z/OS Data Replication as a Driver for Business Continuity

Karen Durward
IBM

August 9, 2011
Session Number 9665

Important Disclaimer

The information contained in this presentation is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided "as is" without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other documentation. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM (or its suppliers or licensors), or altering the terms and conditions of any agreement or license governing the use of IBM products and/or software.

IBM, the IBM logo, ibm.com, InfoSphere, DataStage, MetaStage, QualityStage, Information Agenda, and Information on Demand are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

Other company, product, or service names may be trademarks or service marks of others.

Agenda

- IBM's z/OS Data Replication for DB2 and IMS High Availability
- Typical usage examples
- Shifting to continuous availability …
  Data Replication as part of GDPS Active-Active Sites

InfoSphere Replication Server for z/OS
InfoSphere IMS Replication for z/OS

IBM Replication for High Availability

- **Focus is on**
  - Mirroring the data … minimal or no transformation
  - Very high throughputs … must keep up with enterprise workloads
  - Very low latency … less than one or two second latency is typical

- **Common characteristics**
  - Log-based captures … non-intrusive – no application changes
  - Parallel apply engines … keep up with the workload
  - Recoverable … track where apply "left off" as the point of recovery
  - Asynchronous … unlimited site separation

InfoSphere Replication Server for z/OS

Synchronize like-to-like copies

Distribution, consolidation or synchronization of information in different databases

- **Multi-directional delivery:**
  - Unidirectional
  - Bidirectional
  - Peer-to-Peer
- **Ease-of-use features:**
  - Integrated monitoring & statistics
  - Changed data histories
  - Configuration options:
    - Wizard-driven GUI
    - Command-line processor
    - Script-driven processor

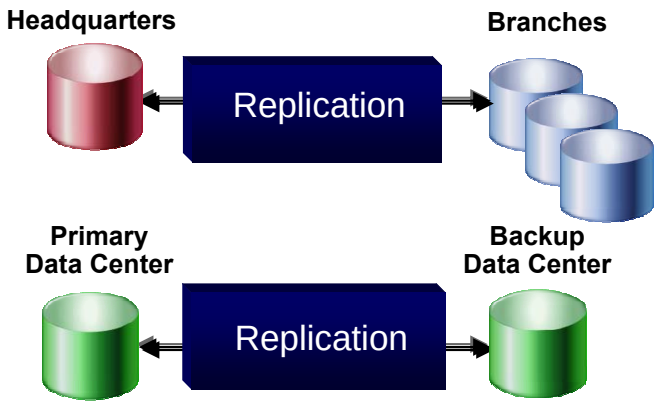

SQL Replication

Enables fan-out and heterogeneous replication
Too many "moving parts" for high availability

