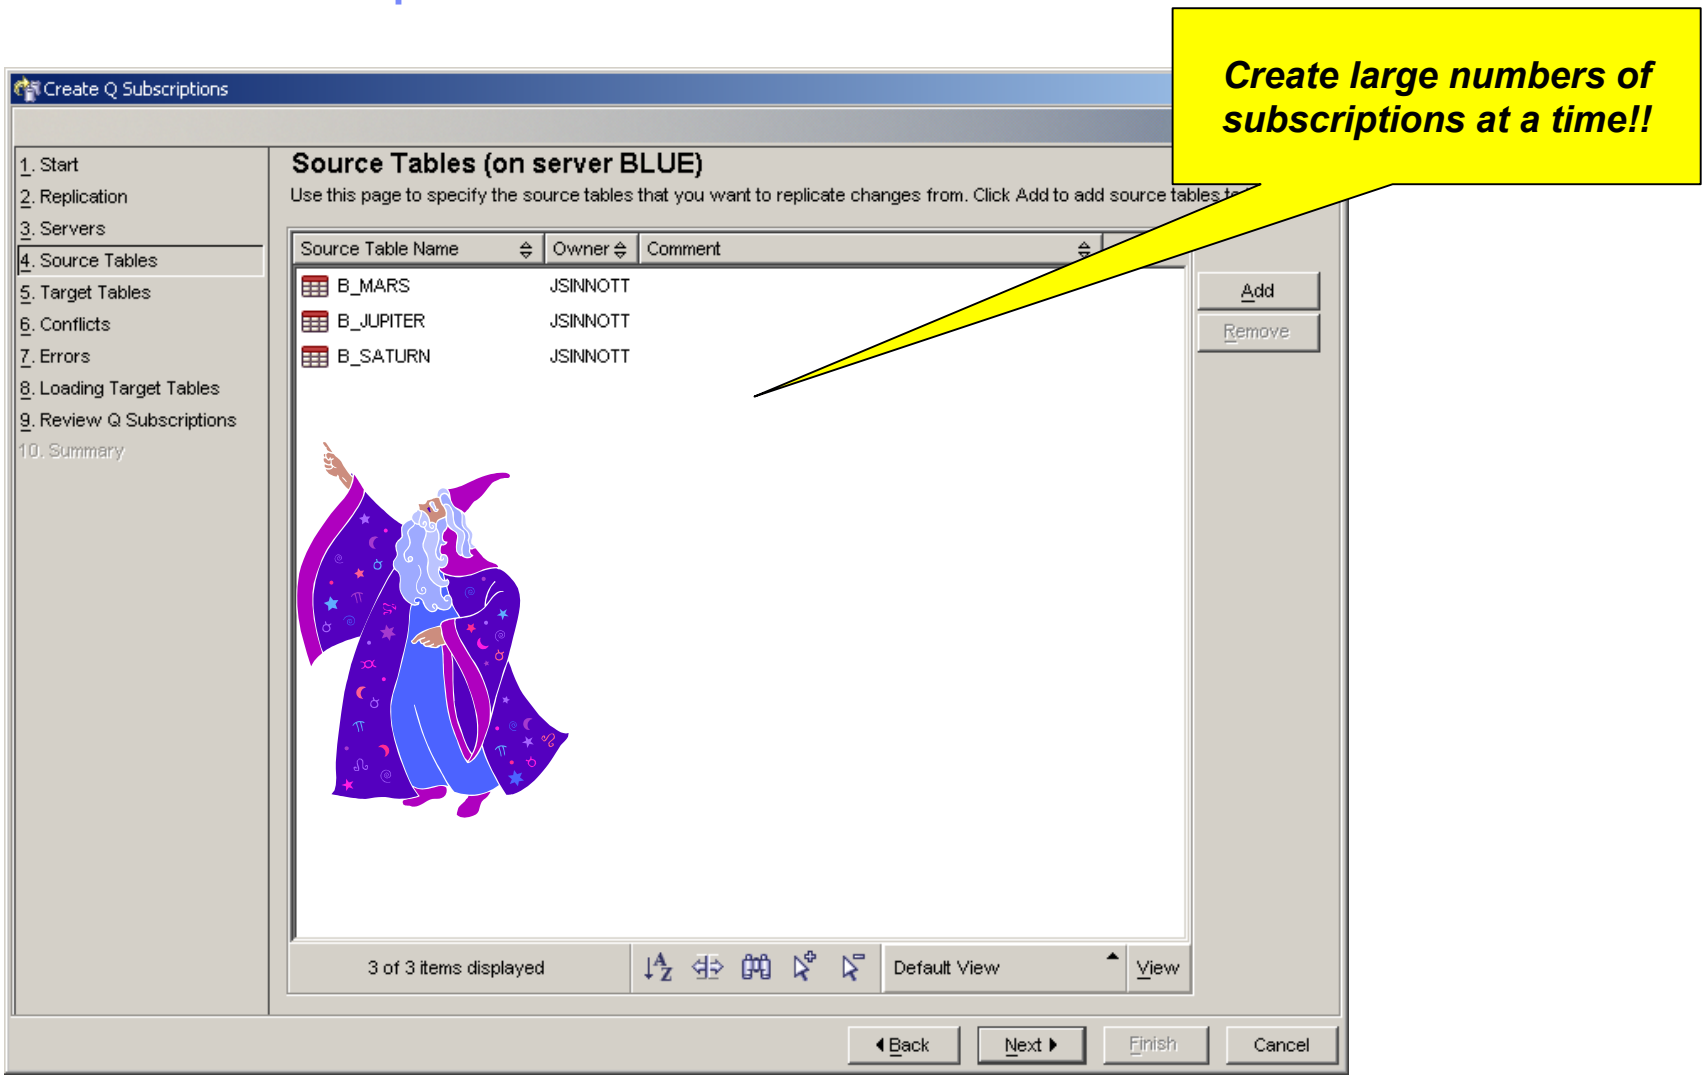

Implementierung der ‚Q Replication' im VM & VSE

- Nur Q Capture, kein Q Apply
- VM hat keinen MQ Server
- Replikation ins VSE nur über ‚Federated Database Support'
- Keine ‚Event Publication'
- Keine SQL-Replikation mehr!

