IBM zSeries Systems and Technology Group

MQ Series der moderne Datenaustausch

GSE Dresden, April 2006

Wilhelm Mild
z/VSE Solution Architect
IBM Boeblingen Laboratory, Germany

IBM zSeries and System z

Trademarks

References in this publication to IBM products or services do not imply that IBM intends to make them available in every country in which IBM operates. Consult your local IBM business contact for information on the products, features, and services available in your area.

AIX*, APPN*, CICS*, CICS/VSE*, CICS, DB2*, DB2 Connect, DB2 Universal Database, DFSORT, DRDA*, e-business logo*, Enterprise Storage Server, FlashCopy, HiperSockets, IBM*, IBM logo*, IBM eServer, iSeries, Language Environment*, MQSeries*, Multiprise*, pSeries, S/390*, S/390, System z, Parallel Enterprise Server, TotalStorage, VSE/ESA, z/VSE, VTAM*, WebSphere*, xSeries, z/OS, z/VM, zSeries and Distributed Relational Database Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds.
Java and all Java-related trademarks and logos are trademarks or registered trademark of Sun Microsystems, Inc.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

Microsoft, Windows, Windows NT, Visual Basic and the Windows flat logo are Trademarks of Microsoft Corporation.

Other trademarks and registered trademarks are theproperties of their respective companies.

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply. This equipment is subject to all applicable FCC rules and will comply with them upon delivery.
Information concerning non-IBM products was obtained from the suppliers of those products. Questions concerning those products should be directed to those suppliers.

All statements regarding IBM’s future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Typical VSE Customer Environment

- Various different servers (zSeries, pSeries, iSeries, xSeries, and competitive)
- VSAM data on VSE (few DB2 environments)
- Relational databases on distributed platforms (DB2, Oracle)
- Data interchange via FTP

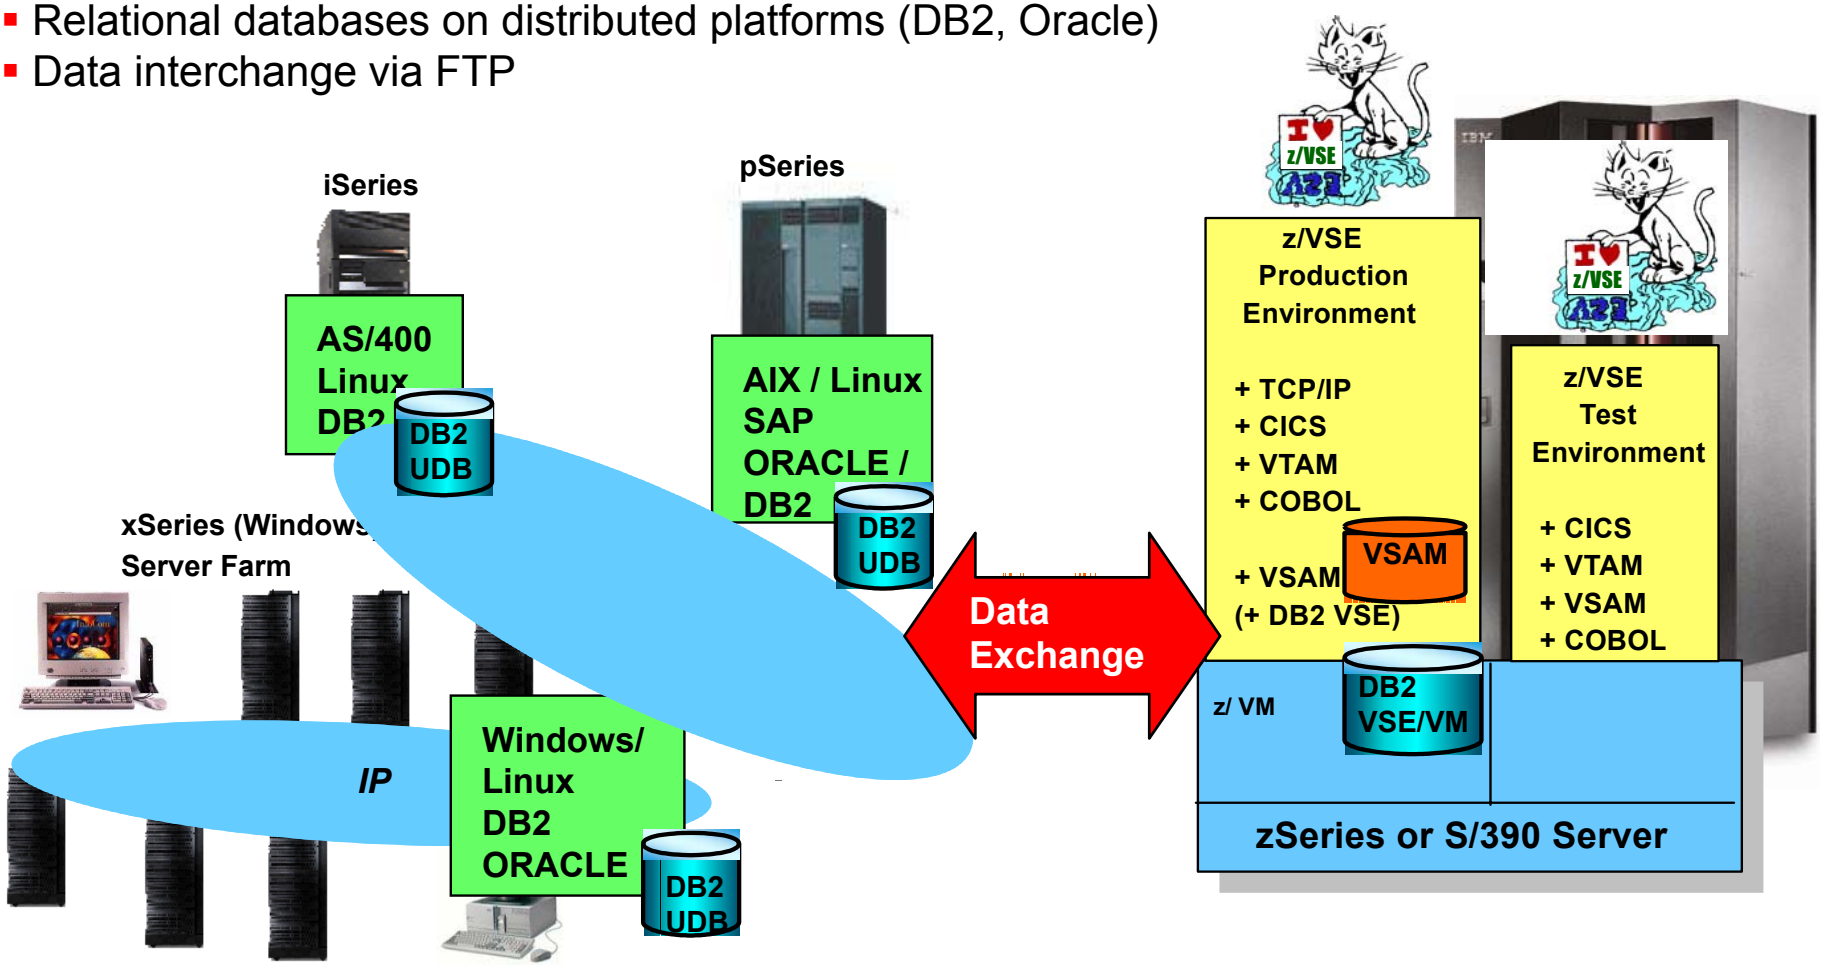

Inter Program Communication

Direct Communication

- Program Call / Response
- RPC Remote Procedure Call

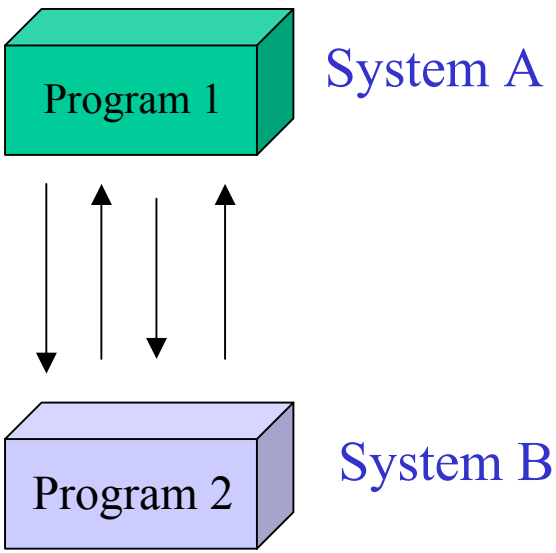

Asynchronous Communication

- Using Messaging
  (no direct program invocation)

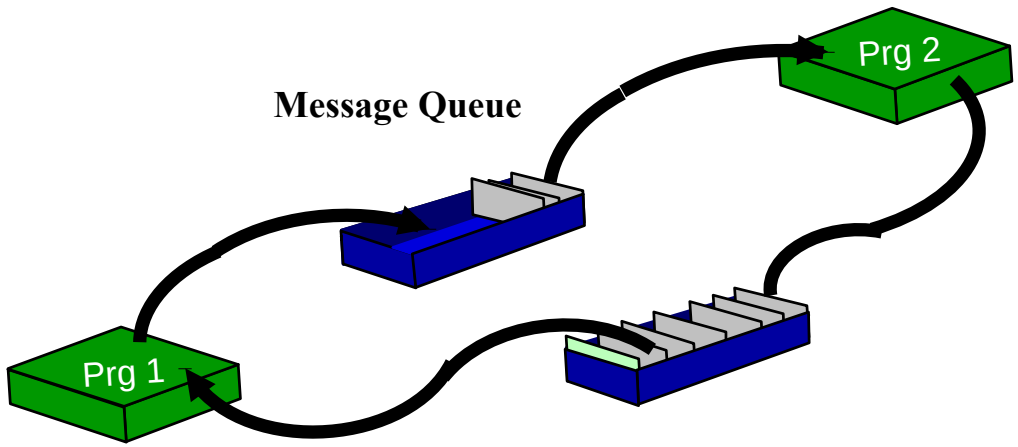

Why are interfaces so expensive to build and maintain?

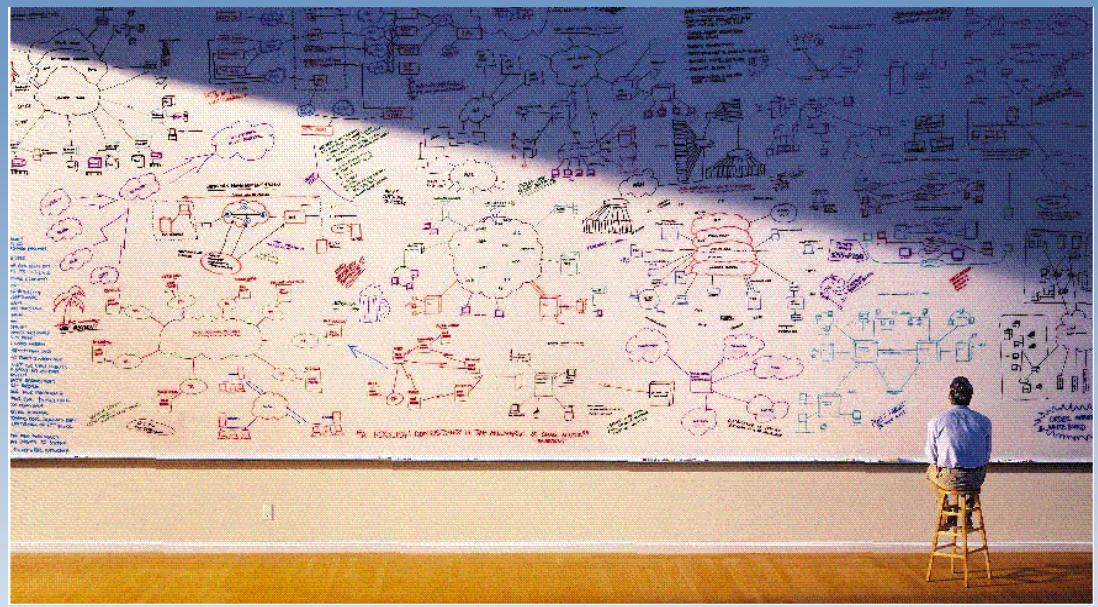

- **Application interface logic is intertwined with business logic.**
- **The more tightly integrated the interface the more difficult the application is to change.**
- **The more interfaces that exist within a set of programs, the more complex the application becomes -- interface logic may, in many cases, exceed business logic.**
- **In such circumstances, re-use becomes difficult and impractical.**

The Message Queuing Interface

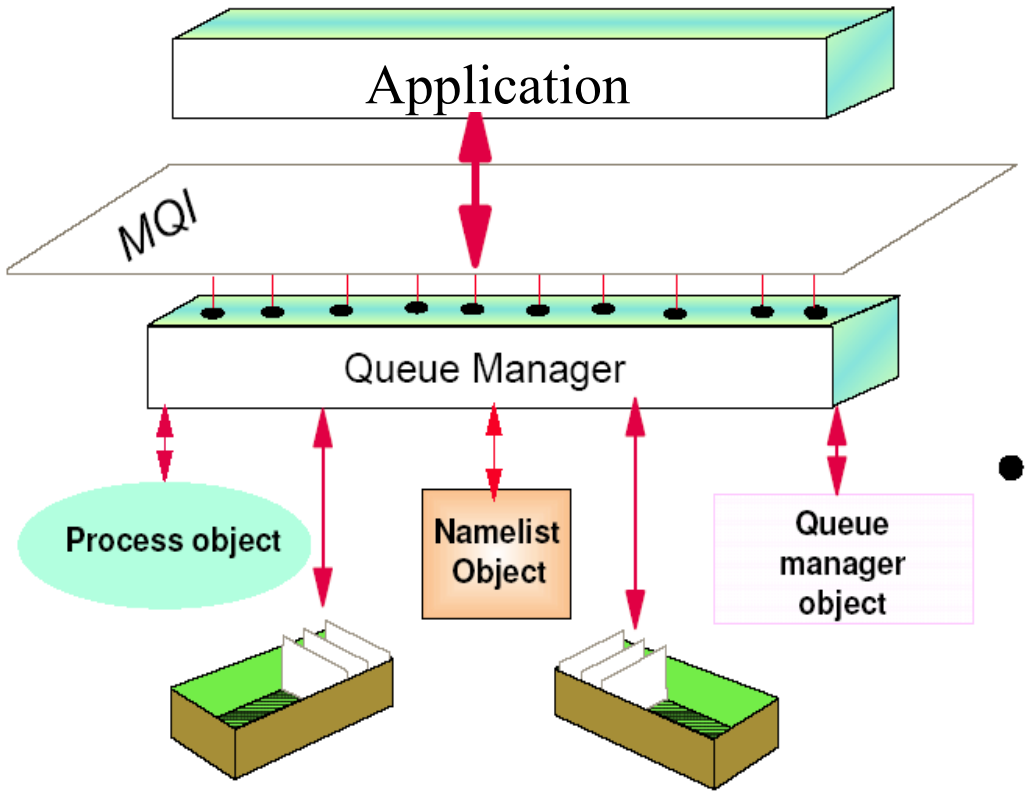

- **Major calls**
  - MQCONN
  - MQCONNX
  - MQDISC
  - MQOPEN
  - MQCLOSE
  - MQPUT
  - MQPUT1
  - MQGET
- **Minor calls**
  - MQBEGIN
  - MQCMIT
  - MQBACK
  - MQINQ
  - MQSET

Messaging Overview
MQ Server – Server scenario

Event Notification (1 way communication),
Request / Response (2 way communication)