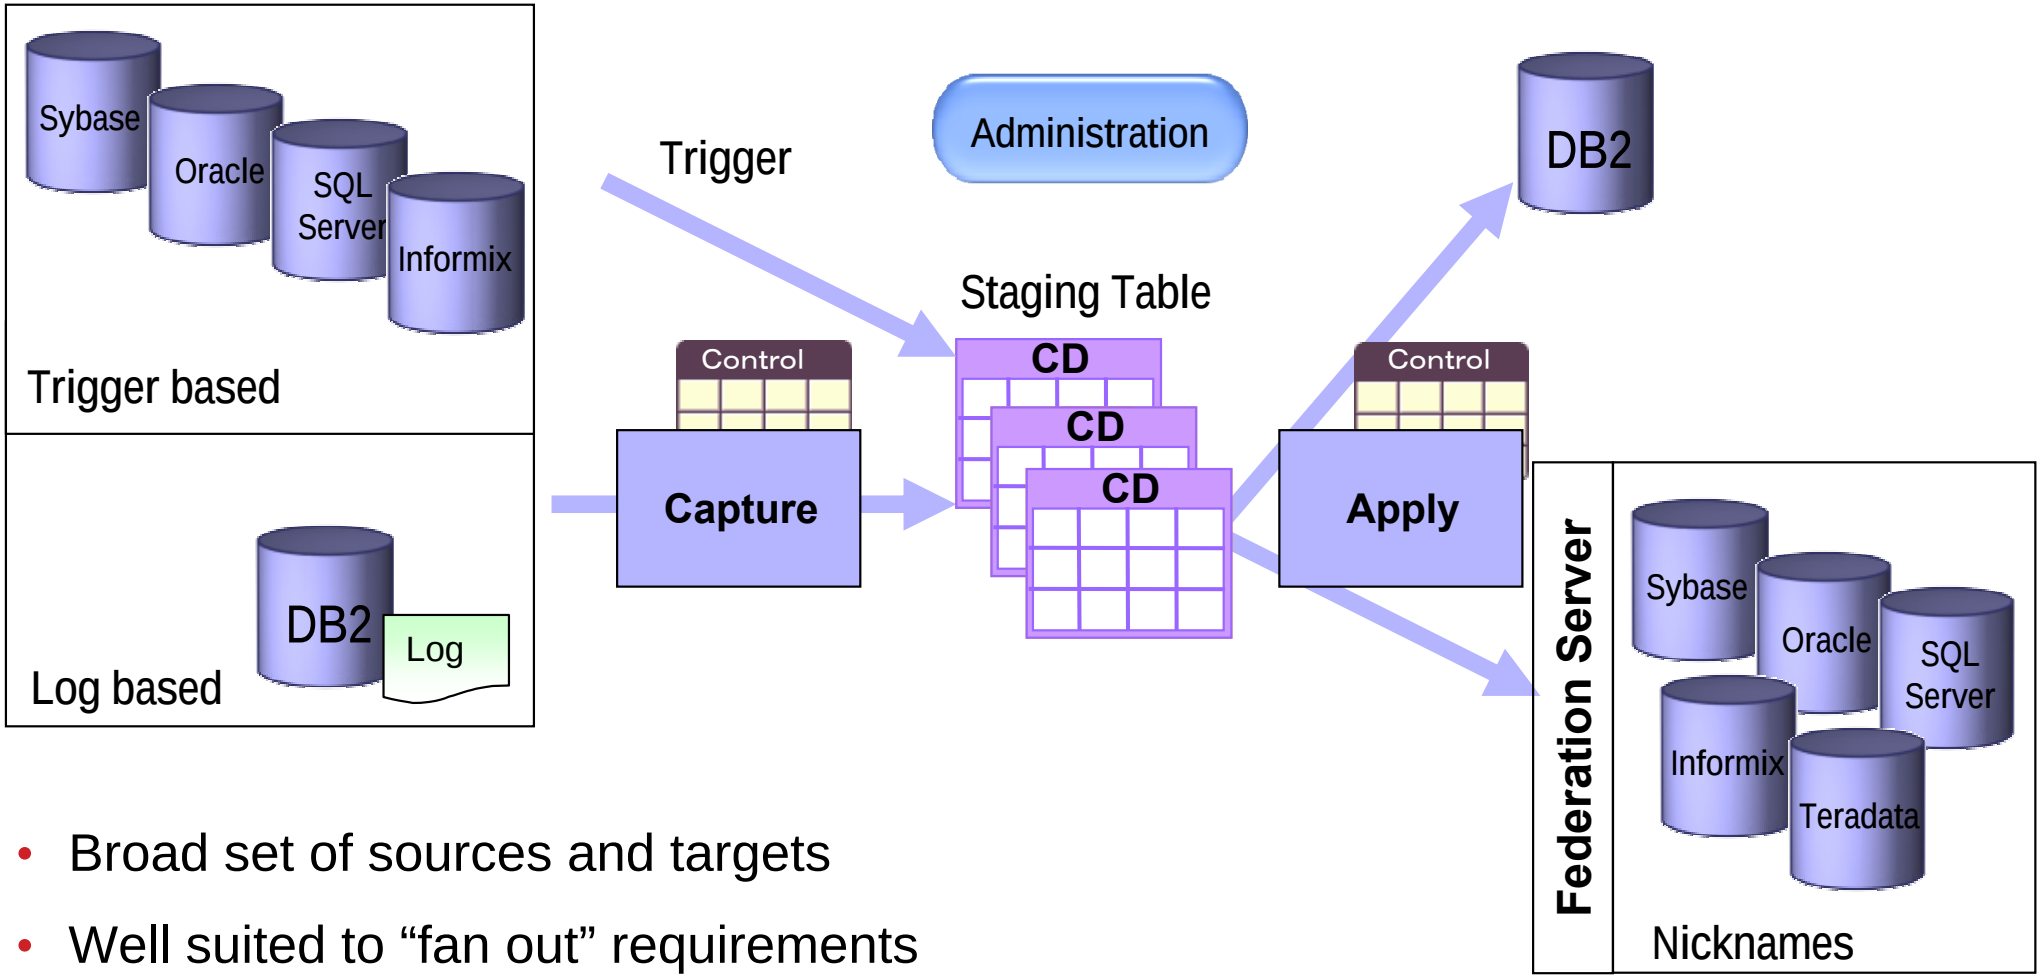

- Broad set of sources and targets
- Well suited to "fan out" requirements
- Flexible scheduling, transformation, distribution