Implementierung der ‚Q Replication' im VM

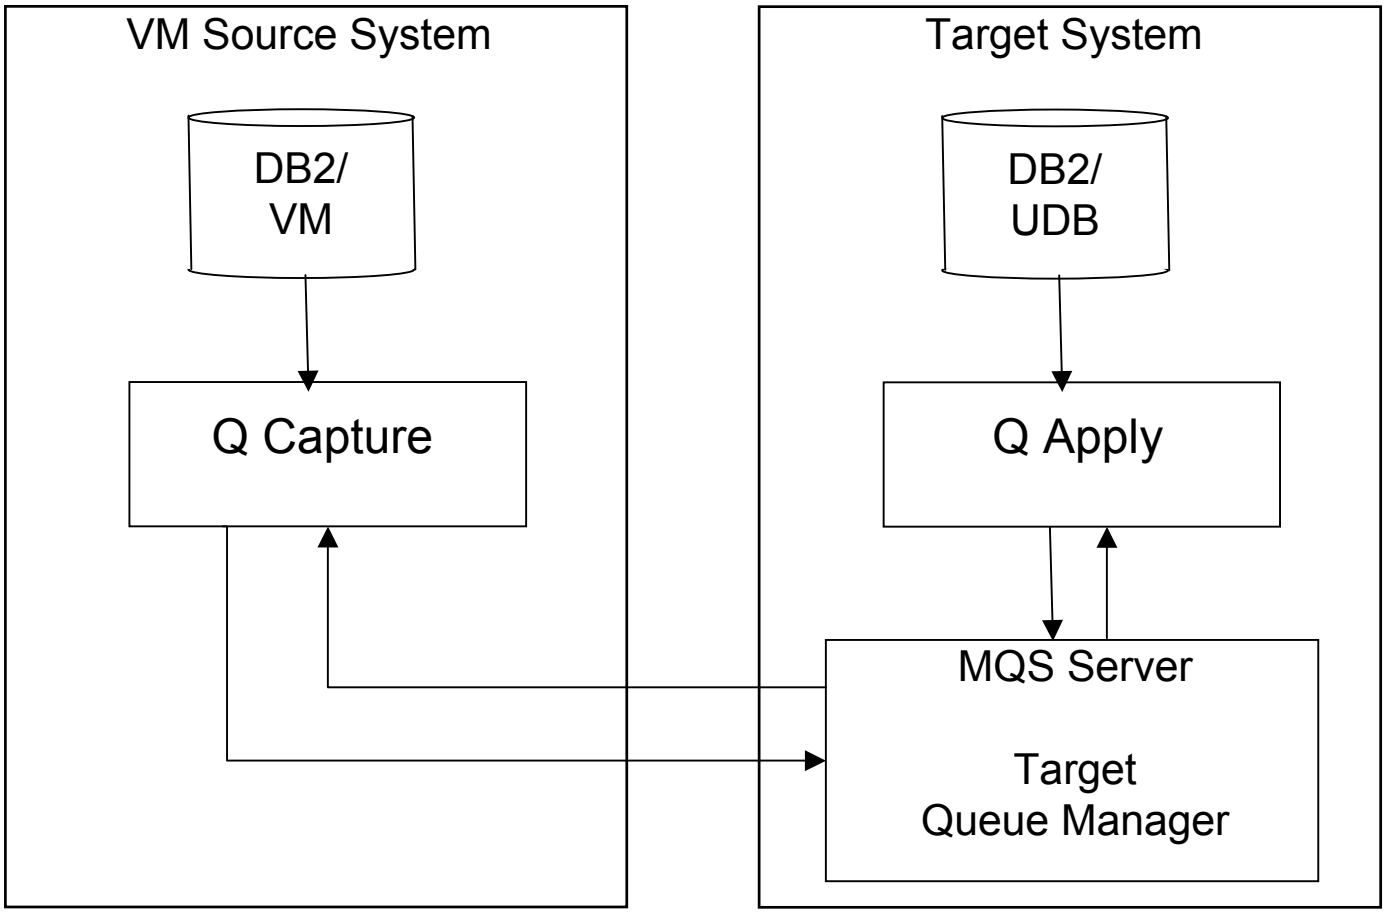

Implementierung der ‚Q Replication' im VSE

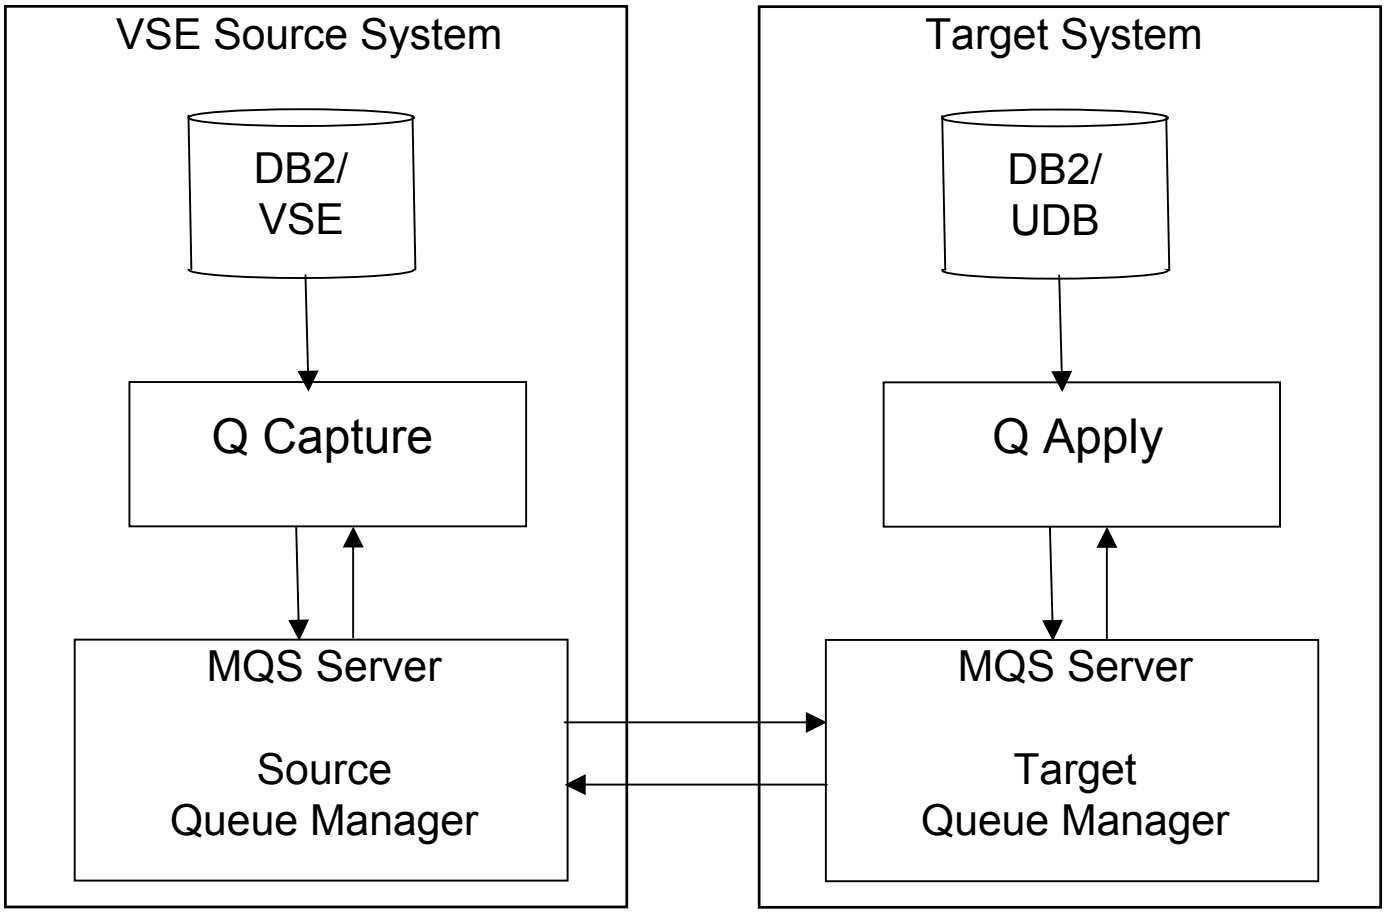

Quo Vadis - DB2 Server for VSE & VM

- Zukunft: DB2 VM & VSE ‚Client only‘?
  - DRDA Performance!
  - UDB-Funktionalität für VSE Anwendungen?
    - Unterstützung aller rein Server-basierenden Funktionen
      - z.B. CASE - Funktion
    - Kein Support für Funktionen die erweiterte Funktionalität des ‚Clients‘ erfordern
      - z.B. CURSOR WITH HOLD

DB2 Server for VSE & VM – weitere Informationen

- Dokumentation: Q Capture Supplement
  - Verfügbar über die DB2 VM/VSE Homepage

    http://www.ibm.com/software/data/db2/vse-vm/

- End of Service für DB2 Server for VM & VSE V **7.1**

  2. Mai 2005

- Kontakt: roeber@de.ibm.com

Haben Sie Fragen?