THE FILLMORE GROUP
Relational Database Solutions

Welcome

Changes and Choices

Premier Business Partner IBM

THE FILLMORE GROUP
Relational Database Solutions

Today's Session

Thursday, February 23, 2012

Premier Business Partner IBM

THE FILLMORE GROUP
Relational Database Solutions

Agenda

1. The Fillmore Group Introduction
2. Reasons to Implement Replication
3. IBM 's Replication Options – How We Got Here
4. The New Single Part Number Strategy
5. How to Choose Which is Best for YOU
6. InfoSphere Change Data Capture
7. InfoSphere Replication Server
8. Q&A

Premier Business Partner IBM

THE FILLMORE GROUP
Relational Database Solutions

The Fillmore Group

Frank C. Fillmore, Jr.

IBM CHAMPION

Premier Business Partner IBM

THE FILLMORE GROUP
Relational Database Solutions

The Fillmore Group, Inc.

- History
- DB2 Technical Support and Consulting
- IBM Authorized Training Partner
- IBM Information Management Software Reseller

Premier Business Partner IBM

THE FILLMORE GROUP
Relational Database Solutions

History

The Fillmore Group, Inc.

- Founded in Maryland, 1987
- IBM Business Partner since 1989
- Delivering IBM Education since 1994
- DB2 Gold Consultant since 1998

Premier Business Partner IBM

THE FILLMORE GROUP
Relational Database Solutions

Representative Replication Customers

- JP Morgan Chase – Q Rep for Oracle migration
- Iron Mountain – ICDC for platform migration
- FBI – Q Rep for continuous availability
- Sears – Q Rep to feed Netezza OLTP data

Premier Business Partner IBM

Reasons to Implement Replication

- Maintain production system availability 24x7
- Disaster Recovery Hot Site
- Database migrations – vendor, version-to-version, platform
- Loading a data warehouse
- Load balancing and facilitating extended use
- Integrating data from disparate systems

THE FILLMORE GROUP
Relational Database Solutions

What exactly IS replication?

From the Information Management Glossary:

The process of copying a portion of a database from one environment to another and keeping the subsequent copies of the data in sync with the original source. Changes made to the original source are propagated to the copies of the data in other environments.

IBM's Replication Options – How We Got Here

- **SQL Replication** – released in the mid-1990's as DataPropagator Relational
- **Queue Replication** - released along with WebSphere Information Integrator V8.2 in 2005, incorporating MQ and SQL Replication to deliver high-speed replication
- **DataMirror/ICDC** – IBM acquired DataMirror in 2007 and rebranded DataMirror Transformation Server as InfoSphere Change Data Capture (ICDC) in 2008.

THE FILLMORE GROUP
Relational Database Solutions

InfoSphere Product Rebranding

- SQL Replication
  - aka SQL Replication
- Q Replication
  - aka InfoSphere Replication Server, Q Replication, Q Rep
- ICDC
  - aka InfoSphere Change Data Capture, ICDC

Premier Business Partner IBM

THE FILLMORE GROUP
Relational Database Solutions

Single Part Number Strategy

December, 2011 IBM Announcement

http://www-01.ibm.com/common/ssi/rep_ca/4/897/ENUS211-504/ENUS211-504.PDF

"IBM InfoSphere Data Replication V10.1.2" (IIDR)

- Consolidates SQL Replication Q Replication & ICDC into a single part number for purchasing

- Same price for all (P/N D0L34LL; $168 per PVU)

THE FILLMORE GROUP
Relational Database Solutions

How to chose which is best for you

- Understand the basic replication terminology
- Evaluate and prioritize your replication goals
- Understand each solution's strengths
- Apply your product understanding to your priorities
- Use the checklist
- Consult IBM and Business Partner experts

Premier Business Partner IBM

Terminology

- **Latency**
  - The time it takes for data to get from one point to another
  - Synonymous with delay; measured in microseconds, seconds, minutes
- **Source and Target databases**
  - The "source" database on which data is initially stored and the "target" is the database to which data is replicated
- **Capture and Apply**
  - Terms used to describe the process of collecting and delivering changed data
- **Transport**
  - The methodology used to move the data from source to target
- **Uni-directional, bi-directional and peer-to-peer**
  - Description of data movement from either source to target (uni) or from both source to target and target to source (bi), or to n-tier

Evaluate your replication goals – consider:

- Source and target databases
  - Are the source and target homogeneous or heterogeneous?
- Uni-directional, bi-directional or peer-to-peer replication?
- Speed and latency
  - Is there an SLA for replication speed?
  - What is the data volume on the source and target databases?
    - GB/TB
  - What is your anticipated transaction volume?
    - INSERTS/UPDATES/DELETES

Evaluate your replication goals – continued:

- Resiliency
  - Are outages acceptable? How often? For how long?
  - Do you need automated failover in the event of an unplanned outage?
- Ease of use
  - Monitoring, installation, administration
  - How frequently do you change your data model?
  - How complex is conflict resolution? Is it based on:
    - Value/Source/Timestamp/Application Logic
- Cost
  - No longer a factor

Understand Each Solution's Strengths

1. SQL Replication
2. Q Replication
3. ICDC

Apply Your Understanding to Your Priorities

1. Are my source and target databases supported?
2. Am I doing uni-directional, bi-directional or peer-to-peer replication?
3. Will my choice deliver the speed and latency required to meet my SLA's?
4. Is my selection resilient enough?
5. Is my selection going to be easy enough to use for the staff supporting replication?

Checklist

When in doubt, use the checklist. Available at:

www.thefillmoregroup.com/blog

THE FILLMORE GROUP
Relational Database Solutions

SQL Replication

- Sources: DB2 (all platforms), Oracle, SQL Server, Sybase, Informix
- Targets: DB2 (all platforms), Oracle, SQL Server, Sybase, Informix

- Uni-directional, bi-directional, peer-to-peer replication supported

- Latency – 3X slower than Q Replication; not recommended for very large transaction volumes

- Resiliency – works with HADR

- Ease of use – Monitoring console

- **Cost - *FREE* with DB2 for LUW; Capture component included with DB2 for z/OS but requires purchase of Apply; may need InfoSphere Federation Server for heterogeneous targets.**