Queue Replication

Ideal for High Availability DB2 Data Synchronization

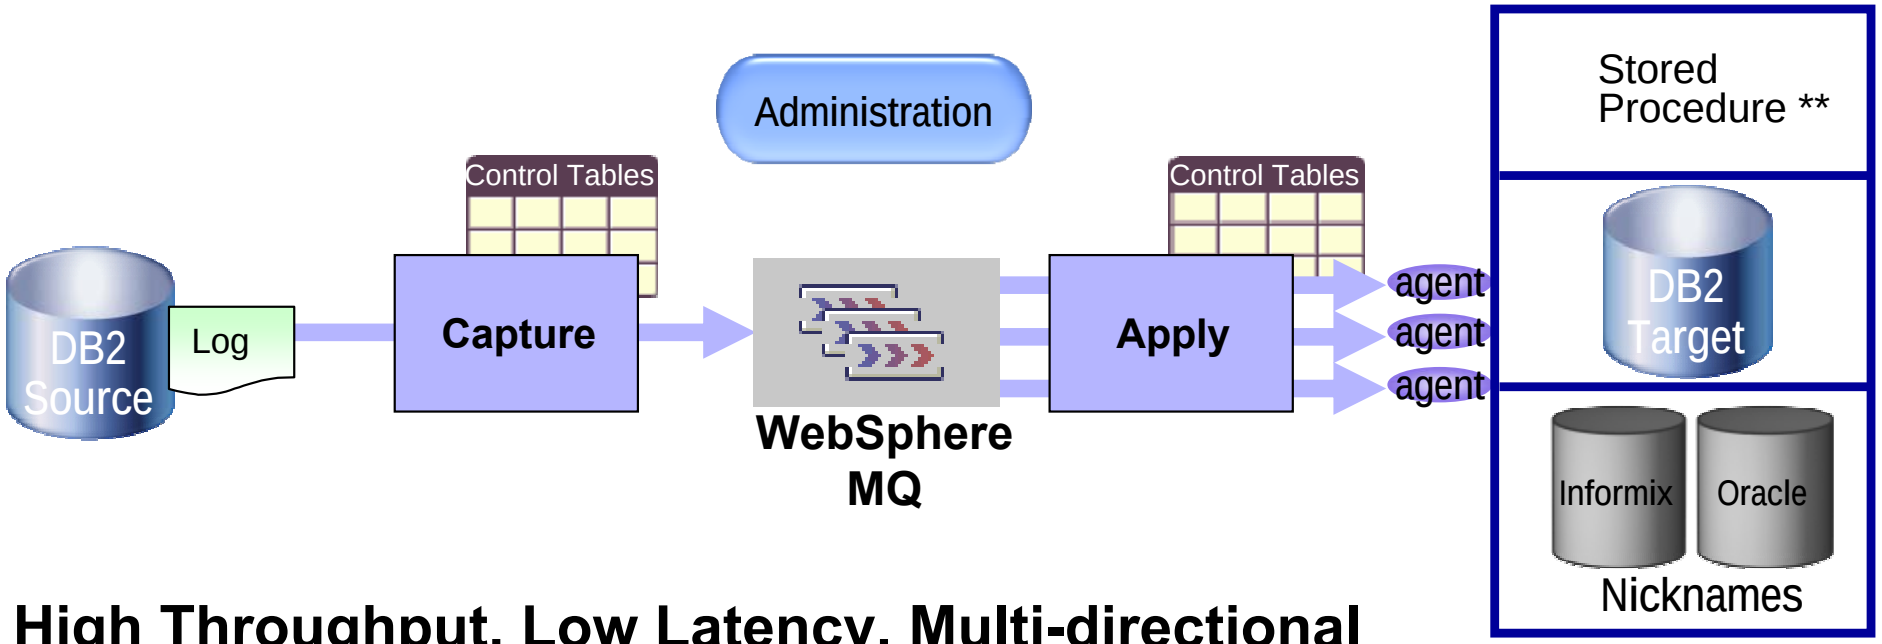

- **High Throughput, Low Latency, Multi-directional**
  - Unidirectional
  - Bidirectional
  - Peer-to-Peer
- **Features:**
  - Log based capture mechanism
  - Highly parallel apply process for high speed and low latency
  - Integrated monitoring & statistics
  - Changed data histories
  - Best of breed conflict detection and resolution