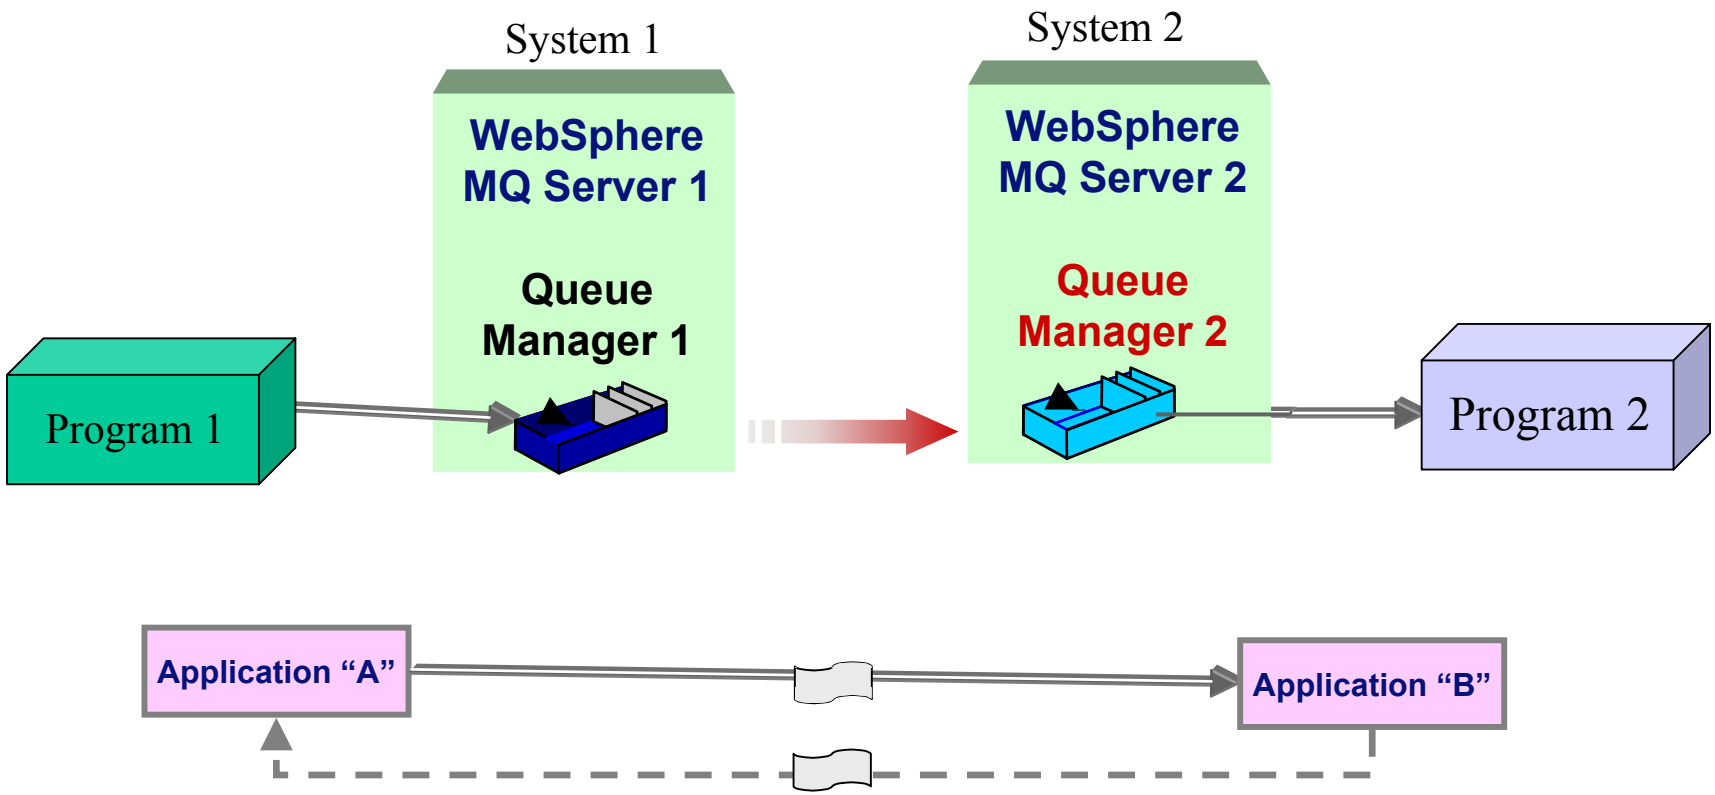

Messaging Overview
MQ Client – Server scenario

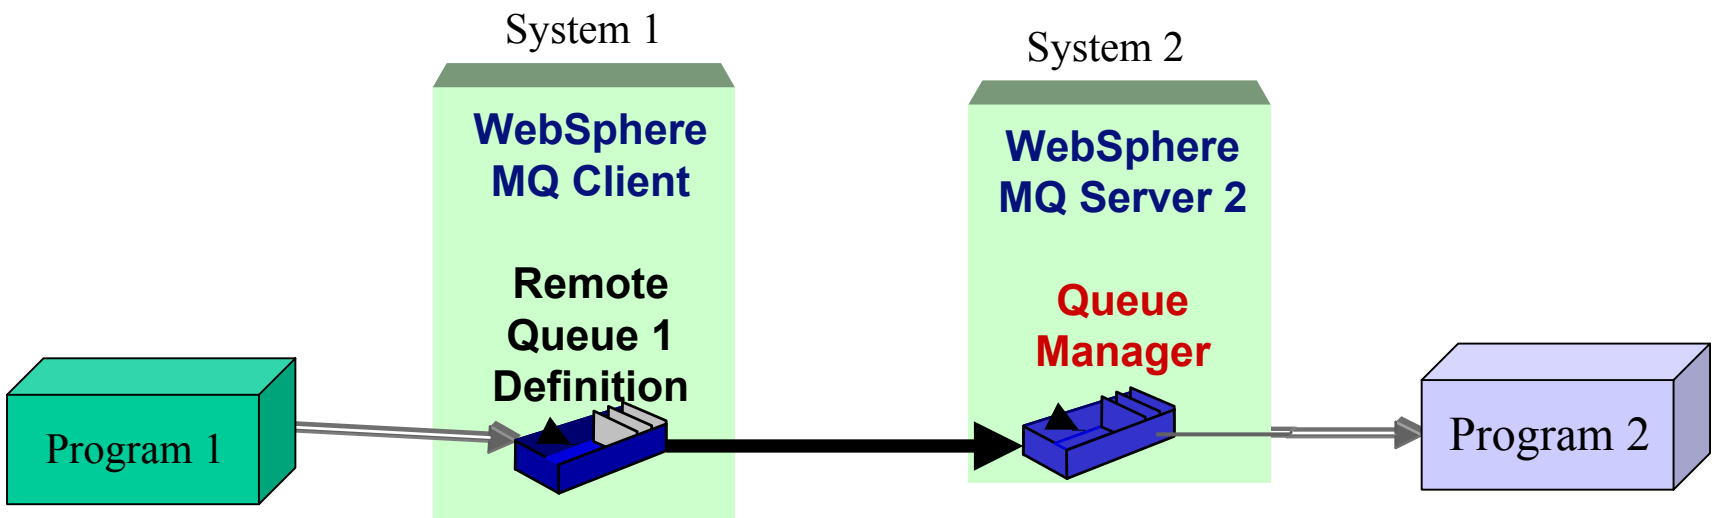

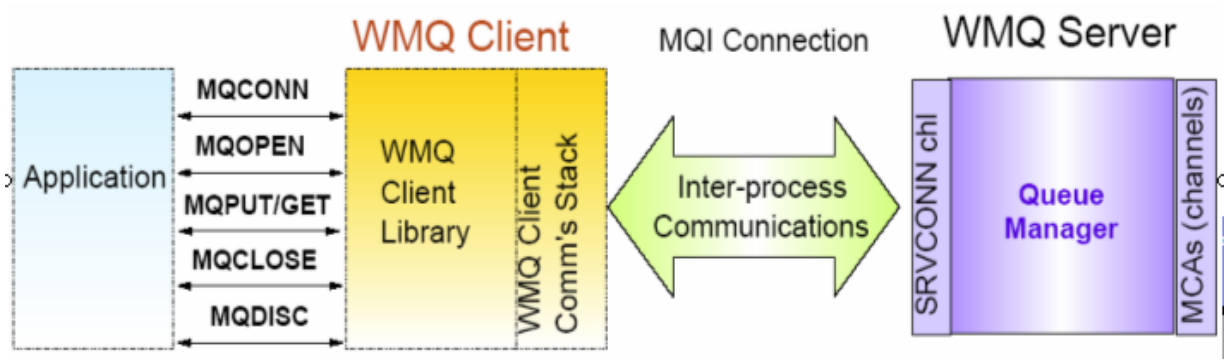

MQ Series servers and MQSeries Clients

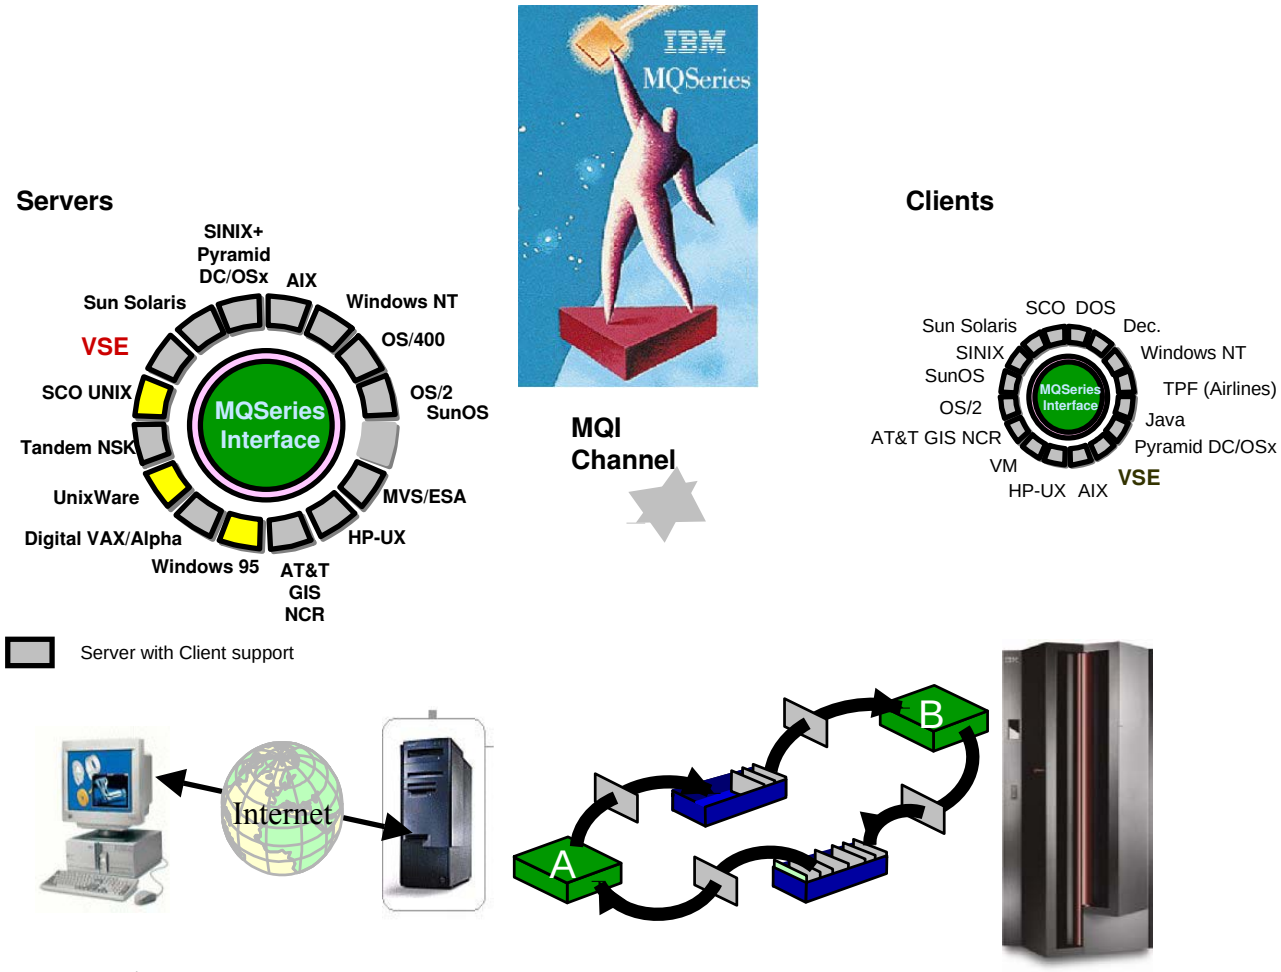

- Various platforms supported
- integration into Web Application servers (WebSphere)

Messaging Transformation Overview

Event Notification (1 way communication), Request / Response (2 way communication)