SQL Replication

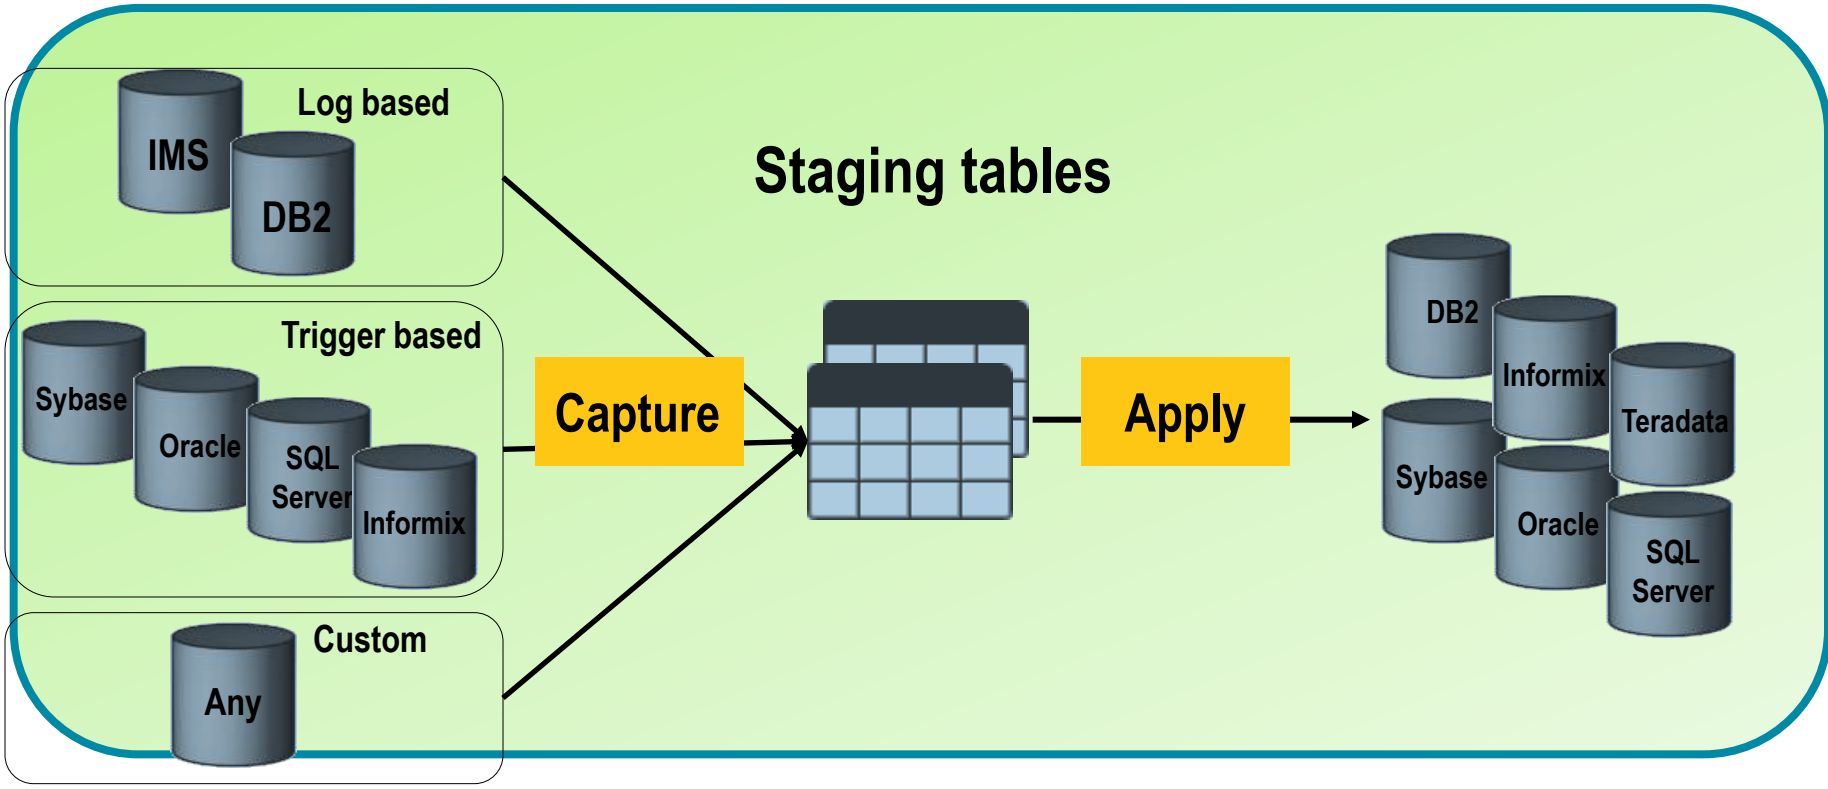

SQL Replication

Important

- Triggers are needed when the source is not DB2.
- InfoSphere Federation Server is needed when the target is not DB2.

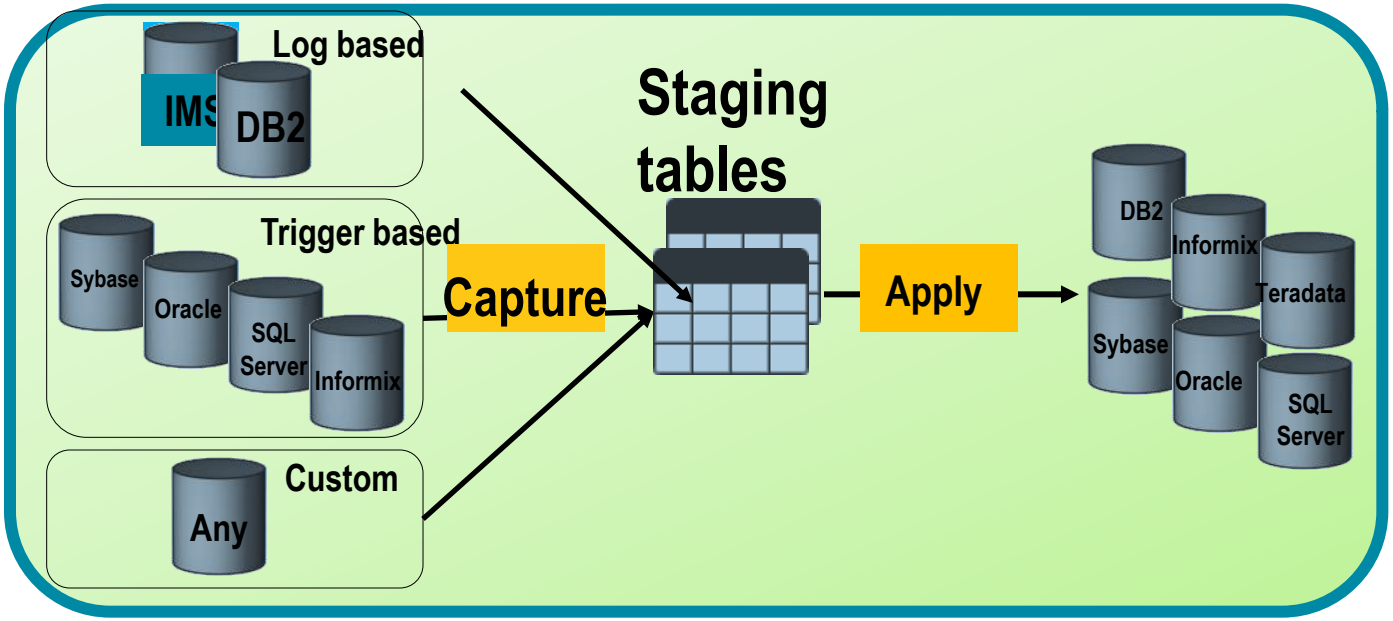

Q Replication

- Sources: DB2 for z/OS and LUW, and Oracle on all platforms
- Targets: DB2 for z and LUW, Oracle, Sybase, others with Federation