Queue Replication

Some Details of Highly Parallel Q Apply

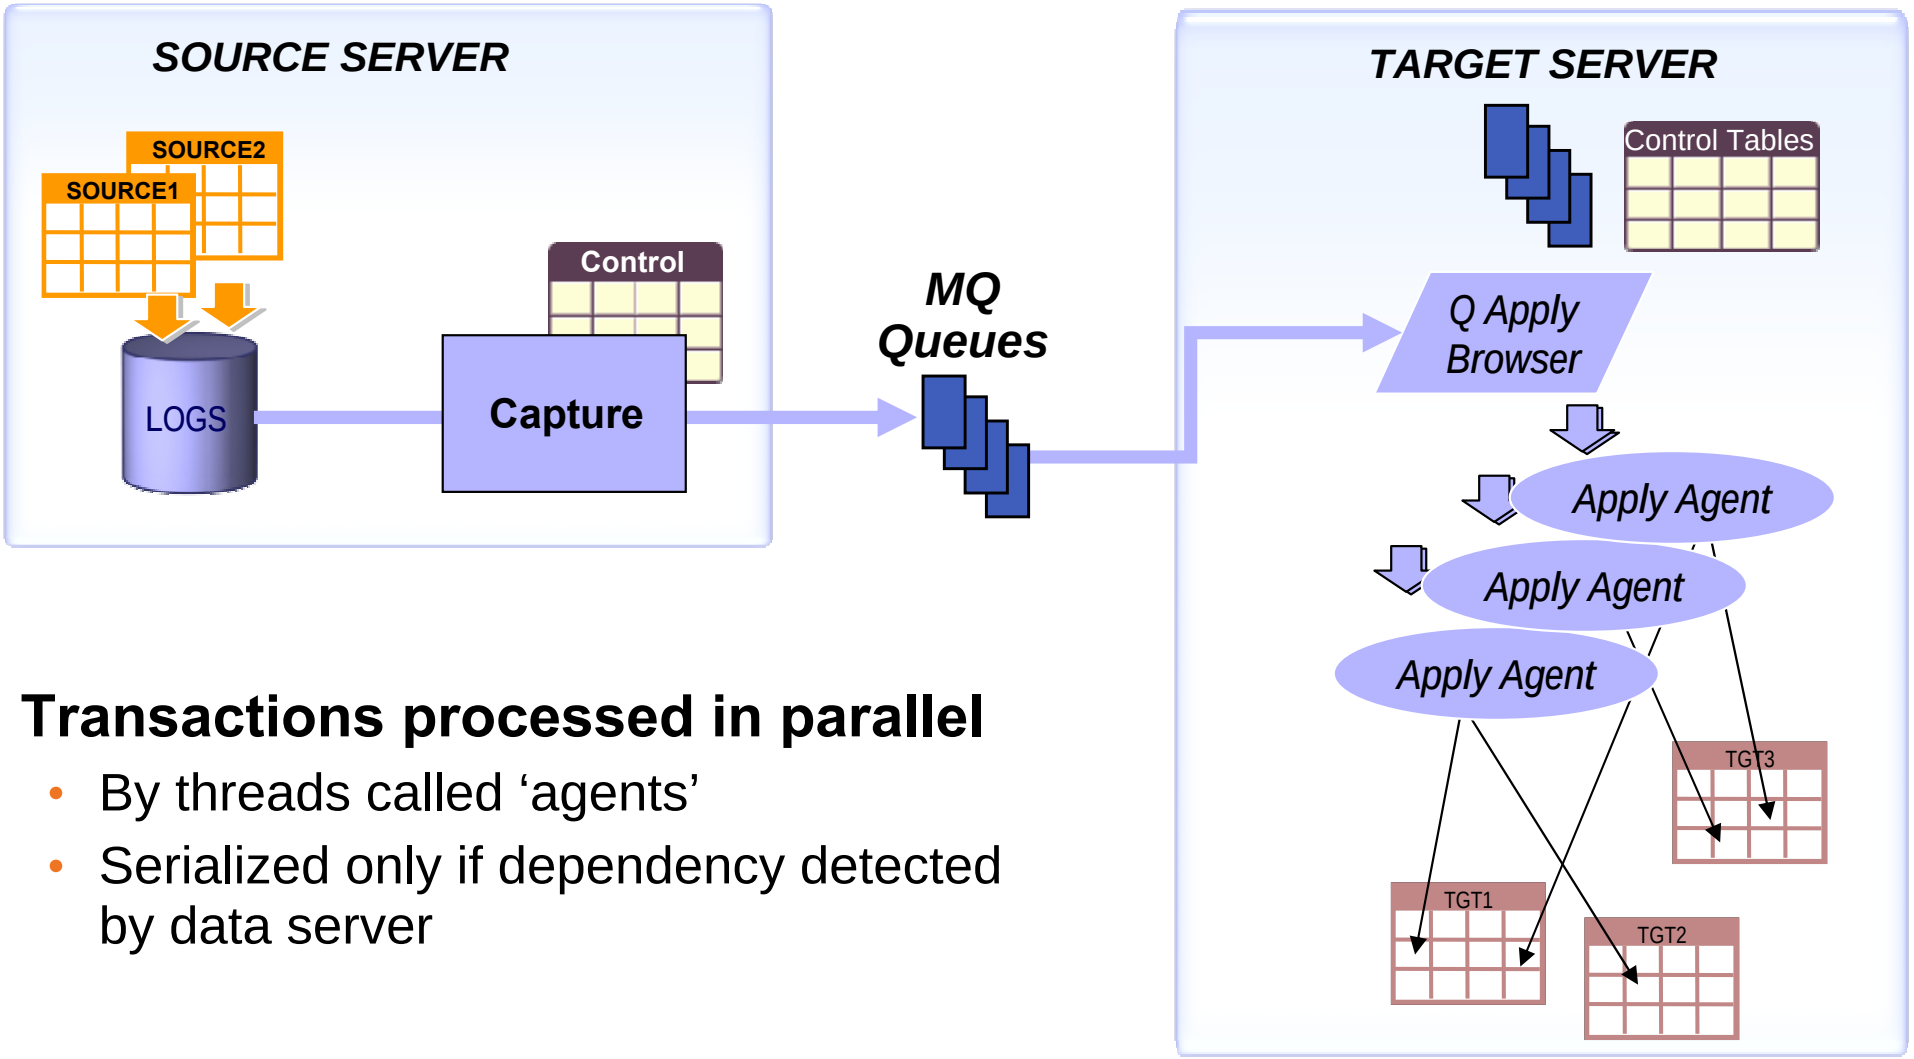

- **Transactions processed in parallel**
  - By threads called ‘agents’
  - Serialized only if dependency detected by data server

InfoSphere IMS Replication for z/OS

Unidirectional Replication of IMS data

- Release 1:
  - Conflicts will be detected
  - Manual resolution will be required
  - External initial load of target DB
  - Basic replication monitoring

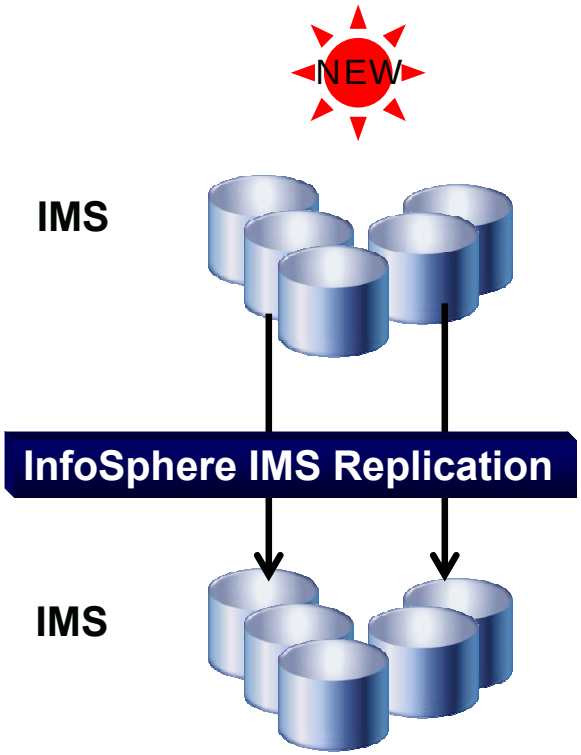

- **Administration via built-in GUI & z/OS console commands**
- **IMS "Capture" supports**
  - DB/TM, DBCTL, Batch DL/I
  - Capture x'99' log records
  - Increase in log volume due to change data capture records
- **IMS "Apply" supports**
  - Serialization based on resources updated by unit of recovery
  - Parallel apply
  - Requires New IMS Replication Restart Database