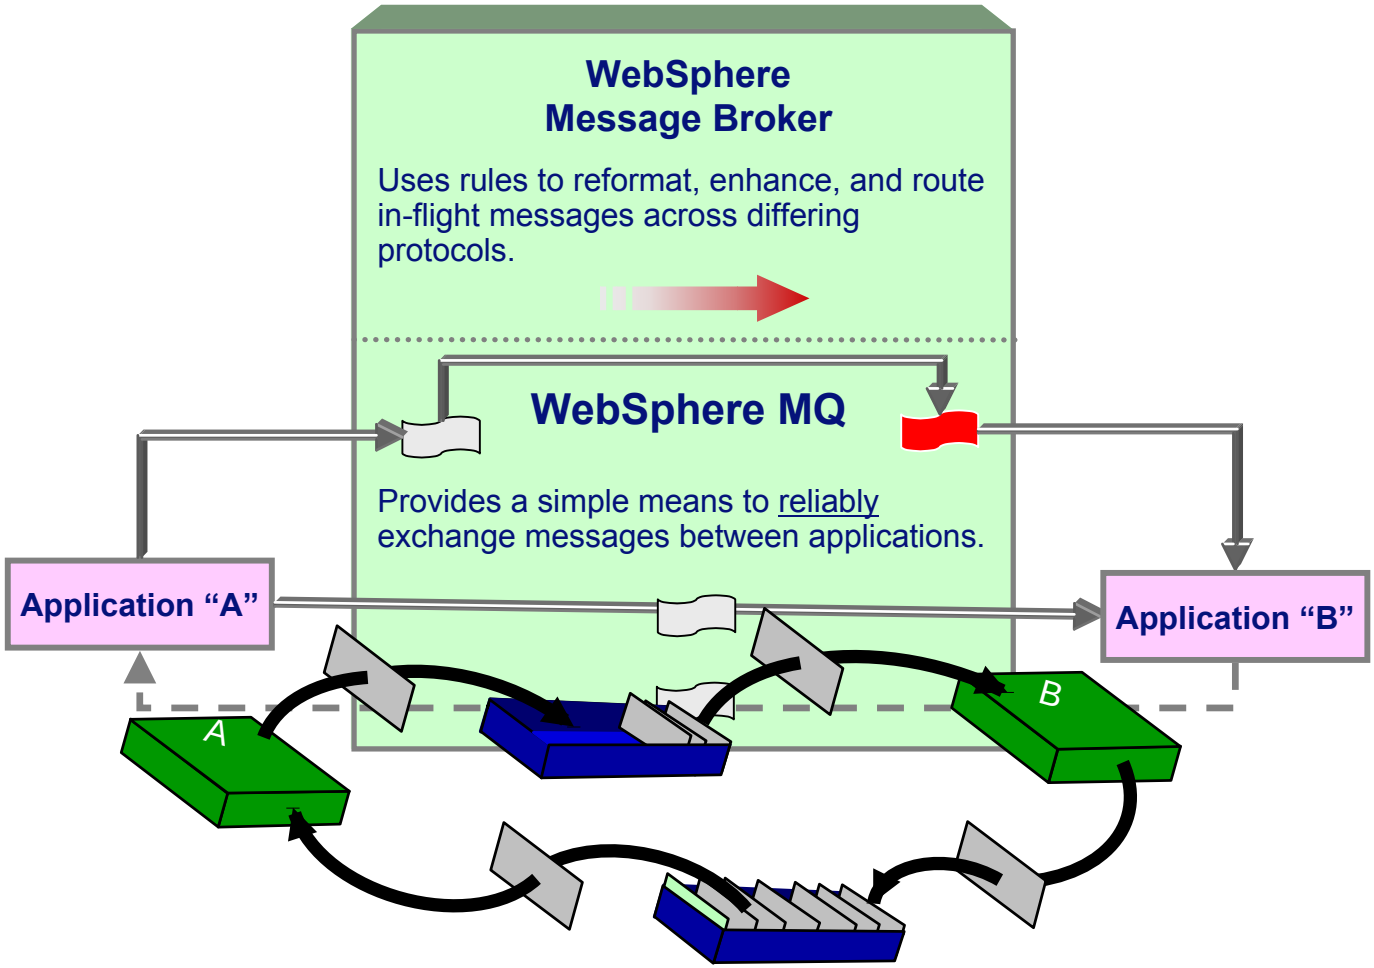

Message Brokering

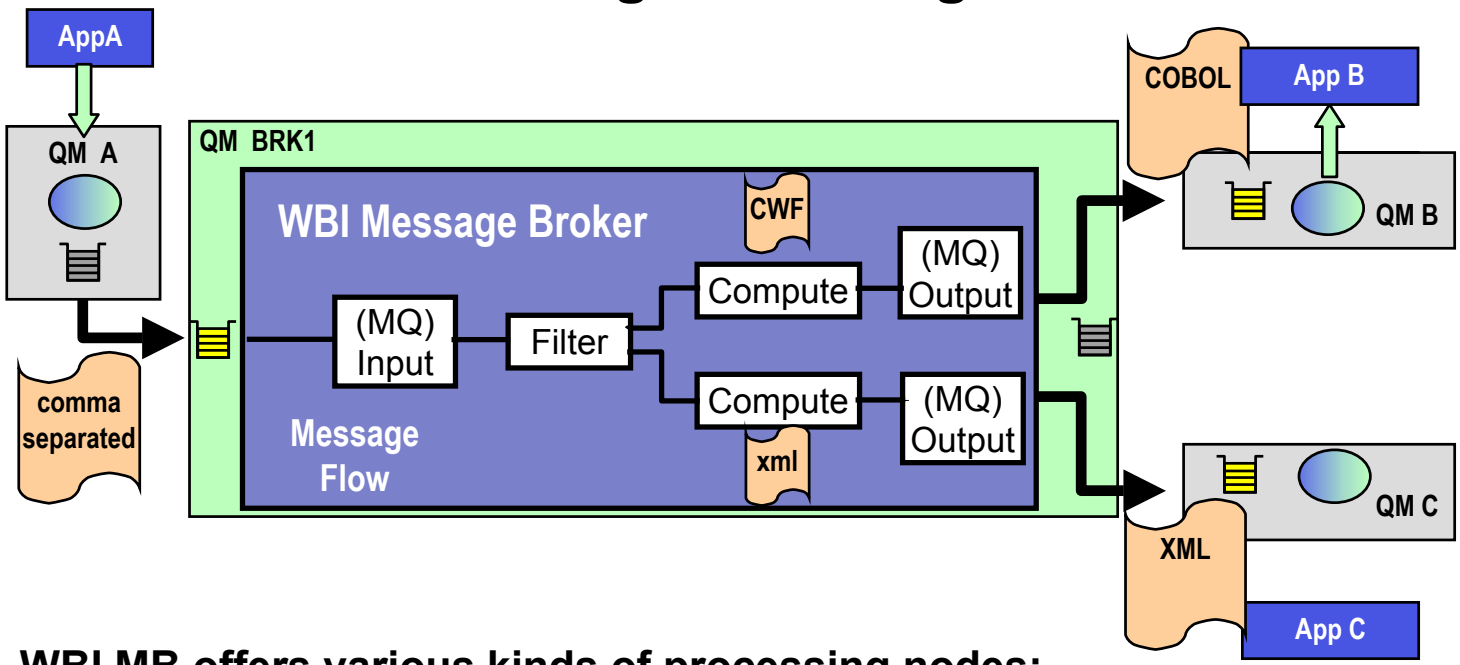

WBI MB offers various kinds of processing nodes:

- ✓ Receiving and routing messages
- ✓Transforming a message to an alternative representation
- ✓Selecting a message for further processing based upon the message's content
- ✓Interacting with an external repository to augment a message or store the whole or part of a message
- ✓Responding to events and errors

What is WebSphere Message Broker?

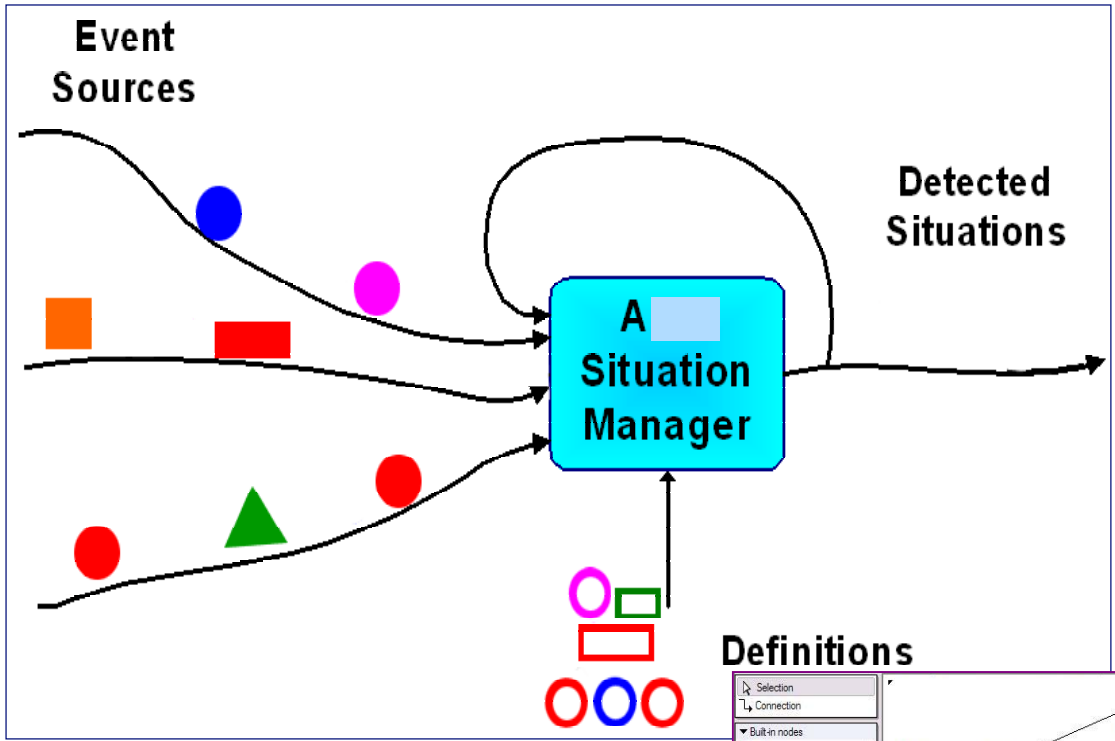

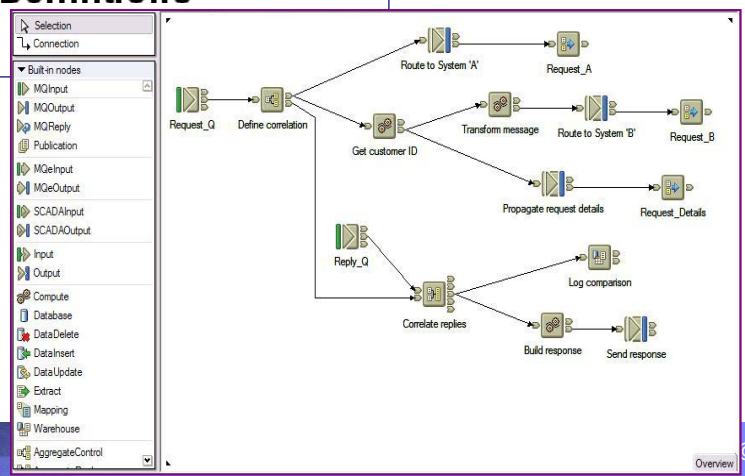

1. **A framework for processing MQ messages**
2. **A robust hosting environment for:**
   - ✓ Transforming data
   - ✓ Enriching data
   - ✓ Interacting with databases
   - ✓ Routing messages based on content
   - ✓ Detecting complex combinations of messages
   - ✓ Interacting existing applications with Web Services

What is an Enterprise Service Bus?

An Enterprise Service Bus (ESB) is a flexible connectivity infrastructure for integrating applications and services.

An ESB powers your SOA by reducing the number, size, and complexity of interfaces.

An ESB performs the following between requestor and service

- **ROUTING** messages between services
- **CONVERTING** transport protocols between requestor and service
- **TRANSFORMING** message formats between requestor and service
- **HANDLING** business events from disparate sources

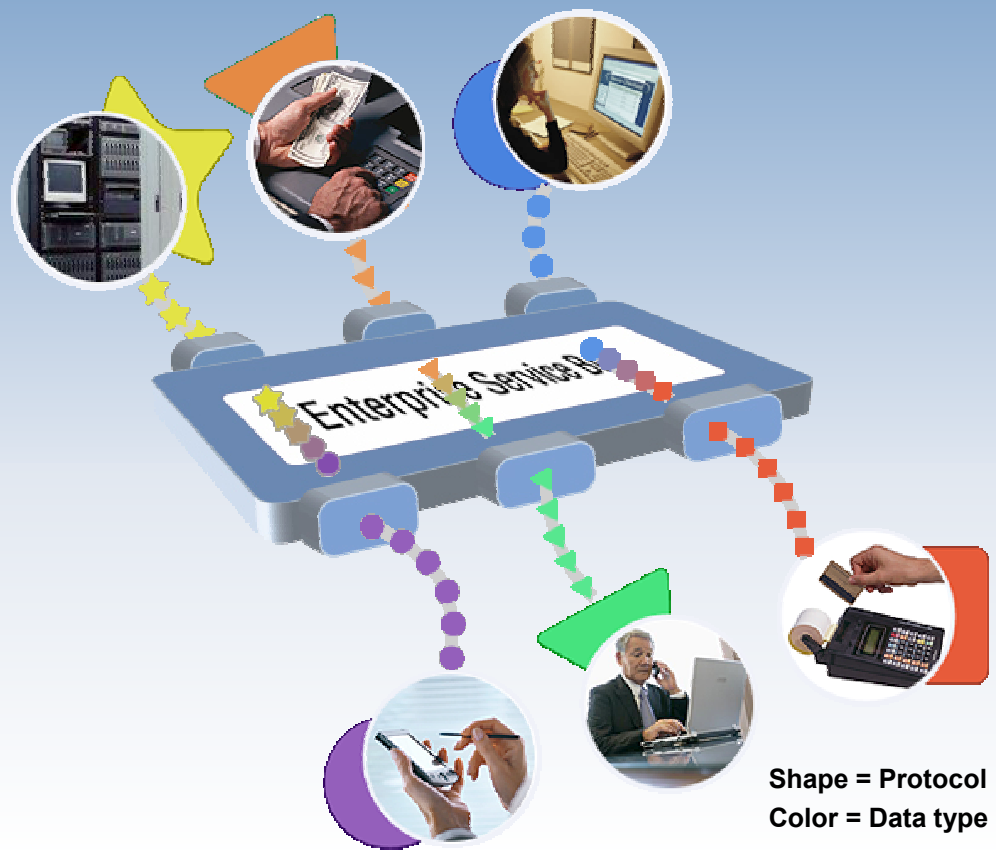

SOA: the next step on the connectivity evolution

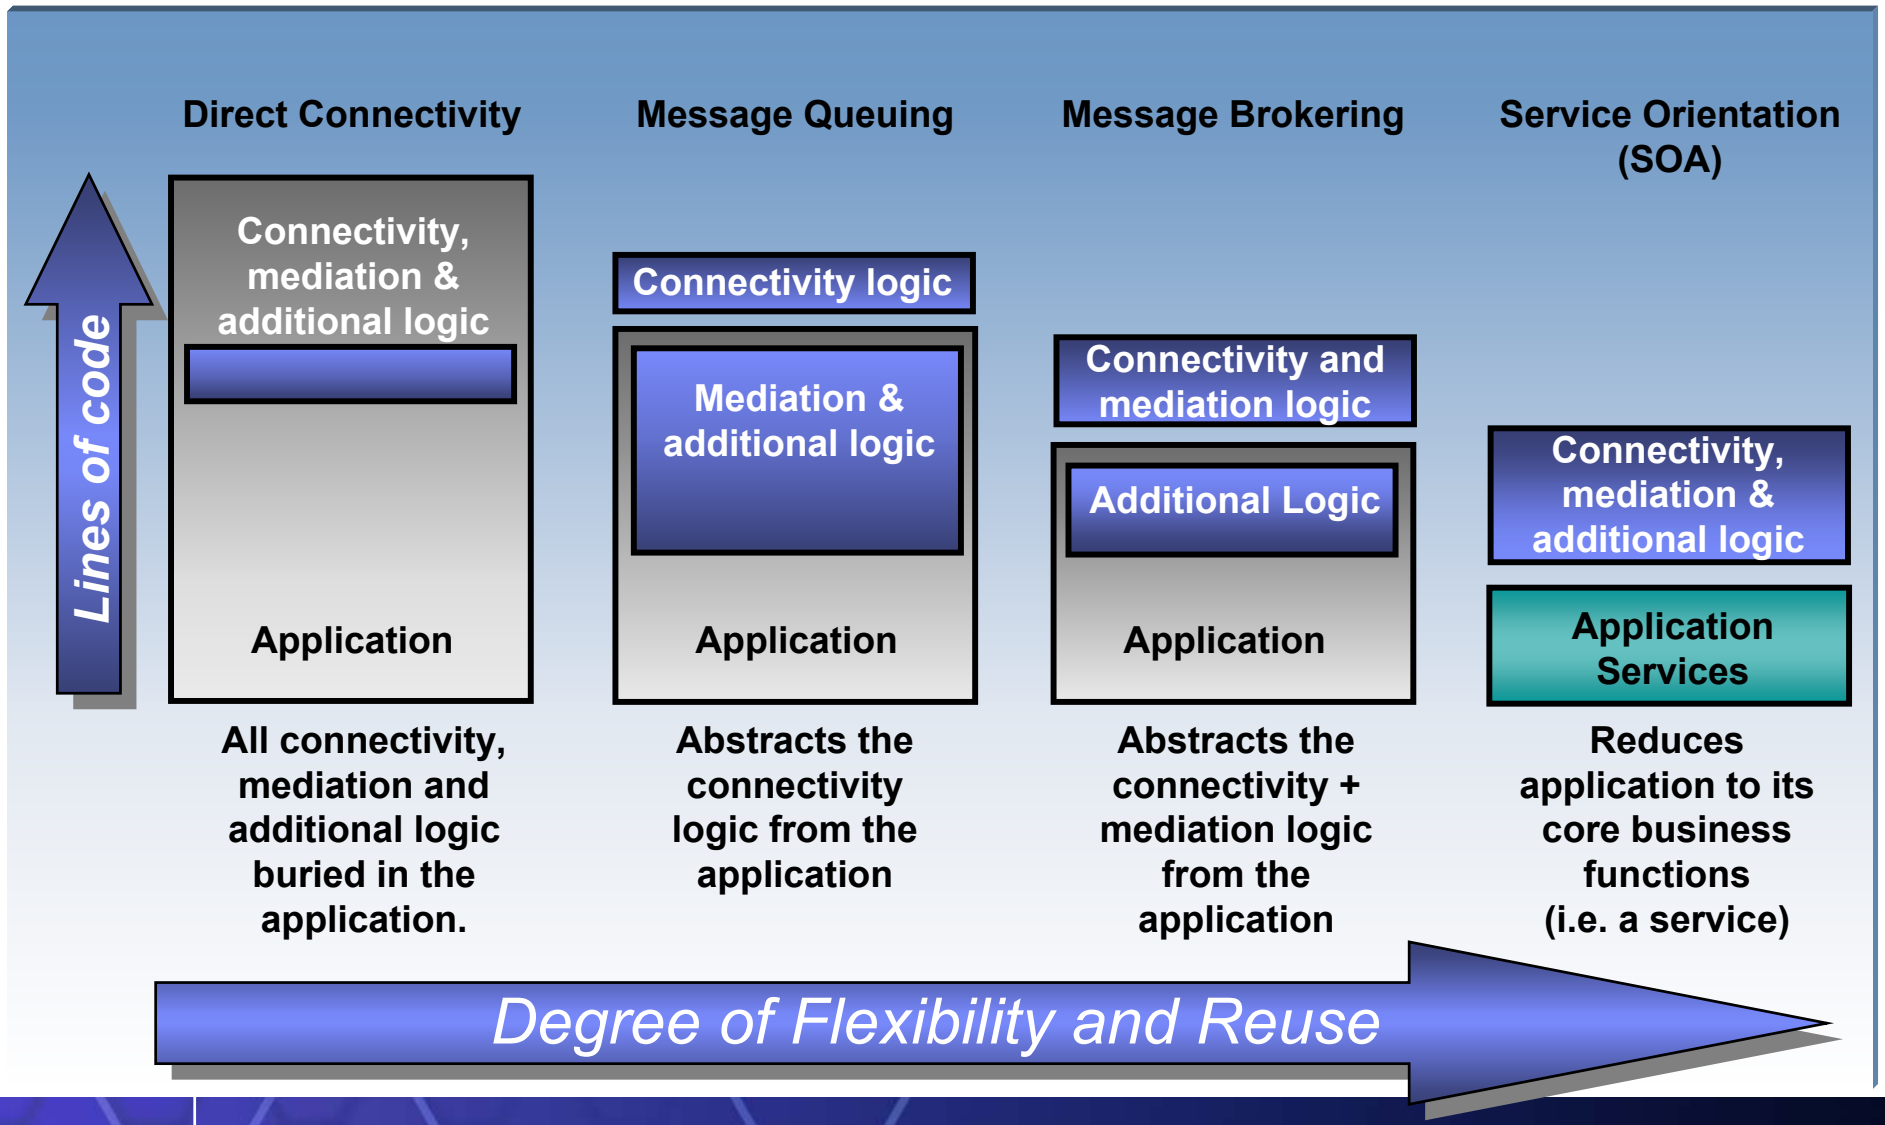

WebSphere Message Broker 6.0

Delivering an advanced ESB to power your SOA

Provides universal connectivity

- Provides Web Services connectivity and non standard interface connectivity
- Unmatched ability in integrating many systems, platforms, devices, and APIs
- Facilitates service oriented integration

Provides universal data transformation

- Advanced message transformation, enrichment, and routing
- Support for industry standard data formats (AL3, HL7, SWIFT, HIPAA, EDI, etc.)