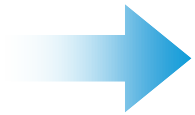

Can invoke Stored Procedures or publish XML

- Uni-directional, bi-directional and peer-to-peer replication supported
- Latency – 3X faster than SQL Replication; recommended for very large transaction volumes; uses MQ Series for speedy delivery
- Resiliency – works with HADR, Q Replication Dashboard monitoring
- Ease of use – ASNCLP scripting language

Q Replication

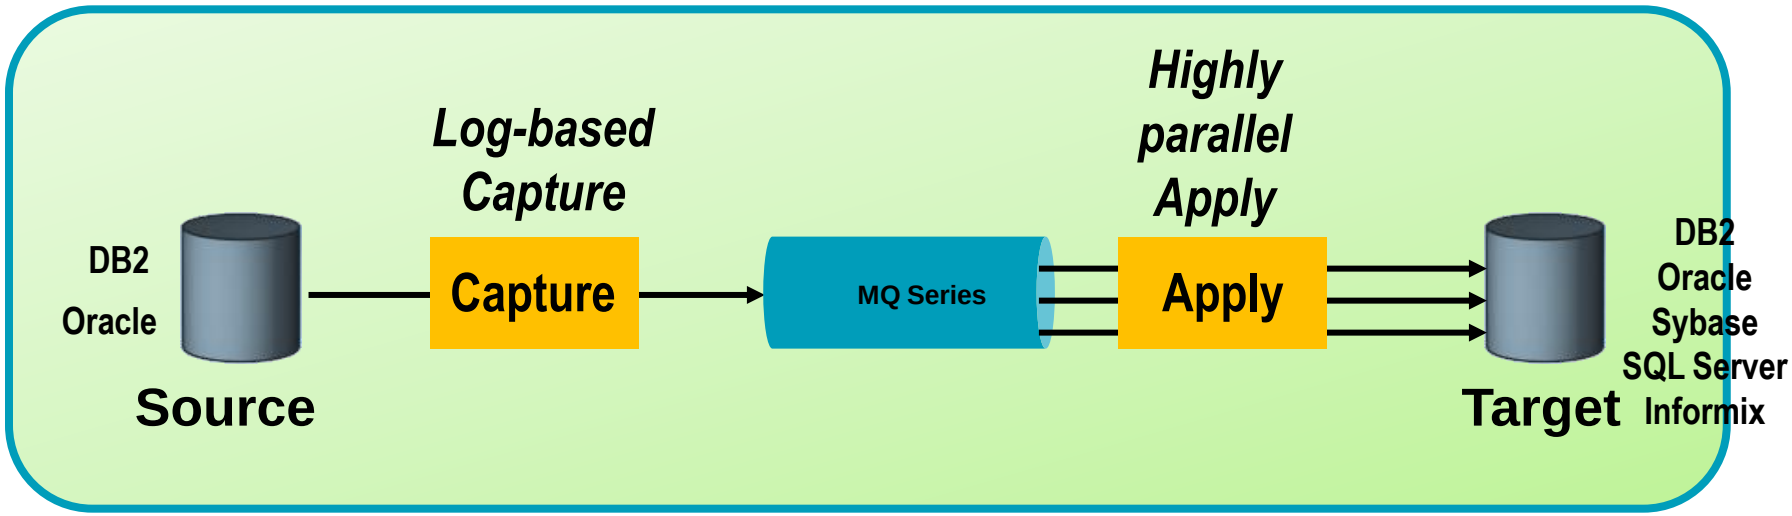

Q Replication

Important

- Additional complexity due to WebSphere MQ
- Fastest; lowest latency
- Replicates compressed DB2 v9.7 data
- No support (source or target) for DB2 on System i