Unidirectional IMS Data Replication

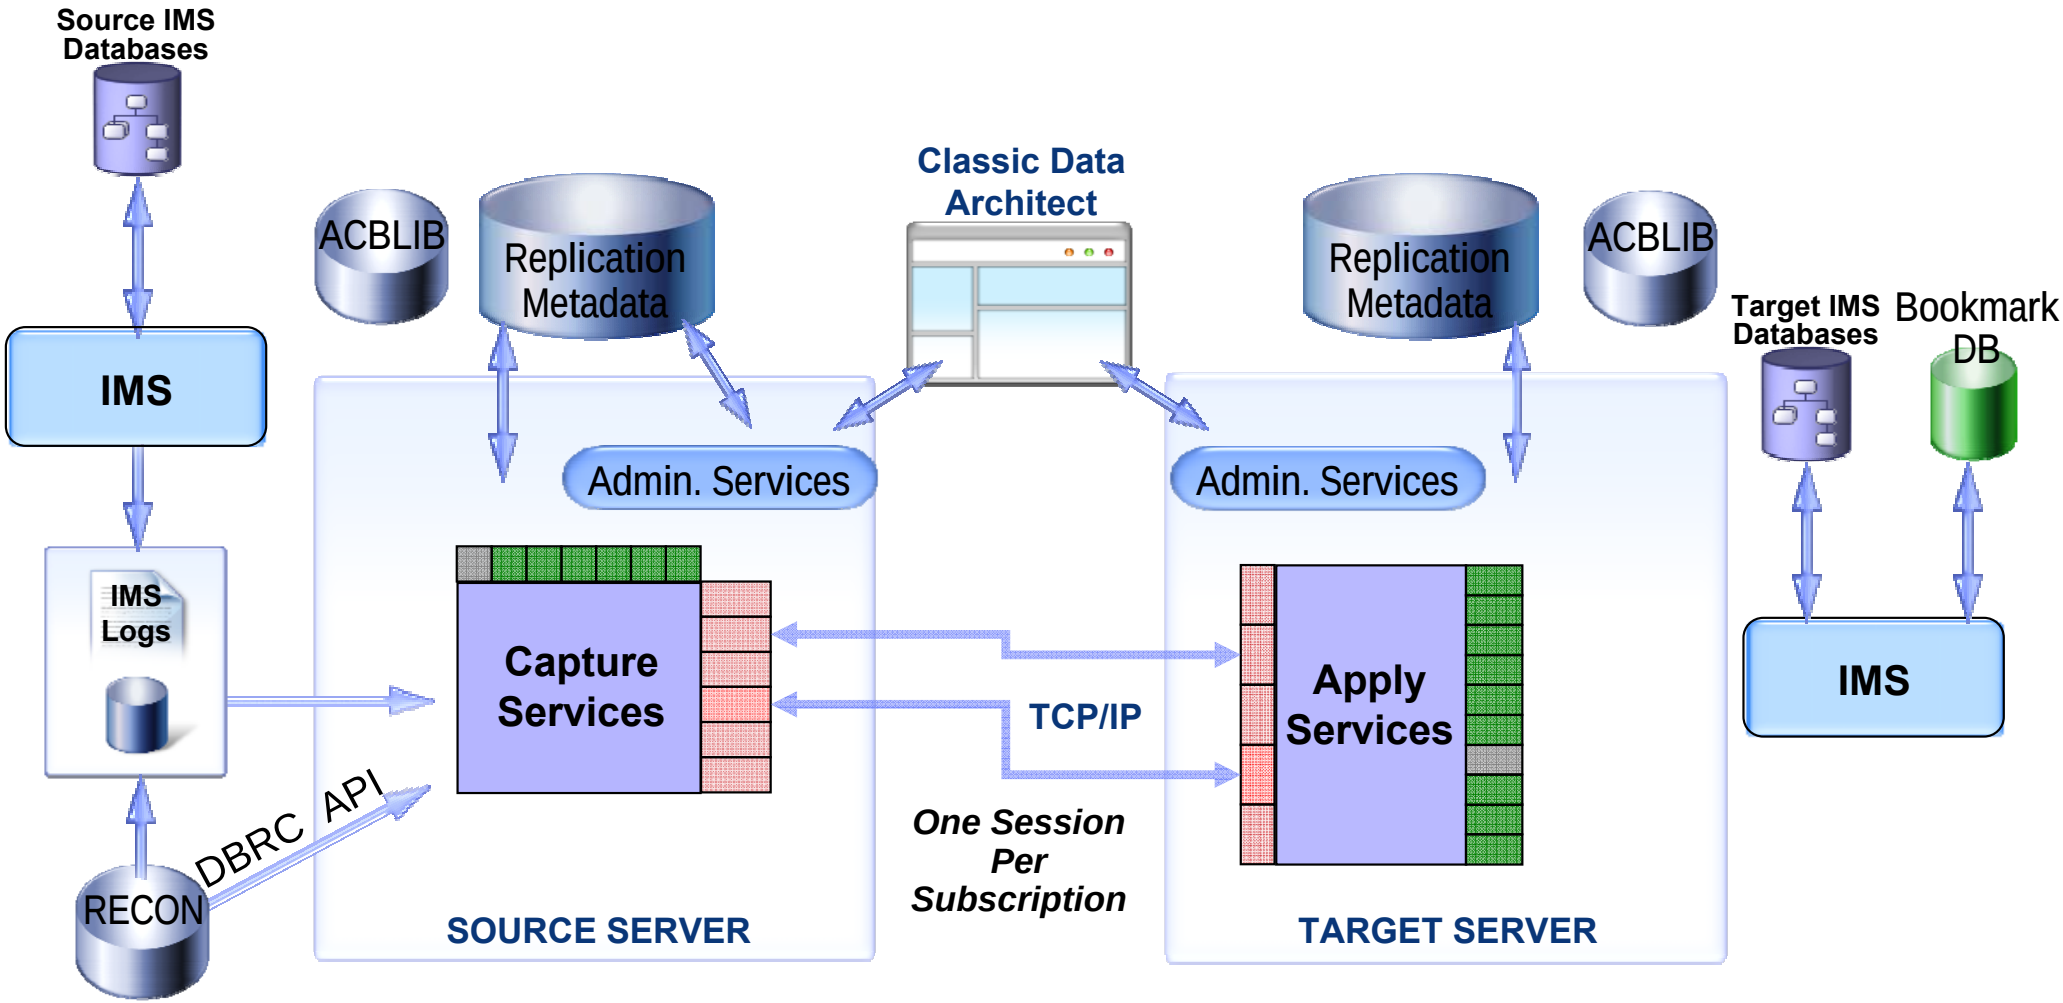

Some Details of IMS Data Replication

Capture Services – Log Merge

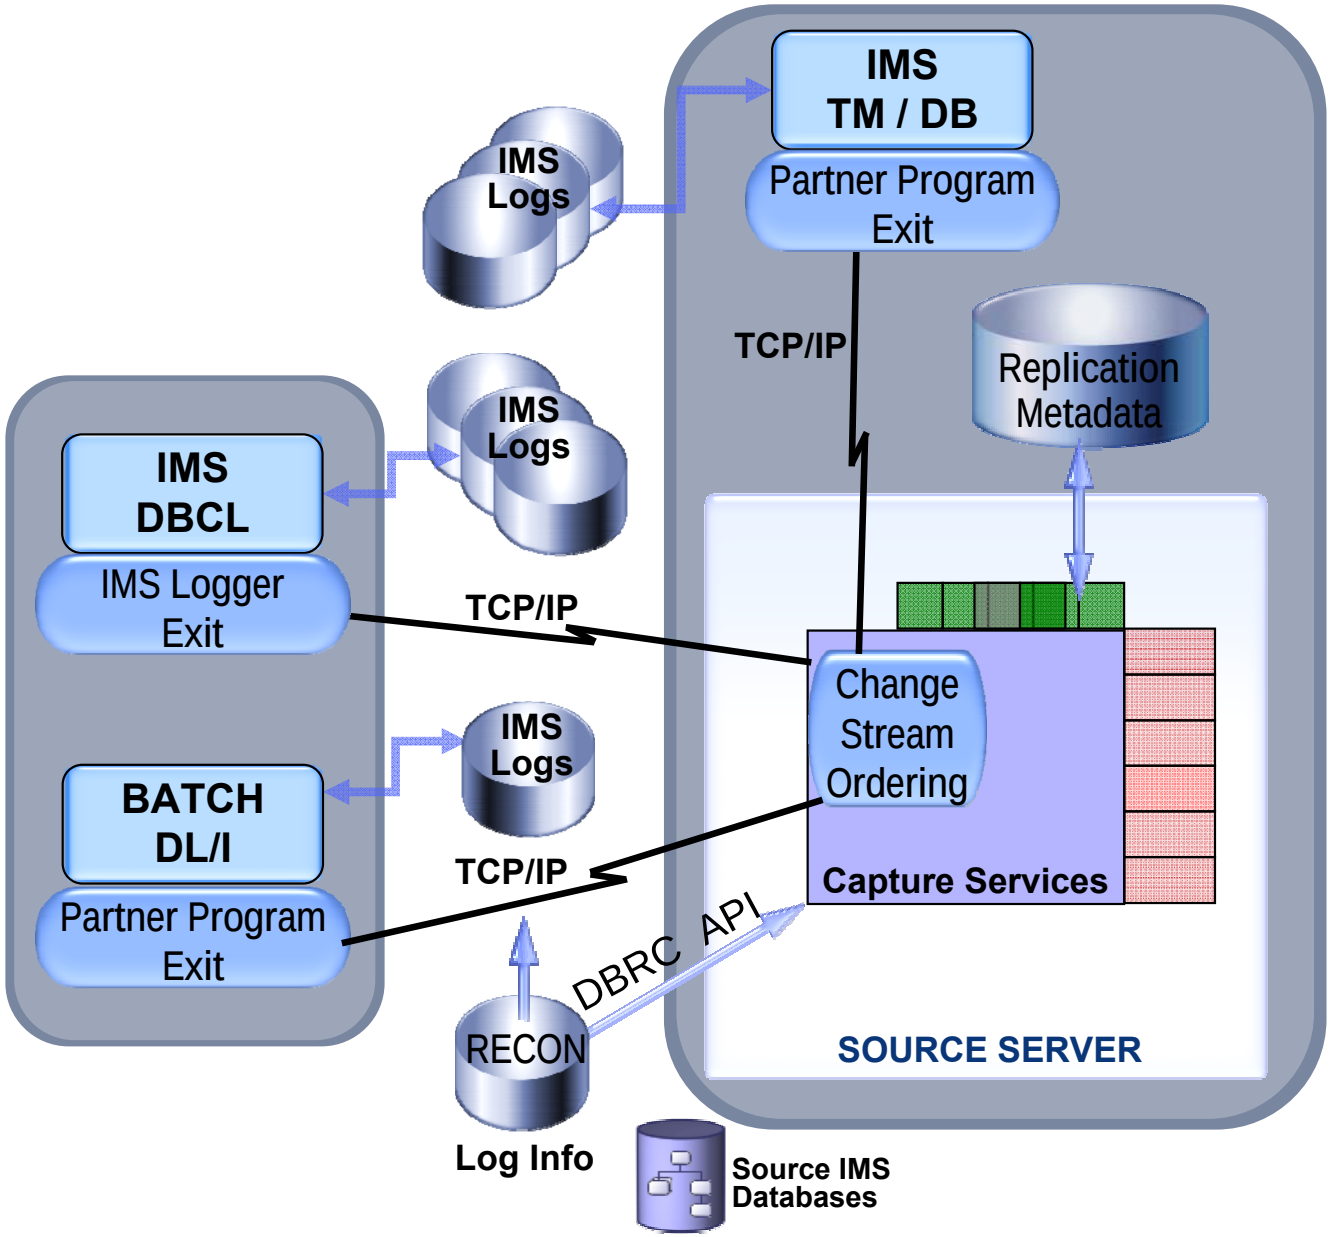

Some Details of IMS Data Replication

Target Services – Parallel Apply

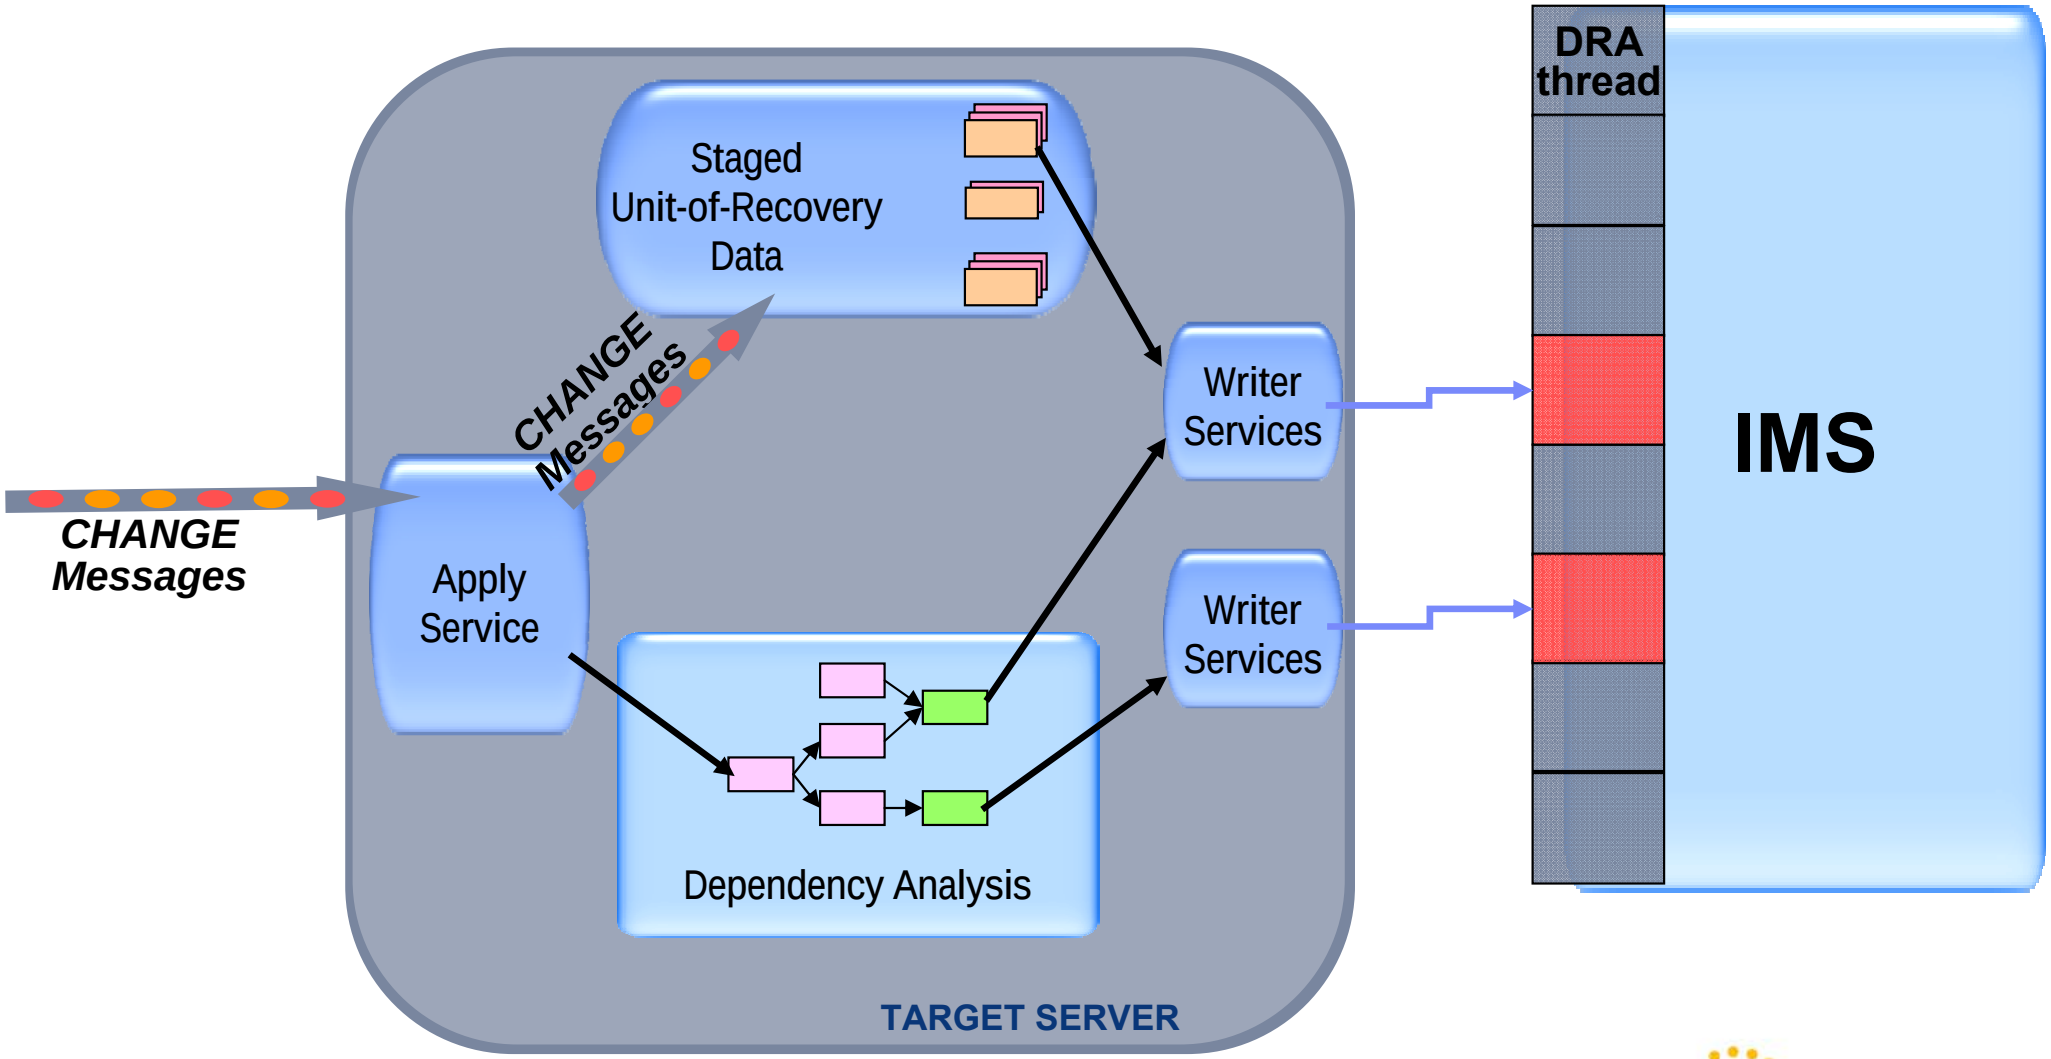

Business Scenarios for Software-Based Data Replication

Customer Scenarios for Replication

- **An automobile company uses a DB2 database to drive the factory floor production. Running reports against that database slows down the manufacturing process. A replicated copy increases manufacturing efficiency while allowing for up to date reports.**
  - *Same to same, low latency*
- **A financial company seeks a database infrastructure that will provide for high availability copies of their database but at the same time provide a real time feed to their information warehouse.**
  - *High availability in addition to ETL*
- **An insurance company distributes data from their central database at headquarters to all branches. At many of these branches the data is further distributed to individual insurance salesmen.**
  - *Many target copies, highly distributed*

CitiStreet

Selective "High Availability"

citistreet
A State Street and Citigroup Company

Challenge

- Support single sign-on access through both Web and IVR applications ensuring 24x7 portal access for plan participants and sponsors

Solution

- Support redundant, active single sign-on applications for failover processing replicating profile changes between them in real time.