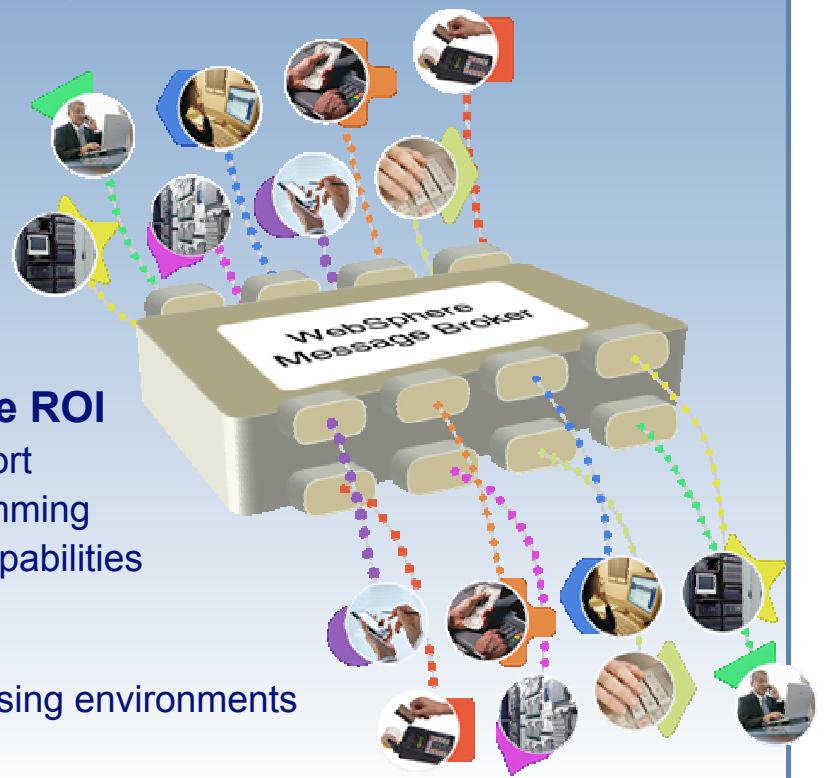

New & improved pre-built capabilities to improve ROI

- Leverage existing skills with rich Java and XML support
- Implement complex event processing with no programming
- Offers simple and easy to use tools with advanced capabilities

Leverage the performance

- Offers performance of traditional transactional processing environments

Integrate your existing environment with the world of web services

Integration of VSE Programs with MQ Series

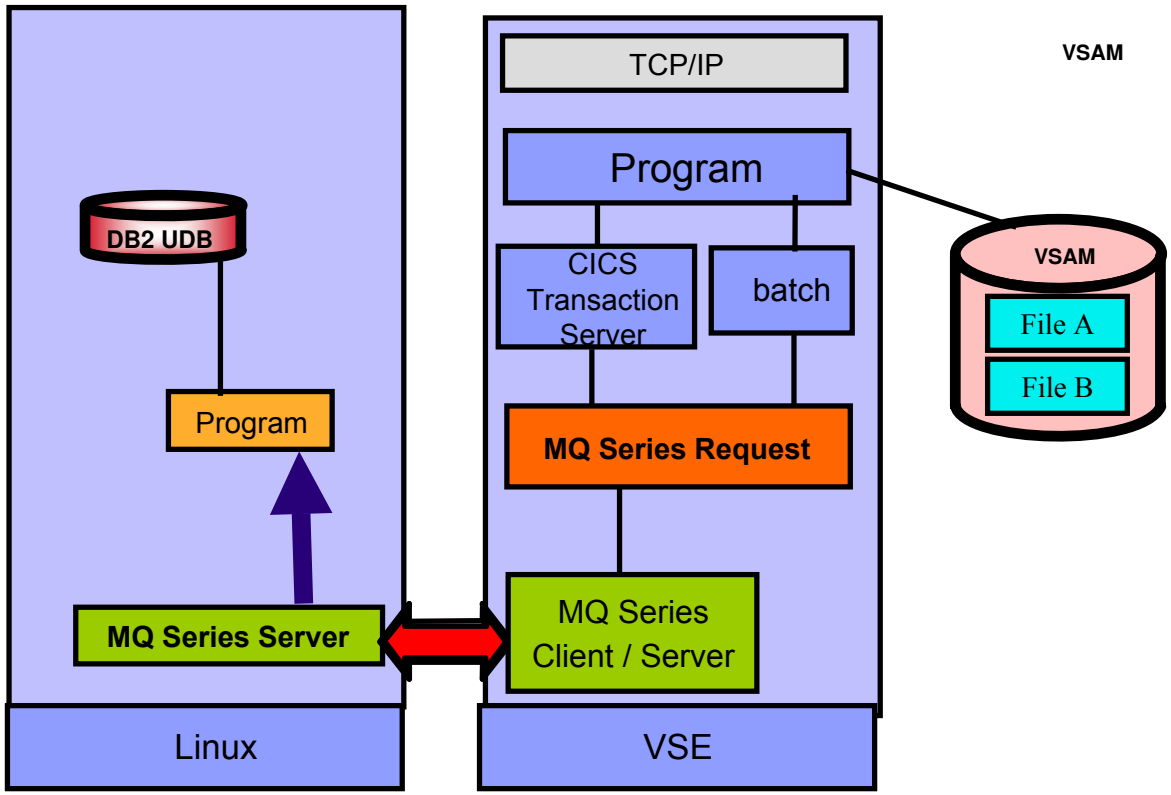

- Data distribution via MQ Series technology
- VSE programs have to write MQ messages – requires changes to existing applications
- NEW: WebSphere MQ Series Client for VSE – free of charge – enablement for MQ environments and modern solutions

VSAM Data collection / transformation / journaling on VSE

Vendor Exit

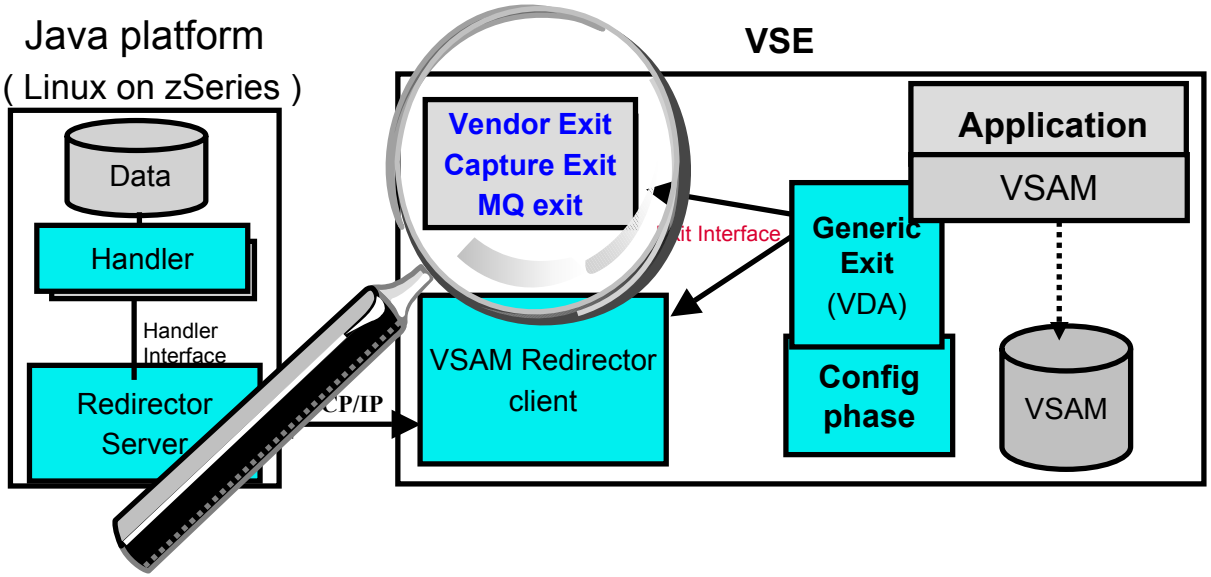

- Vendor Exit
  - user (vendor) written phase for data collection/transformation
  - has to comply with the documented Exit Interface

Note: No chaining of Vendor Exit with VSAM Redirector client supported

Architectural View

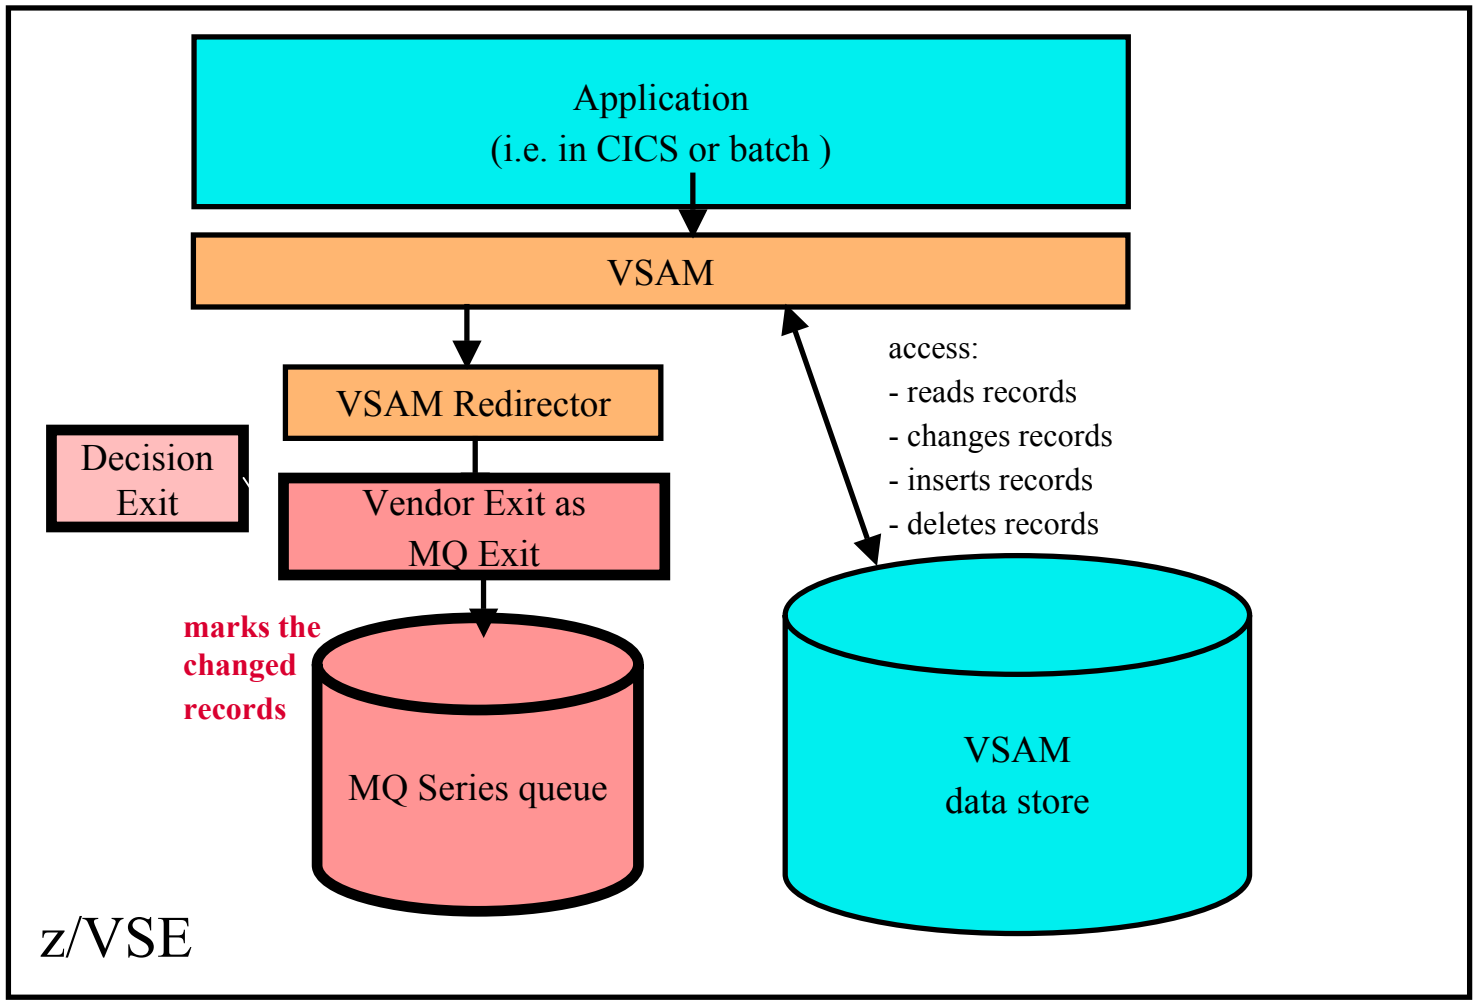

Integration of VSE Programs (w/o changes) with MQ Series

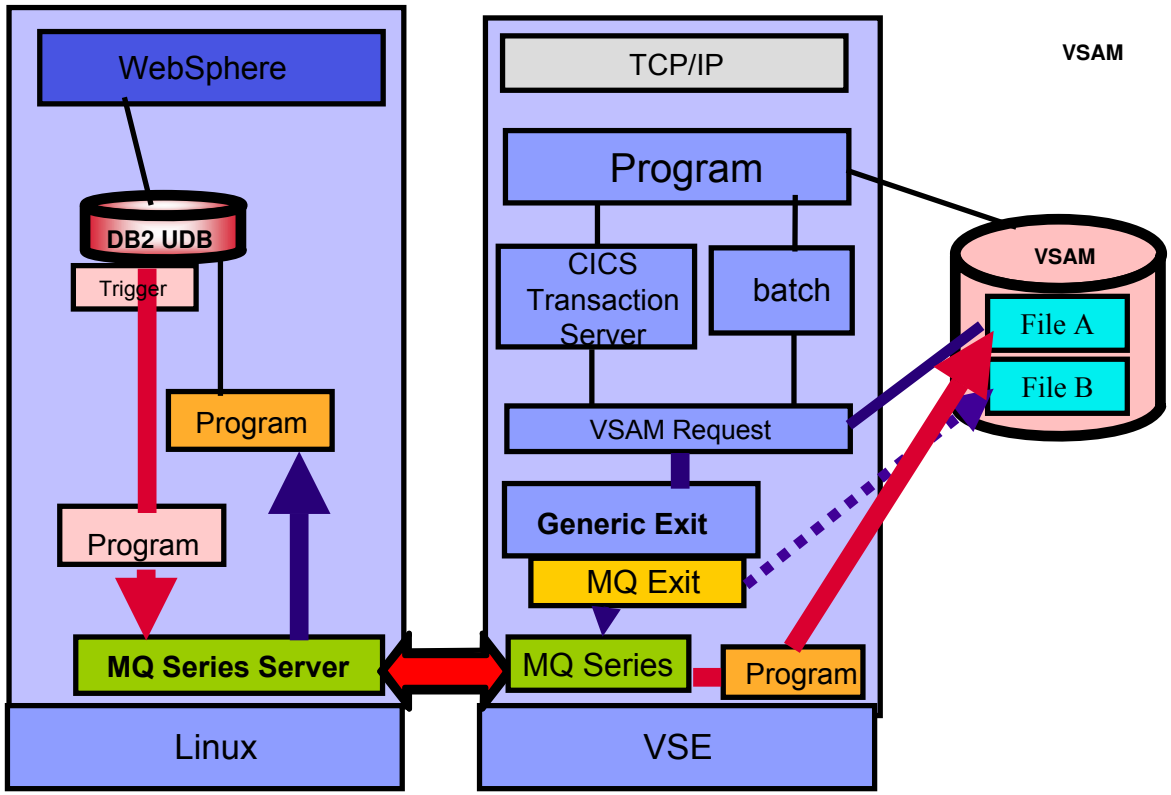

- Data distribution via MQ Series technology
- NO changes to VSE programs – using MQ Exit and VSE VSAM Redirector
- NEW: WebSphere MQ Series Client for VSE – free of charge – enablement for MQ environments and modern solutions