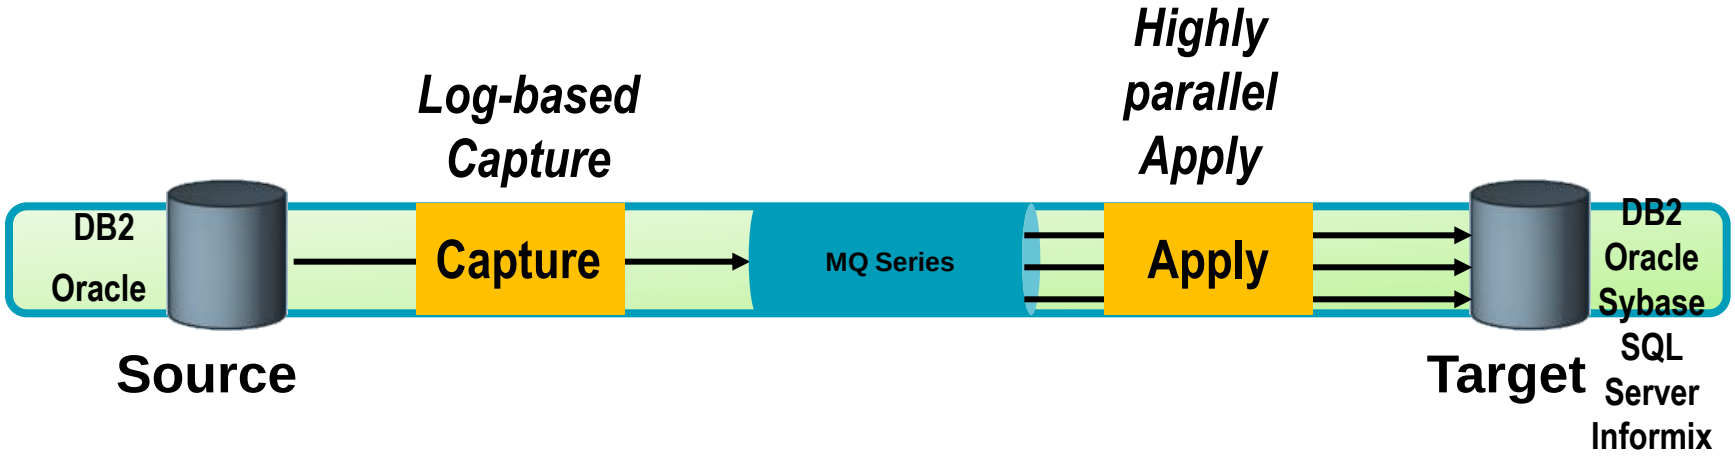

THE FILLMORE GROUP
Relational Database Solutions

InfoSphere Change Data Capture

•Sources: DB2 (all platforms), Oracle, SQL Server, Sybase, Informix
•Targets: DB2, Oracle, SQL Server, Sybase, Teradata, Netezza

And non-relational data targets such as DataStage and MQ

•Uni-directional and bi-directional replication supported

•Latency – faster than SQL Rep; vs. Q Replication scalability issues may impact performance at very high volumes

•Resiliency – works with HADR for DB2

•Ease of use – great GUI interface; no equivalent scripting to ASNCLP

InfoSphere Change Data Capture Supported Sources and Targets

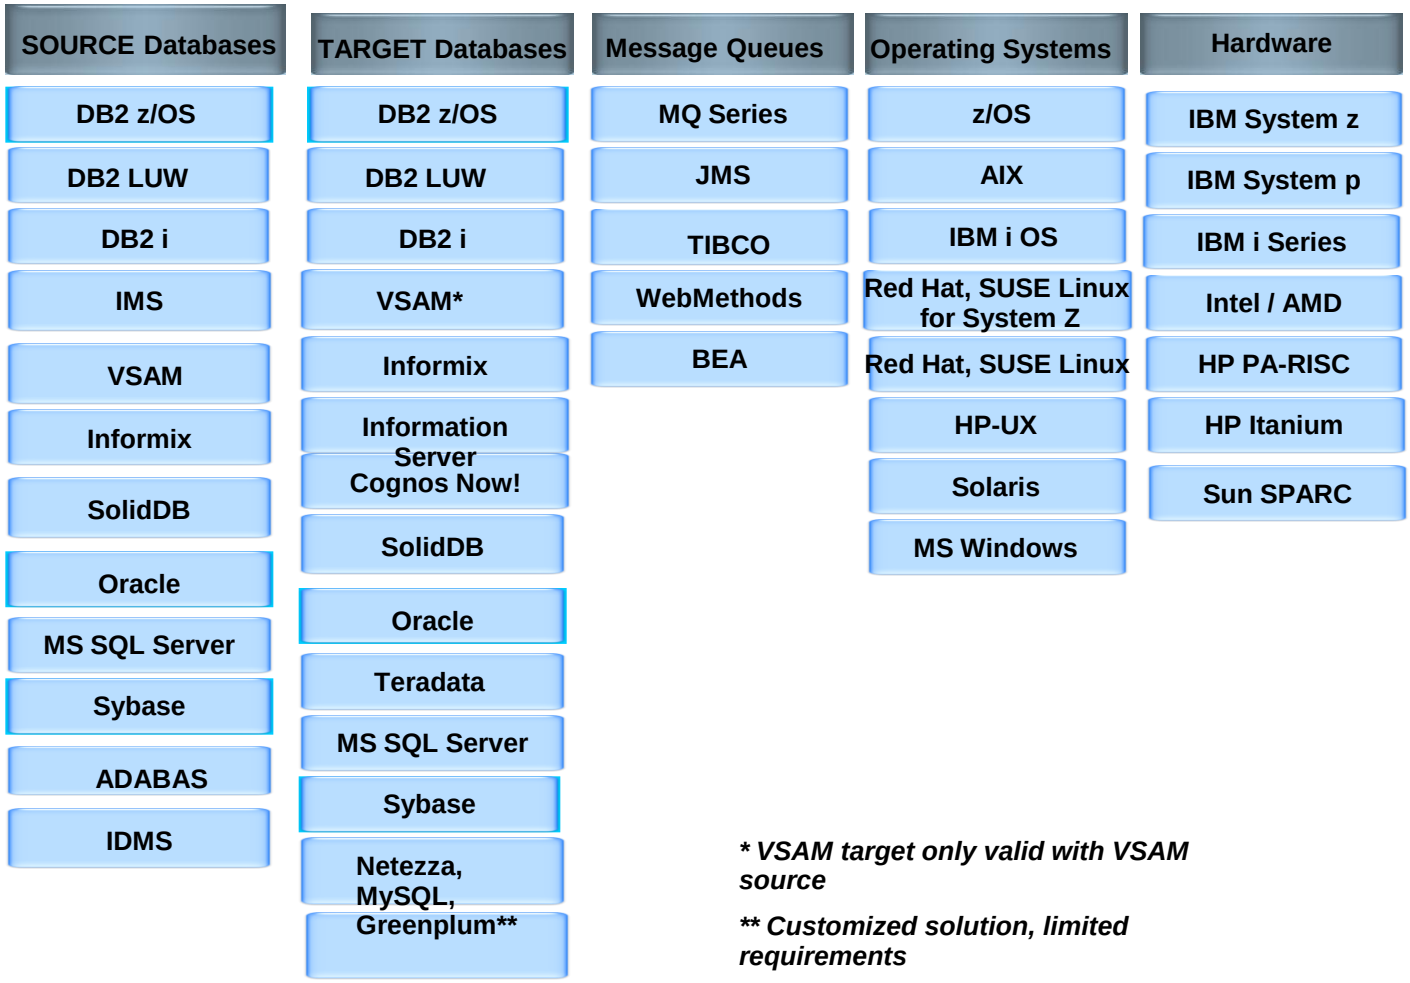

** VSAM target only valid with VSAM source*

*** Customized solution, limited requirements*

InfoSphere Change Data Capture

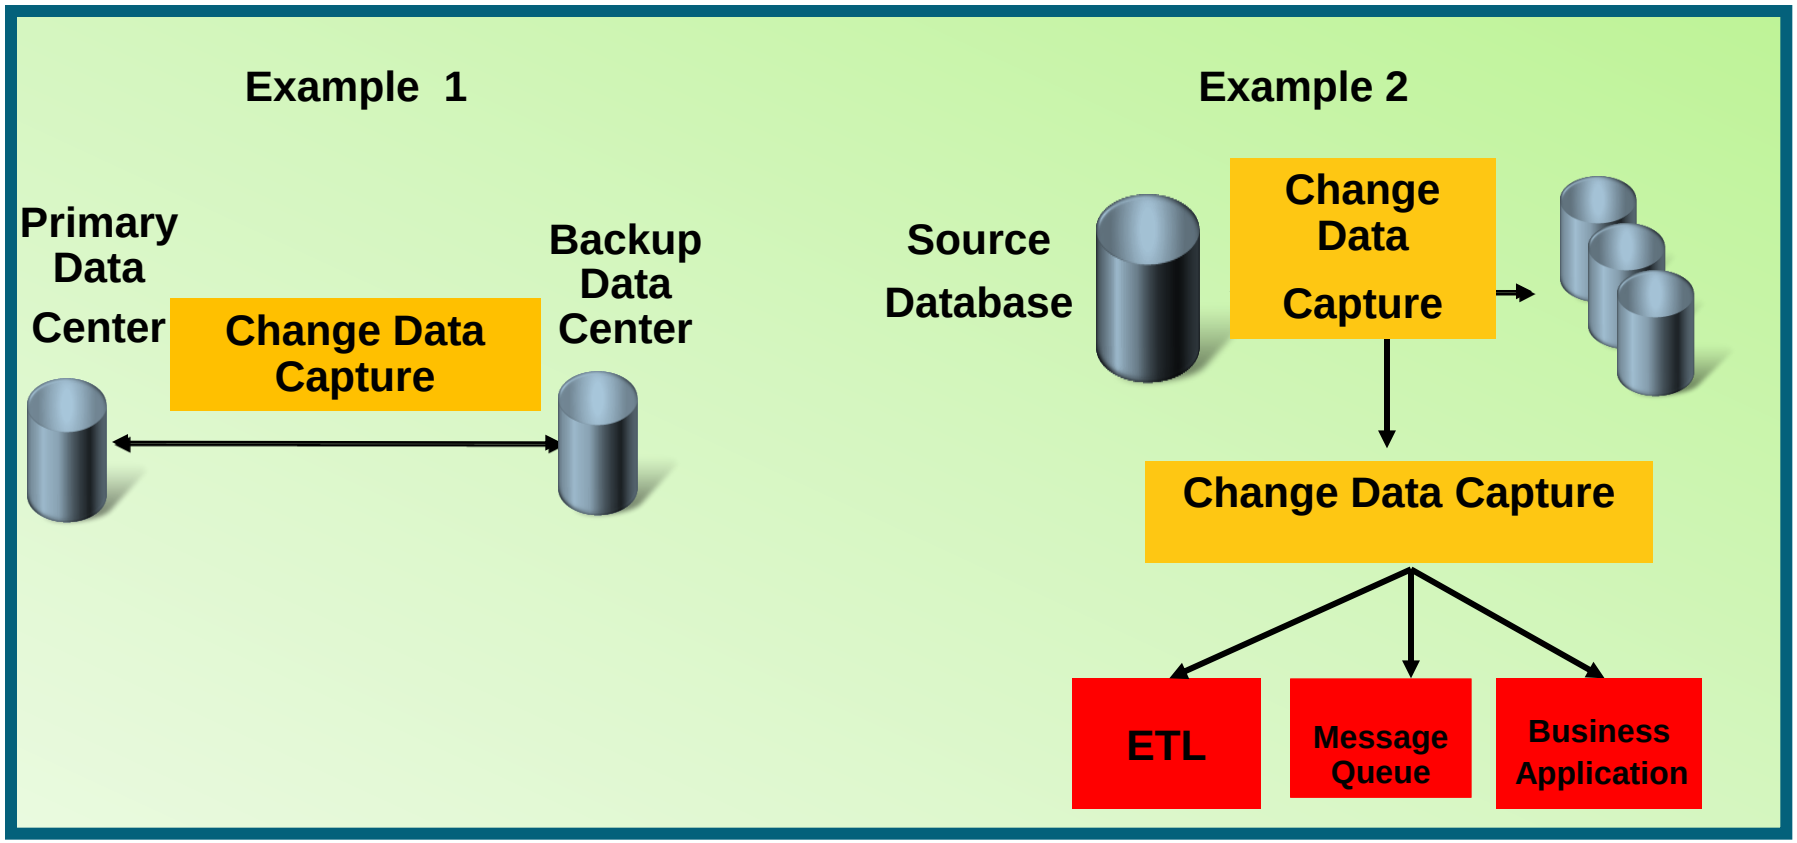

InfoSphere Change Data Capture

Important

- Updates a bookmark with any change being delivered to the target.
- ICDC uses a small, proprietary database to manage the log position of each change.
- ICDC's performance is dependent on this database and it may constrain the volume and speed of transactions.

→ And non-relational data sources such as DataStage and MQ

- Does not require triggers on a heterogeneous source.
- Does not require InfoSphere Federation Server for heterogeneous targets.