MQ Series - asynchronous transactions

- functional characteristics
  - guaranteed, secured asynchronous data access for remote systems
  - same API for all supported MQ Series platforms
  - transaction security, therefore appropriate for e-business processes
  - integration with WebSphere Application Server
  - works well for Business-to-Business (B2B) environments

- software requirements
  - For VSE/ESA :
    - VSE/ESA 2.6/2.7
    - MQ Series Server
    - Program that interfaces with MQ on VSE or VSAM Redirector
  - On the remote system:
    - MQ Series Client / Server
    - Program that interface with MQ Series

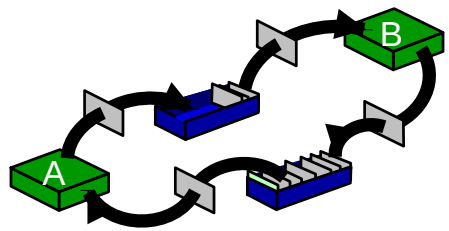

WebSphere MQ Series Solutions with z/VSE

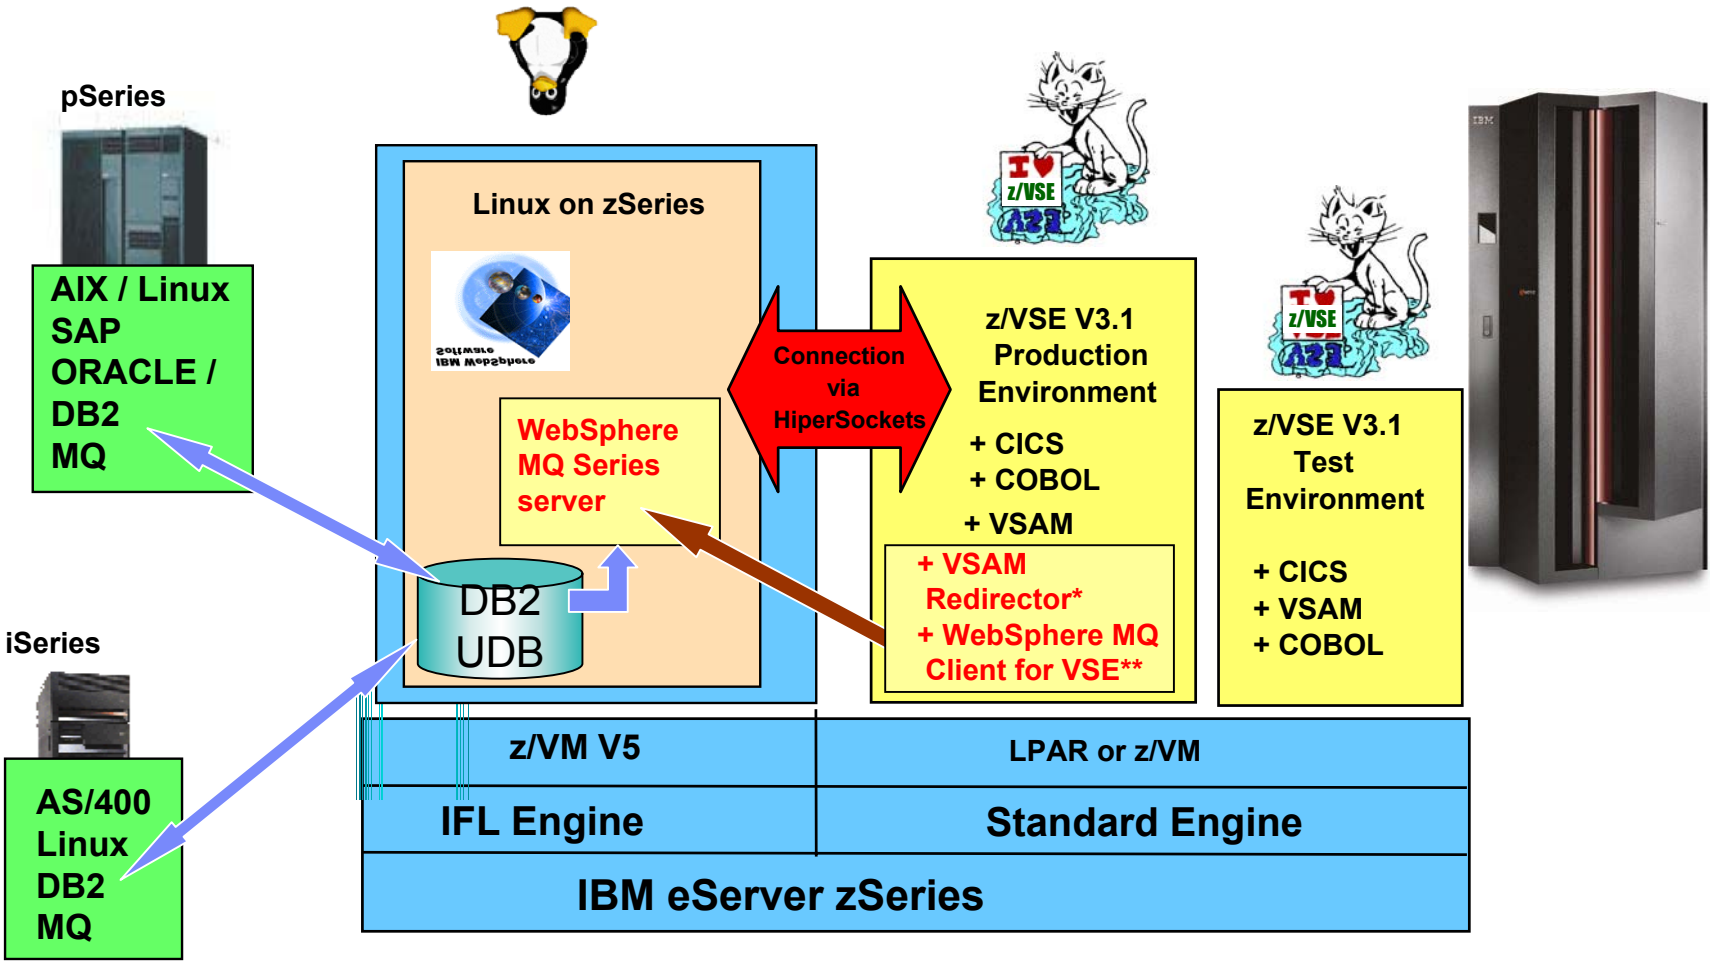

(*) **VSAM Redirector + Redirector MQ Exit** allows MQ Solutions without changes to VSAM programs
(**) **WebSphere MQ Client for VSE** is brand-new and free of charge

DB2/VSE Q-Replication als Anwendungs-Beispiel für Message Queuing Services

IBM Software Group

Torsten Roeber

GSE Frühjahrstagung, April 2006

Current SQL Replication Architecture

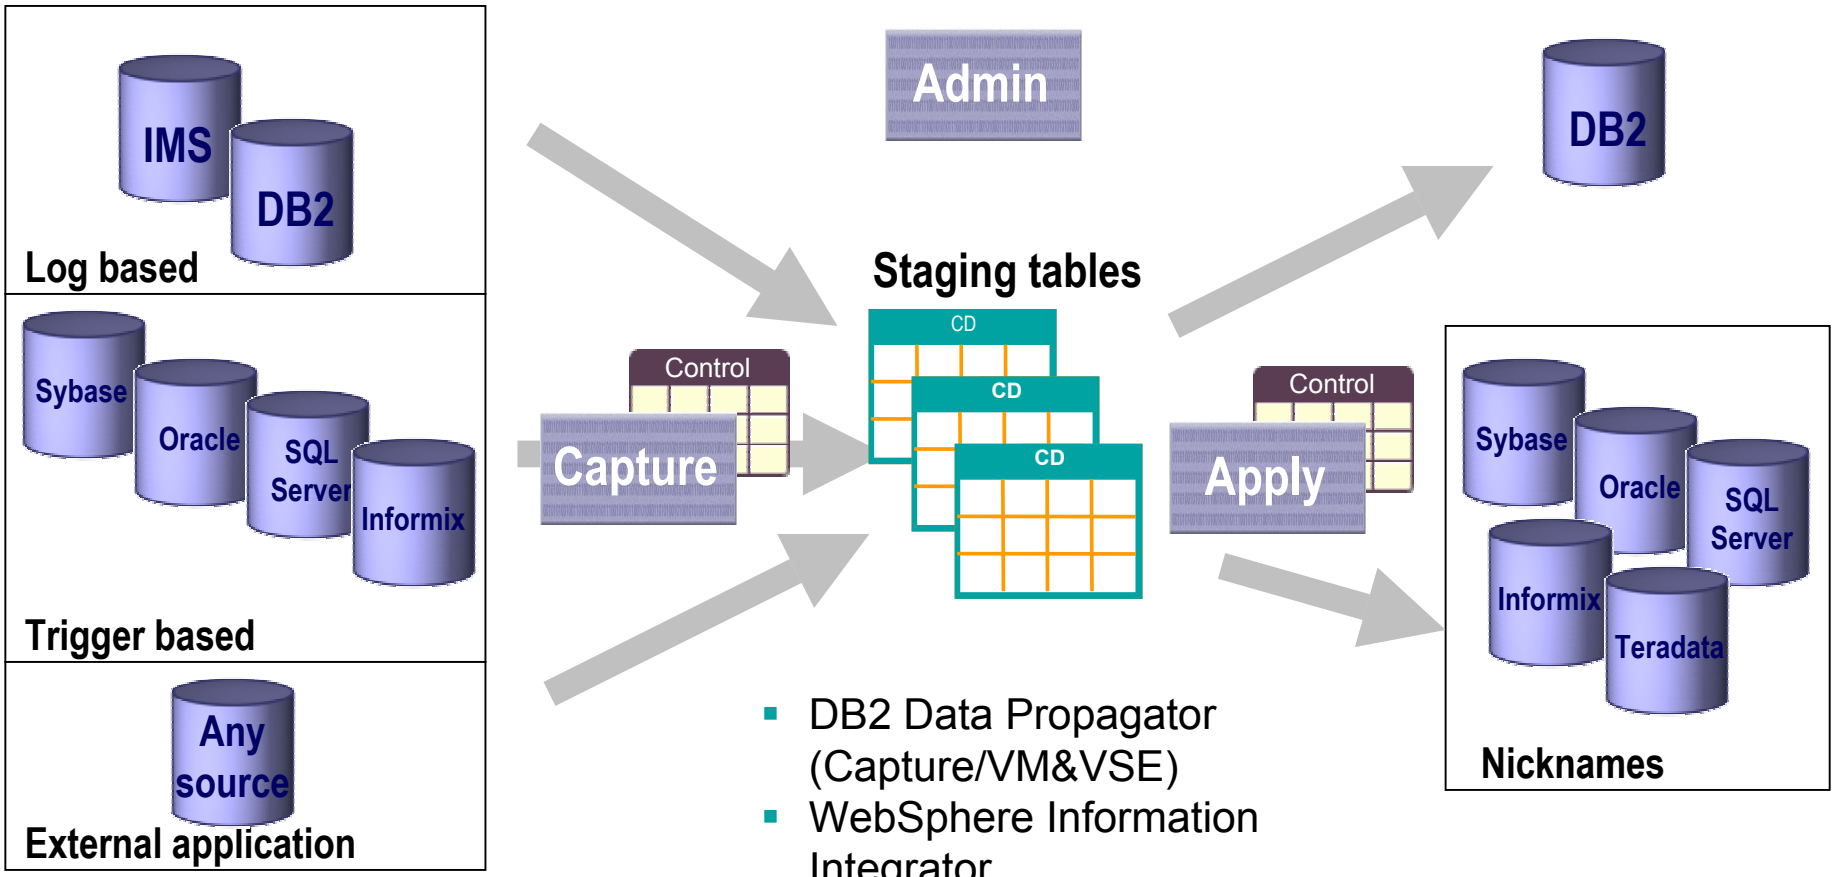

- DB2 Data Propagator (Capture/VM&VSE)
- WebSphere Information Integrator
- IMS Data Propagator

Why Create Another Replication Architecture?