ICDC Architecture

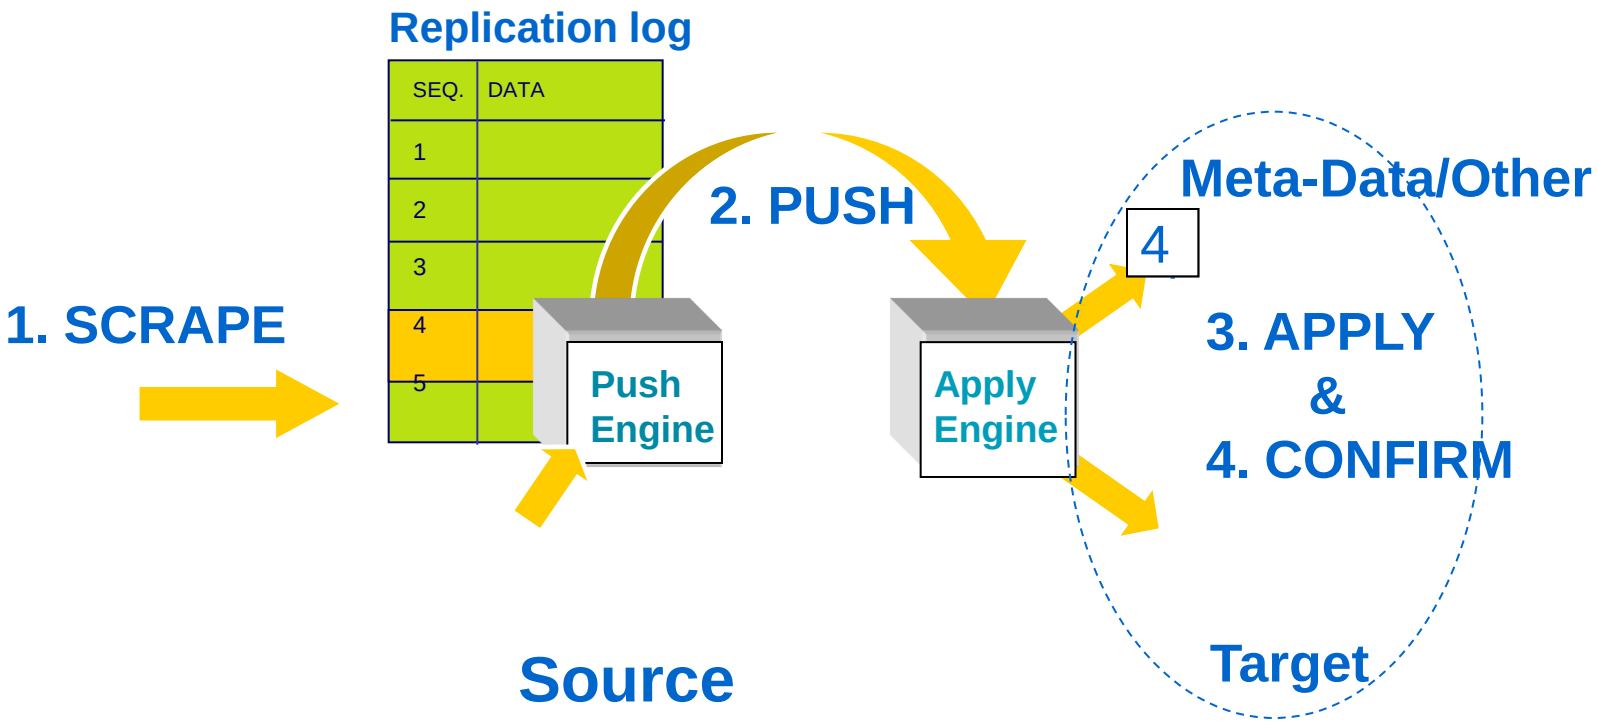

Customer Example – SQL Replication

- State government web portal
  - Source and target = DB2 Enterprise Server Edition for LUW – both source and target on z/Linux
  - Uni- directional replication from a database behind a firewall in the state capital to a public portal database located outside the firewall
  - Speed and latency – Not a critical issue; data volume low
  - Resiliency – Not a critical issue; replacing an old system with a planned 2 hour outage each night
  - Ease of use – relatively important; heritage mainframe DBA staff already being tasked with developing new DB2 for LUW skills and frequent changes need to be incorporated
  - Cost – Agency had a minimal IT budget

Customer Example – Q Replication

- National government customs and immigration
  - Source and target = DB2 for z/OS
  - Bi-directional replication for warm failover
  - Speed and Latency - built to ensure continuous availability during anticipated peak traffic times
  - Resiliency – critical particularly during high traffic periods
  - Ease of use – not as important as performance
  - Cost – well funded customer

Customer Example – ICDC

- Appliance Parts Retailer extending web presence
  - Source = DB2 on System i, target = DB2 Express C
  - Bi-directional, using replication to populate OLTP to web portal
  - Speed and Latency – up-to-the-second replication not critical
  - Resiliency – important during peak web use hours but not a priority
  - Ease of use – critical as data inventory items change constantly
  - Cost – budget < $25K

THE FILLMORE GROUP
Relational Database Solutions

Consult IBM and Business Partner Experts

- David Tolleson, IBM Data Replication Product Manager
  tolleson@us.ibm.com
- Mark Ketchie, IBM Replication Client Technical Specialists
  mketchie@us.ibm.com
- Frank Fillmore, The Fillmore Group
  frank.fillmore@thefillmoregroup.com

Premier Business Partner IBM

THE FILLMORE GROUP
Relational Database Solutions

Resources

- PVU Explanation

  http://www-01.ibm.com/software/lotus/passportadvantage/pvu_licensing_for_customers.html
- Link to IBM Replication Page

  http://www-01.ibm.com/software/data/infosphere/data-replication/
- developerWorks

  http://www.ibm.com/developerworks/data/roadmaps/qrepl-roadmap.html

Premier Business Partner IBM

Resources

- IBM Redbooks:
  www.redbooks.ibm.com

- IBM Authorized Training:
  www.ibm.com/training

The Fillmore Group website
www.thefillmoregroup.com

IBM Learning Services website
www.ibm.com/training