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

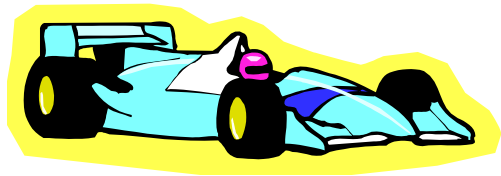

- **Capability:** Significantly improve multi-directional replication support

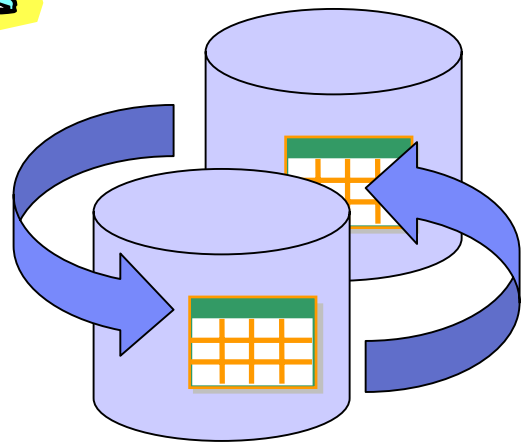

- **New function:** Event publishing, table difference utility

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

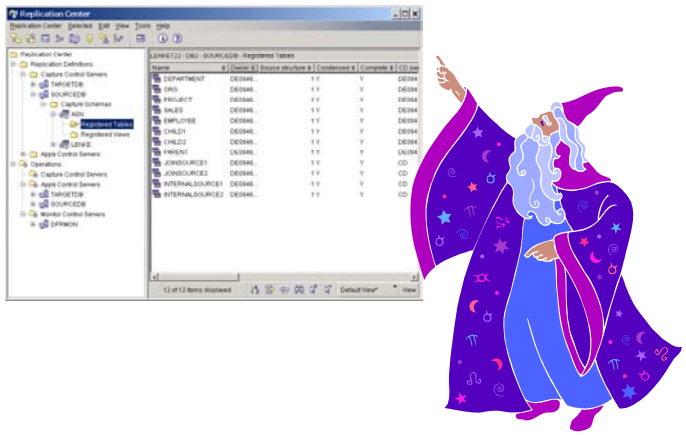

Q Replication – Q Subscription Process

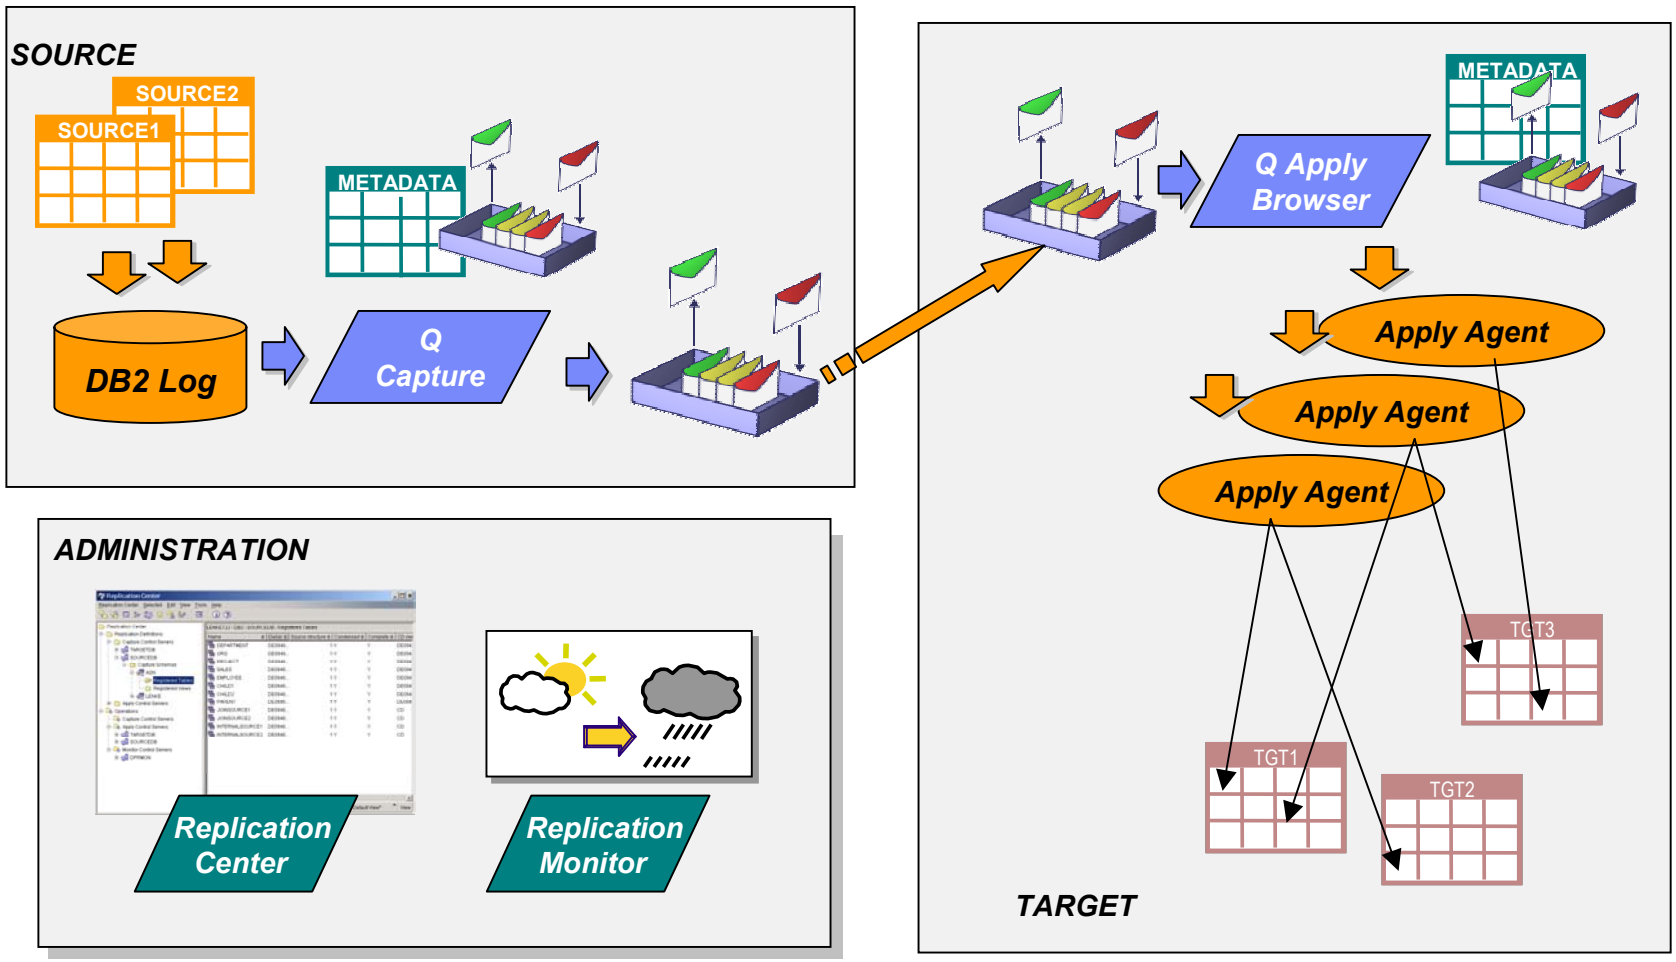

Q Replication – Event Publication Process

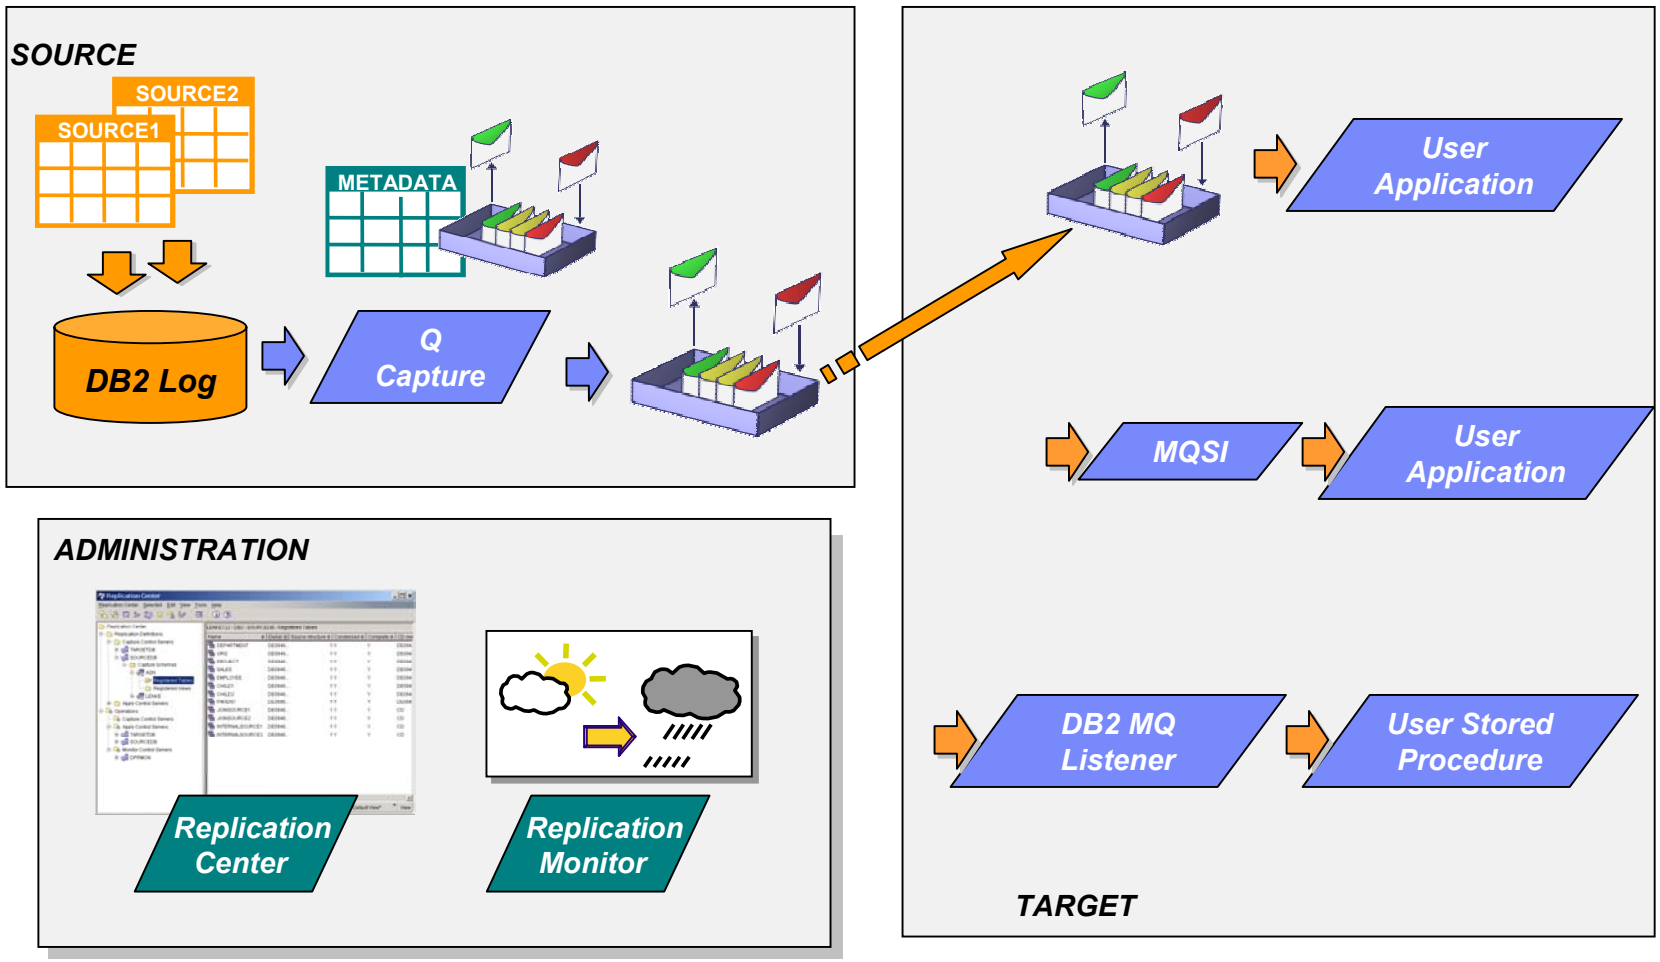

Subscription Types

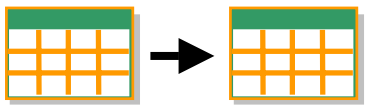

- **Unidirectional**
  - Changes are replicated in one direction between two servers (i.e. from source to target)
  - Changes can be filtered and transformed

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

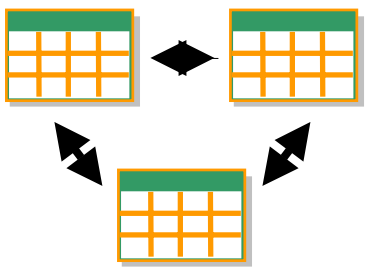

- **Peer to peer**
  - Changes are replicated between 2 or more servers
  - Utilizes **VERSION** based conflict detection

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

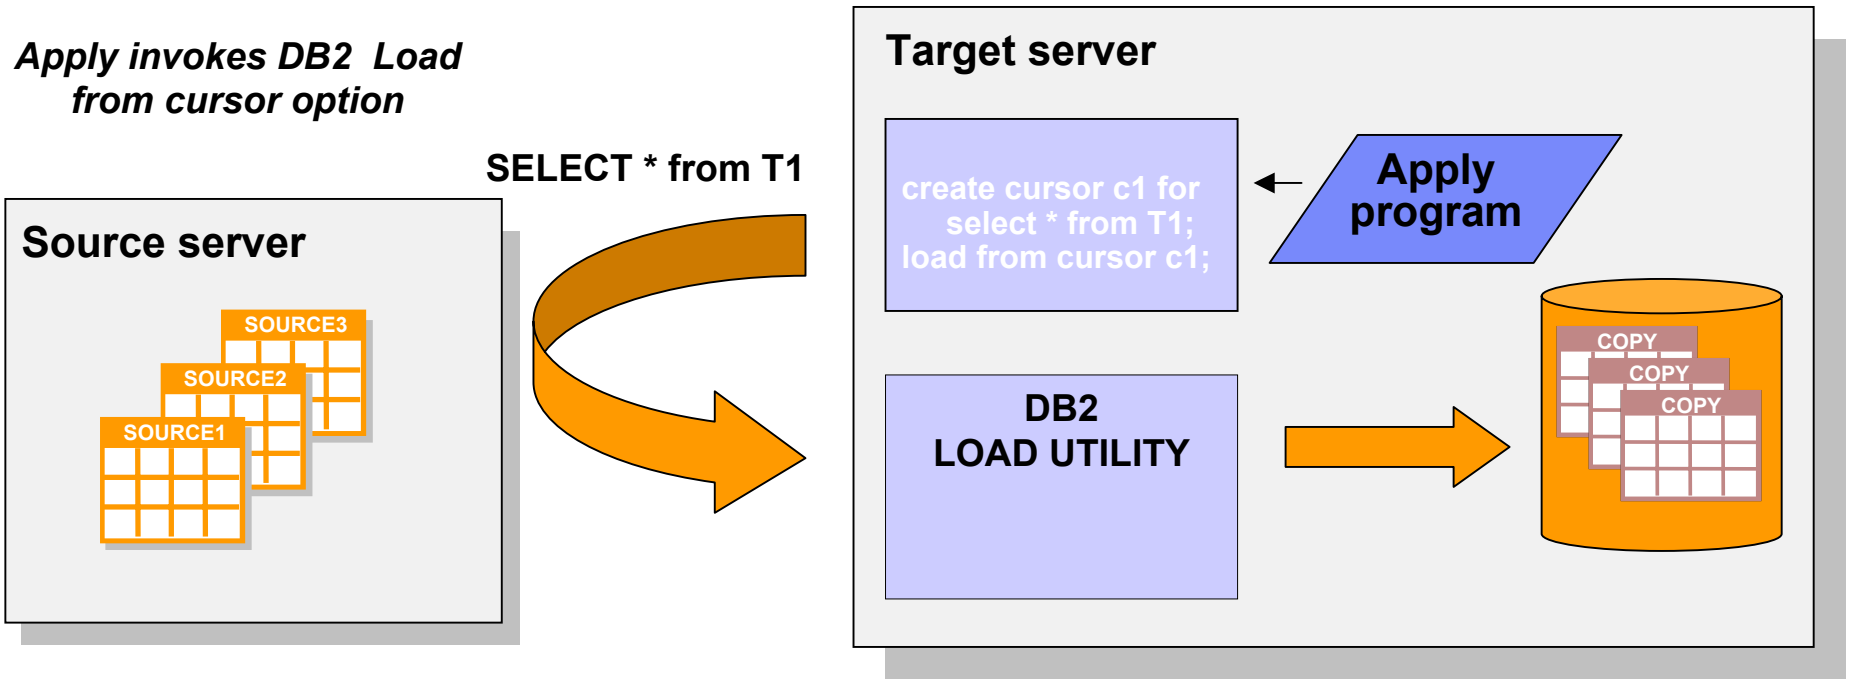

Selected data blocks feed directly into load utility

Subscription Start with Load Processing

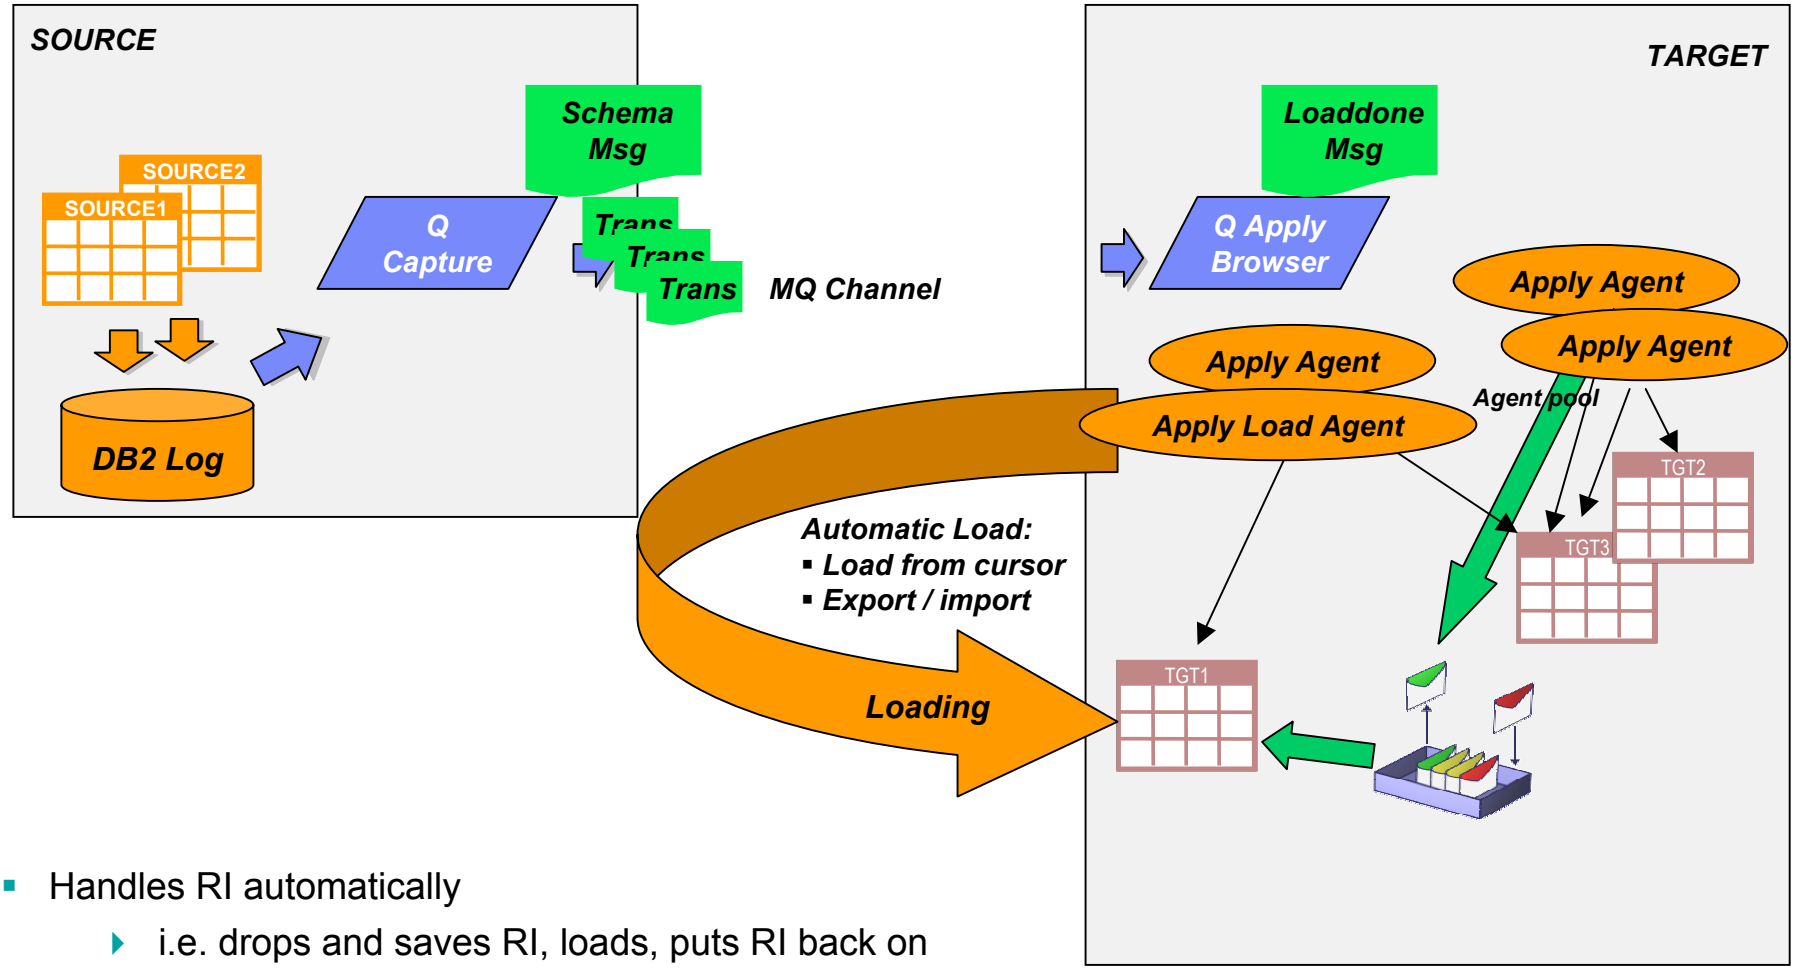

- Handles RI automatically
  - i.e. drops and saves RI, loads, puts RI back on
- Can handle Peer to Peer loading of n peers
  - Includes starting a new peer into an existing active configuration

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.

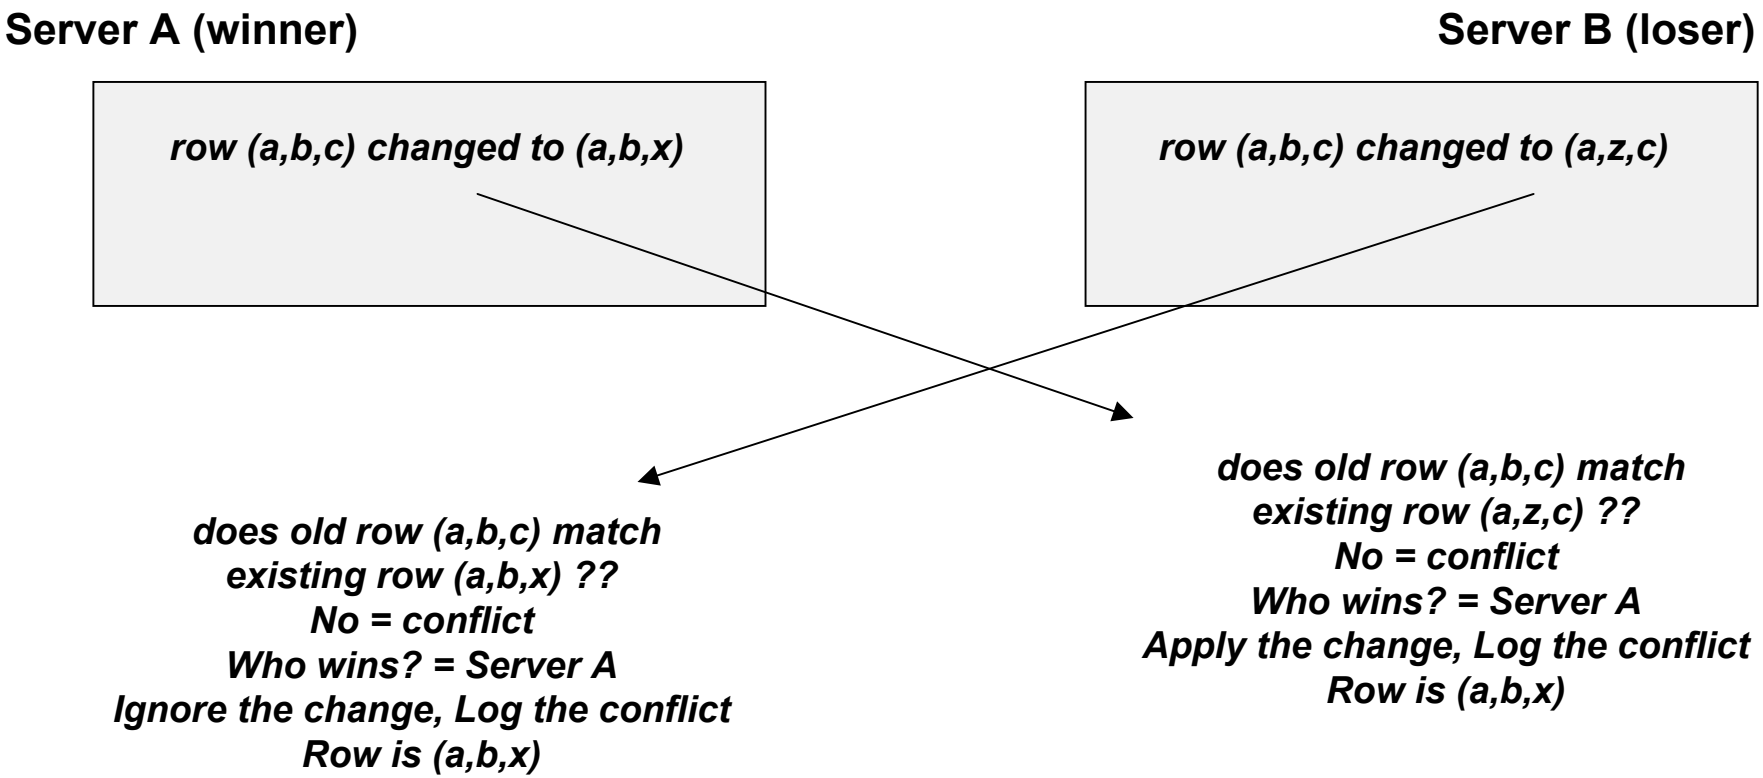

Value Based Conflict Detection and Resolution

- 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.

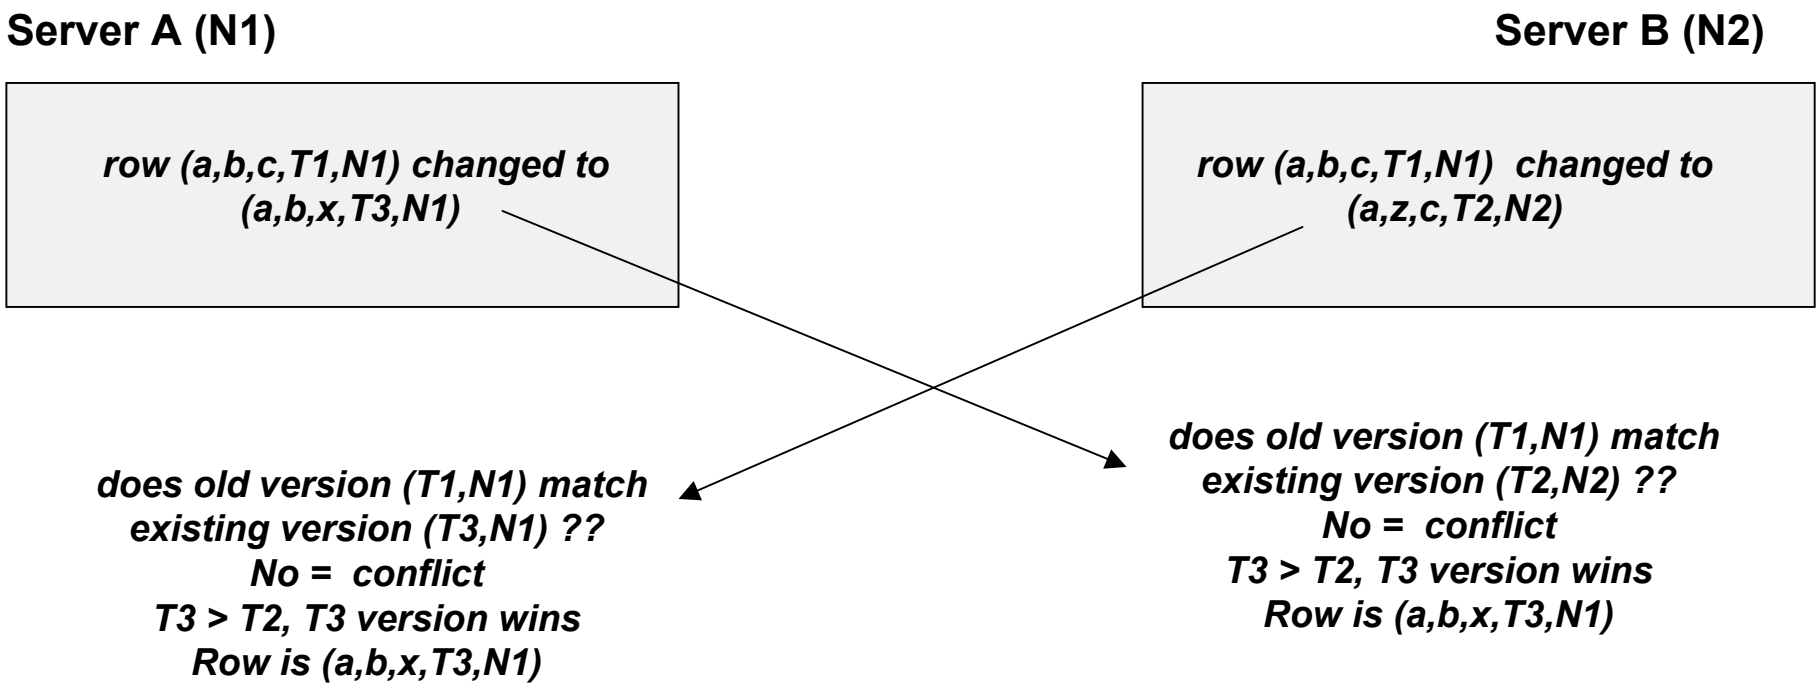

Replication Administration

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

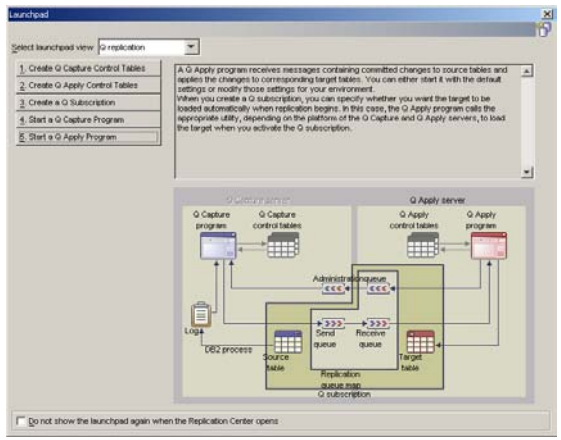

- Command Line Interface
  - Scripts or interactive mode
  - Example:

```
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

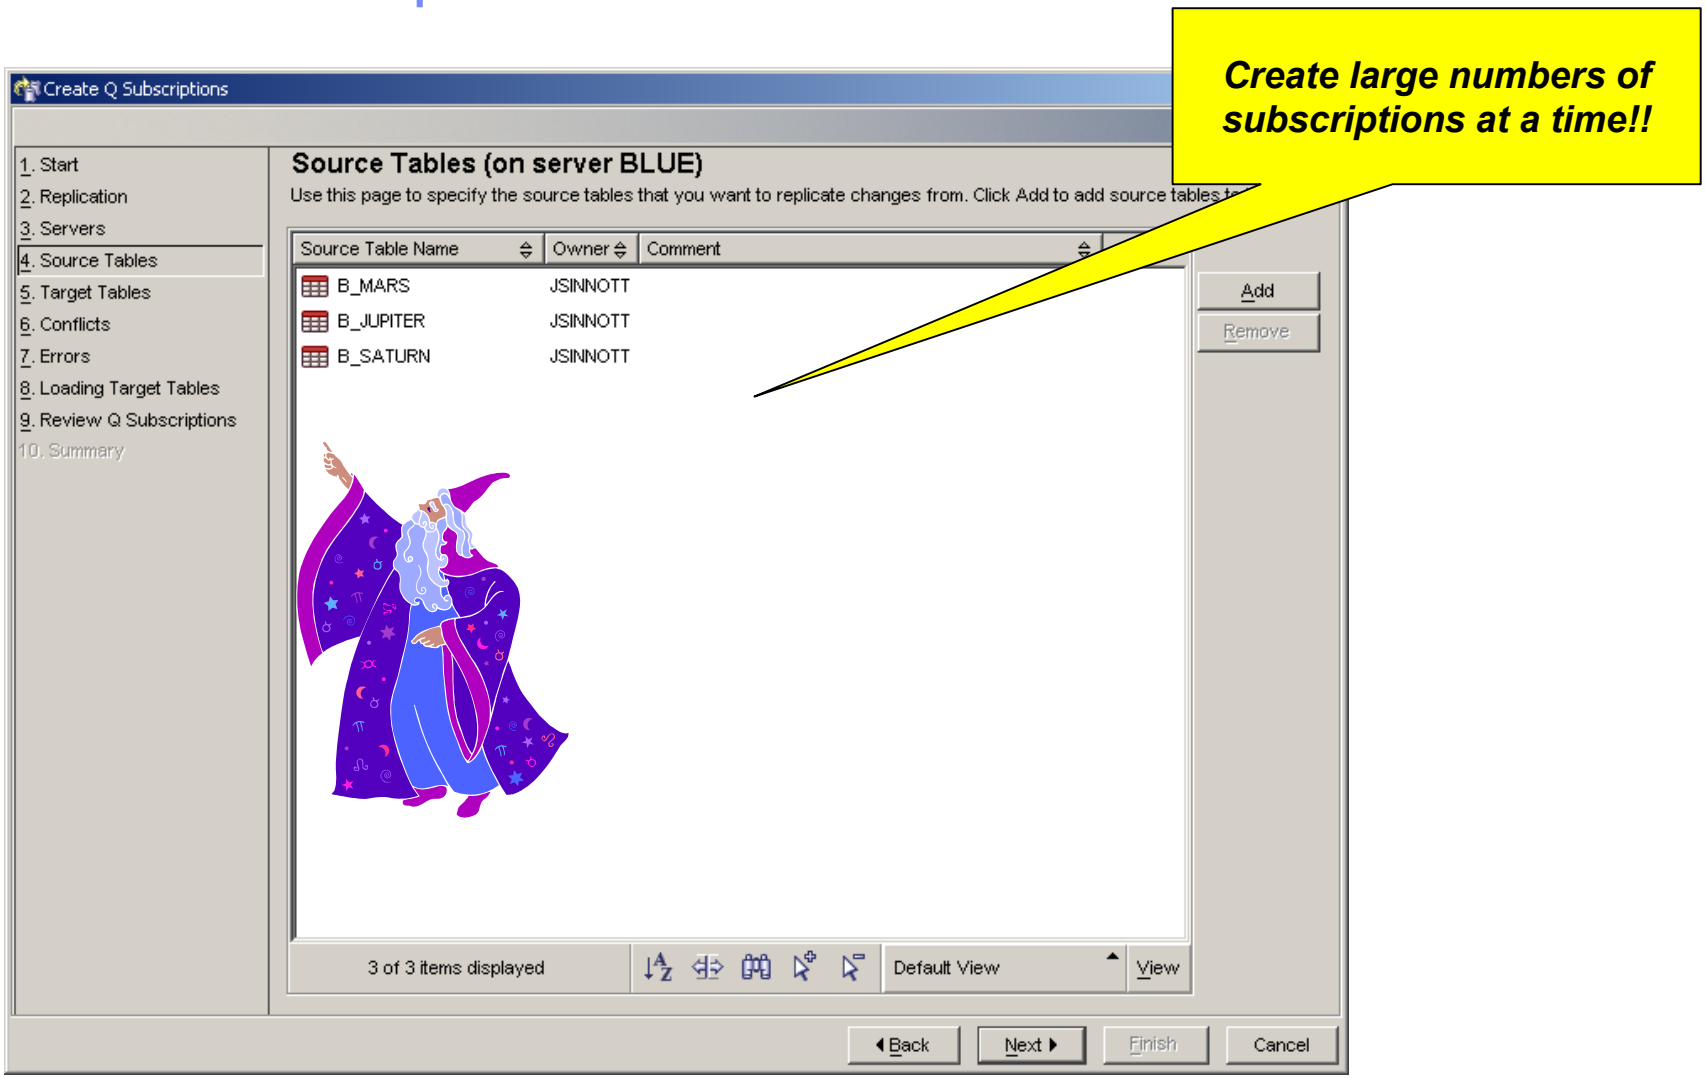

Implementierung der ‚Q Replication' im VM & VSE

- MQ Server im VSE
- MQ Client im VM und VSE
- Ab 7.4 Q-Capture/VM&VSE (kein Q-Apply!)
- Replikation ins VSE (Apply) über ‚Federated Database Support'
- ‚Event Publication' noch nicht verfügbar
- Ab 7.4 keine SQL-Replikation mehr!

Implementierung der ‚Q Replication' im VSE

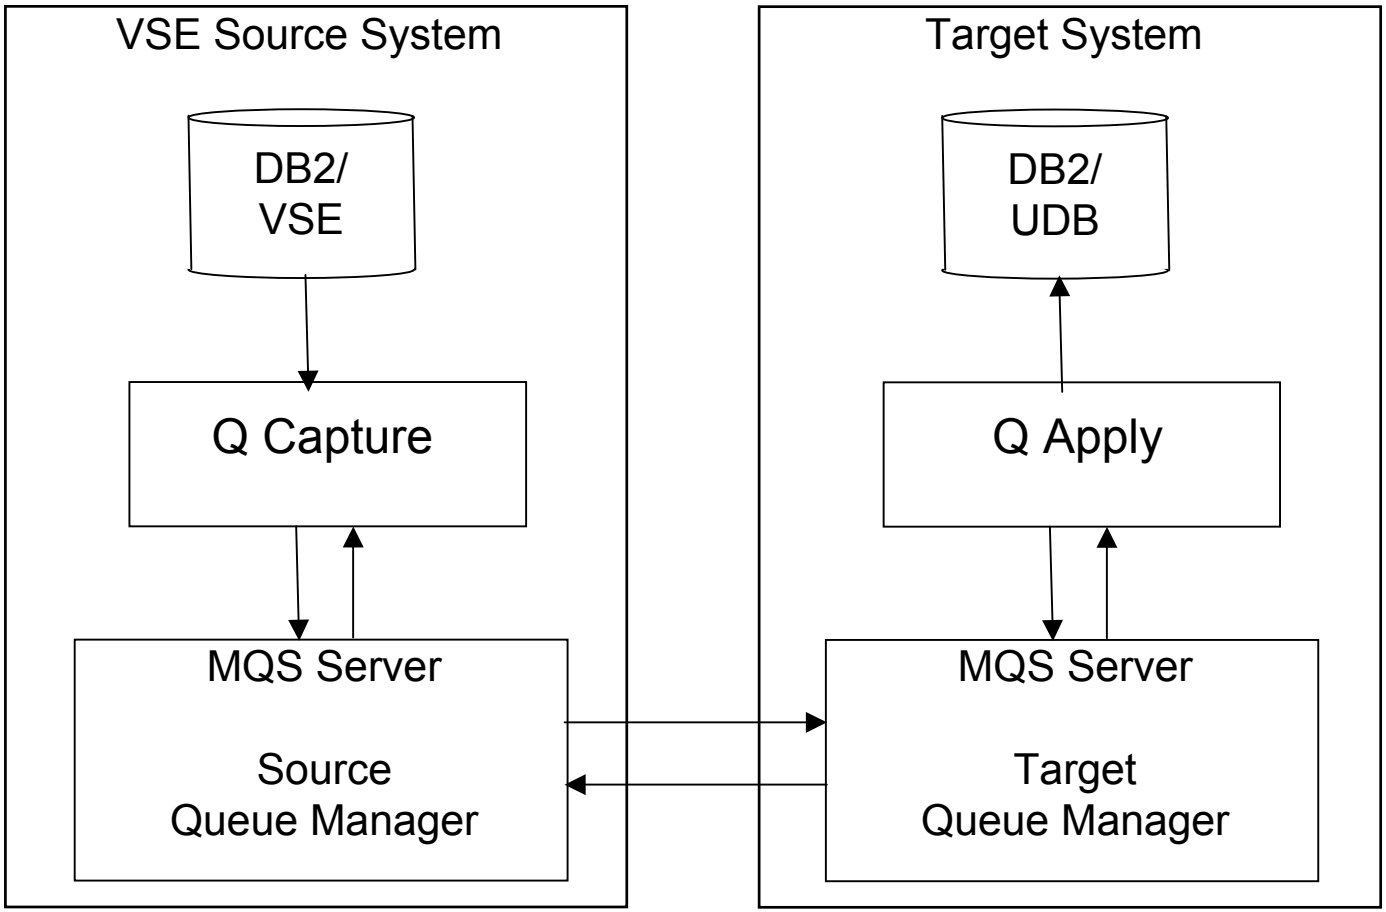

Implementierung der ‚Q Replication' im VM & VSE

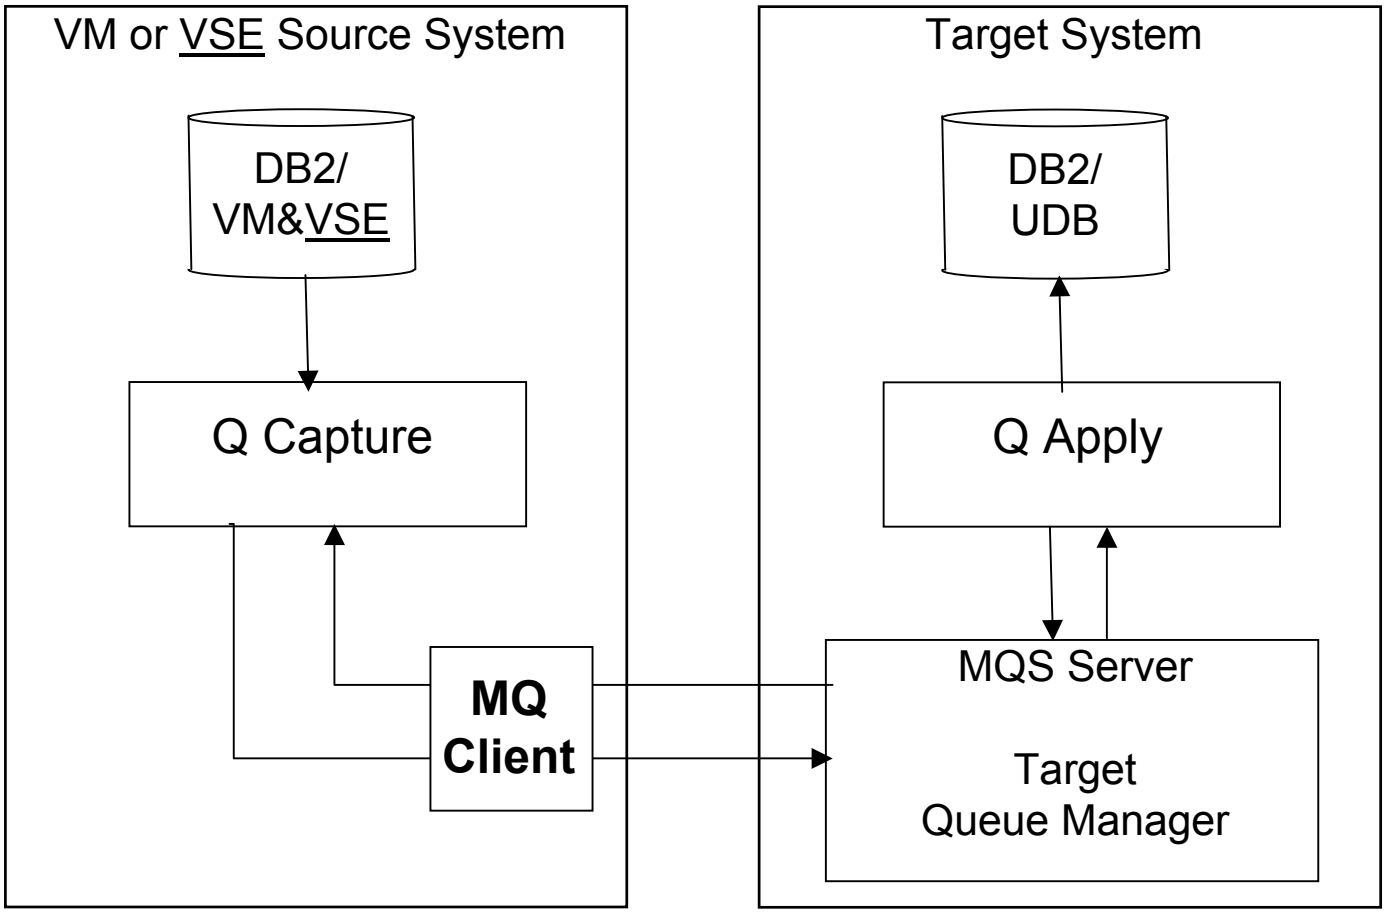

DB2 Server for VSE & VM – weitere Informationen

- 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 **7.2**

  31. März 2006

- Kontakt: roeber@de.ibm.com

Solutions on the new z/VSE homepage

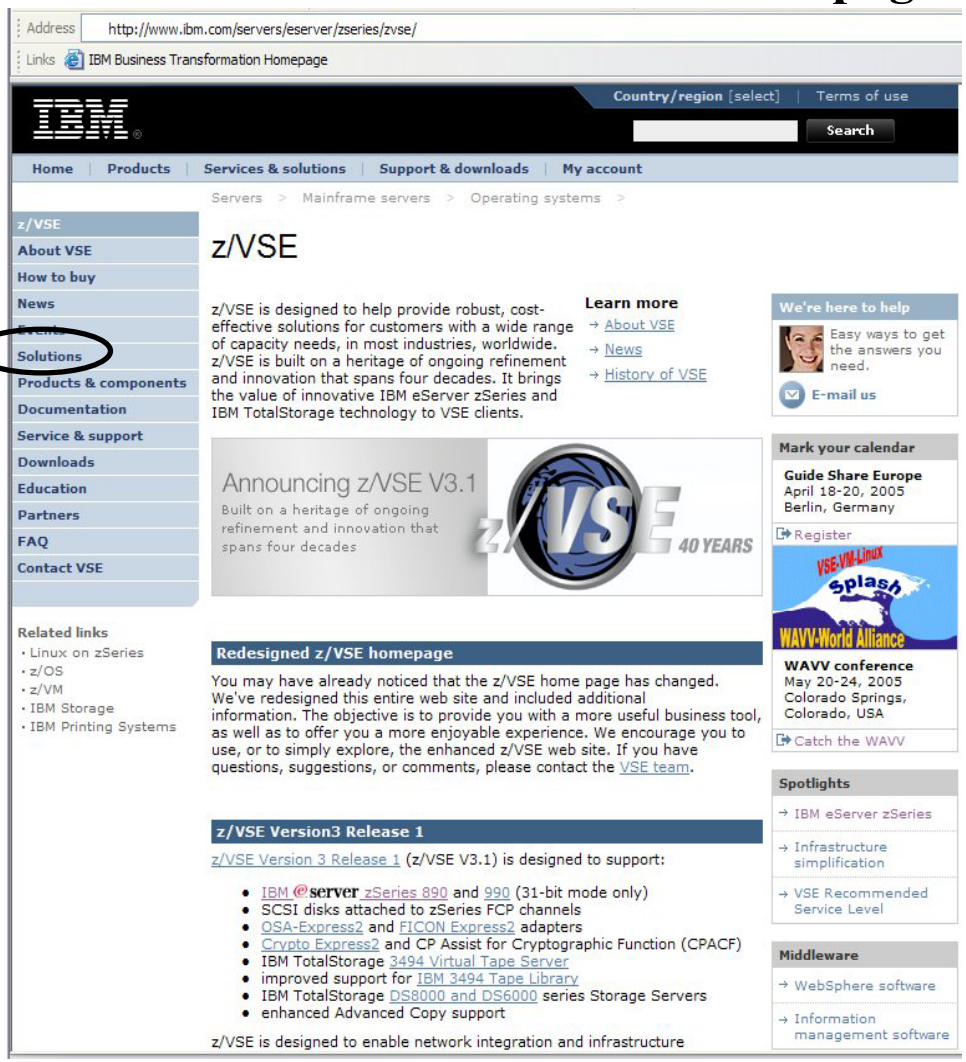

http://www.ibm.com/servers/eserver/zseries/zvse/

Additional Information

- **z/VSE Home Page**
  http://www.ibm.com/servers/eserver/zseries/zvse/

- **Solutions for VSE**
  http://www.ibm.com/servers/eserver/zseries/zvse/solutions/

- **e-business Connectivity for VSE/ESA** **SG24-5950**
- **e-business Solutions for VSE/ESA** **SG24-5662**
- **Servlet and JSP Programming** **SG24-5755**
- **Linux Web Hosting with WebSphere, DB2, and Domino** **SG24-6007**
- ***NEW: Websphere Handbook (Connectors to z/OS and VSE) SG24-7042***

We appreciate your comments at zvse@de.ibm.com

Haben Sie Fragen?