"Since nearly 10 million of CitiStreet customers are offered 24-hour access to their retirement accounts, the company can't afford downtime and must be able to replicate data changes when they happen. We fully replicate our database over redundancy data lines, so to us the stability and speed of that asynchronous replication is strategic for us." Barry Strasnick , CIO CitiStreet

Overview

- CitiStreet is one of the largest and most experienced global benefits providers servicing over 9 million plan participants across all markets. CitiStreet was formed in partnership between subsidiaries of State Street Corporation and Citigroup

Business benefits

- Ensure application availability for plan participants and sponsors
- The new solutions from IBM will improve data integrity with a reduced level of maintenance

Technology benefits

- Maintain bi-directional synchronization of profile updates in real time (approx 175,000 updates daily)

SHARE in Orlando 2011

International Financial & Investment Services

Roll Your Own Continuous Availability

Challenge

- **Corporate initiative to provide customers better performing real-time queries by utilizing multiple sites.**
- **Replication of critical order processing details for core business functionality**

Solution

- **Q Replication for high speed movement of up to *10 Million transactions* to secondary site several thousand miles away.**

Business benefits

- Replicating 5-10 Million transactions with less than 2 seconds latency.
- More efficient and cost-effective resource utilization
- Secondary platform services reporting and business intelligence queries and acts as backup to primary

Technology benefits

- Real-time back up of secondary system provides results in increased capacity for peak workloads.

Today's Automated High Availability Solutions
GDPS PPRC/XRC/GM

Business Continuity Evolution with GDPS

GDPS/PPRC	GDPS/XRC **or** GDPS/GM	GDPS/Active/Active NEW
Failover Model	**Failover Model**	**Continuous availability model**
Recovery Time ≈ 2 min	**Recovery Time < 1 hour**	**Recovery time < 1 minute**
Distance < 20 miles	**Unlimited distance**	**Unlimited distance**

Continuous Availability w/ Disaster Recovery within a Metropolitan Region

GDPS/PPRC

RPO = 0 / RTO <1 hr

Two Data Centers

Systems remain active

Multi-site workloads can withstand site and/or storage failures

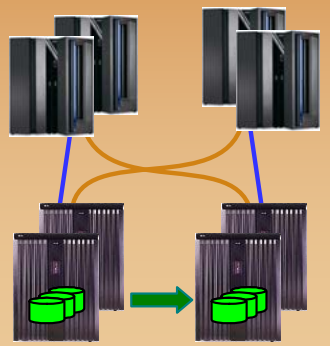

Disaster Recovery at Extended Distance

GDPS/GM & GDPS/XRC

RPO secs / RTO <1 hr

Two Data Centers

Rapid Systems Disaster Recovery with "seconds" of Data Loss

Disaster recovery for out of region interruptions

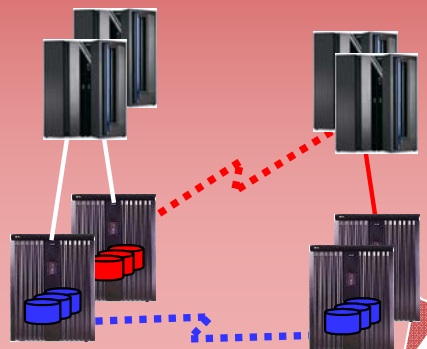

Continuous Availability, Disaster Recovery & Cross-Site Workload Balancing at Extended Distance

GDPS Active-Active Sites

RPO secs / RTO <1 min

Two or More Data Centers

All Sites Active

Continuous Availability for planned and unplanned interruptions

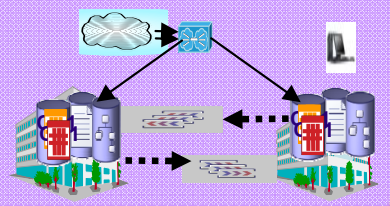

Regional Continuous Availability

GDPS/PPRC

- **Built on a multi-site Parallel Sysplex and synchronous disk replication**
- **Provides both:**
  - Metro-area Continuous Availability (CA),
  - Disaster Recovery solution (DR)
- **Supports two configurations:**
  - Active/standby
  - Active/active
- **Active/active customer configurations:**
  - All critical data must be PPRCed and HyperSwap enabled
  - All critical CF structures must be duplexed
  - Applications must be parallel sysplex enabled
  - Signal latency will impact OLTP thru-put and batch duration resulting in the sites being separated by no more than a ~20-30 of KM of fiber network

Issue: Insufficient site separation for some workloads

Disaster Recovery at Extended Distances

GDPS/XRC and GDPS/GM

- **Asynchronous disk replication**
- **Unlimited distance Disaster Recovery solutions**
- **Require the failed site's workload to be restarted in the recovery site and this typically will take 30-60 min**
  - Power fail consistency
  - Transaction consistency

Issue: Can NOT achieve RTO of seconds needed for some workloads

Customer Requirements

RTO near zero, Replace roll-your-own, Leverage all resources

Shift focus from failover to nearly-continuous availability

- **"Recover my business rather than my platform technology"**
  - Multi-sysplex, multi-platform solution
  - No application changes
  - Access data from any site with unlimited distance between sites
  - Provide application level granularity rather than the current "all-or-nothing" model
    - Some workloads may require immediate access from every site
    - Some workloads may only need to update other sites every 24 hours

- **Minimize costs and Optimize resource utilization**
  - Automated recovery processes (similar to GDPS technology today), minimizing operator learning curve
  - Provide workload distribution between sites
    - Dynamically select sites based on their ability to handle workload
    - Route around failed sites

GDPS Active/Active Sites Configurations

- **Configurations**
  - **Active/Standby – Announced June, 2011**
  - Active/Query – Stated Direction
  - Active/Active – Customer Defined Goal
- **A configuration is specified on a workload basis**
  - Mixed configurations can be used to handle the diverse recovery requirements
- **A workload is the aggregation of these components**
  - Software – user written applications (e.g., COBOL program) and the middleware run time environment (e.g., CICS region & DB2 subsystem)
  - Data - related set of objects that must preserve transactional consistency and optionally referential integrity constraints (e.g., DB2 Tables)
  - Network connectivity – one or more TCP/IP addresses & ports (e.g., 10.10.10.1:80)

Active/Active concepts

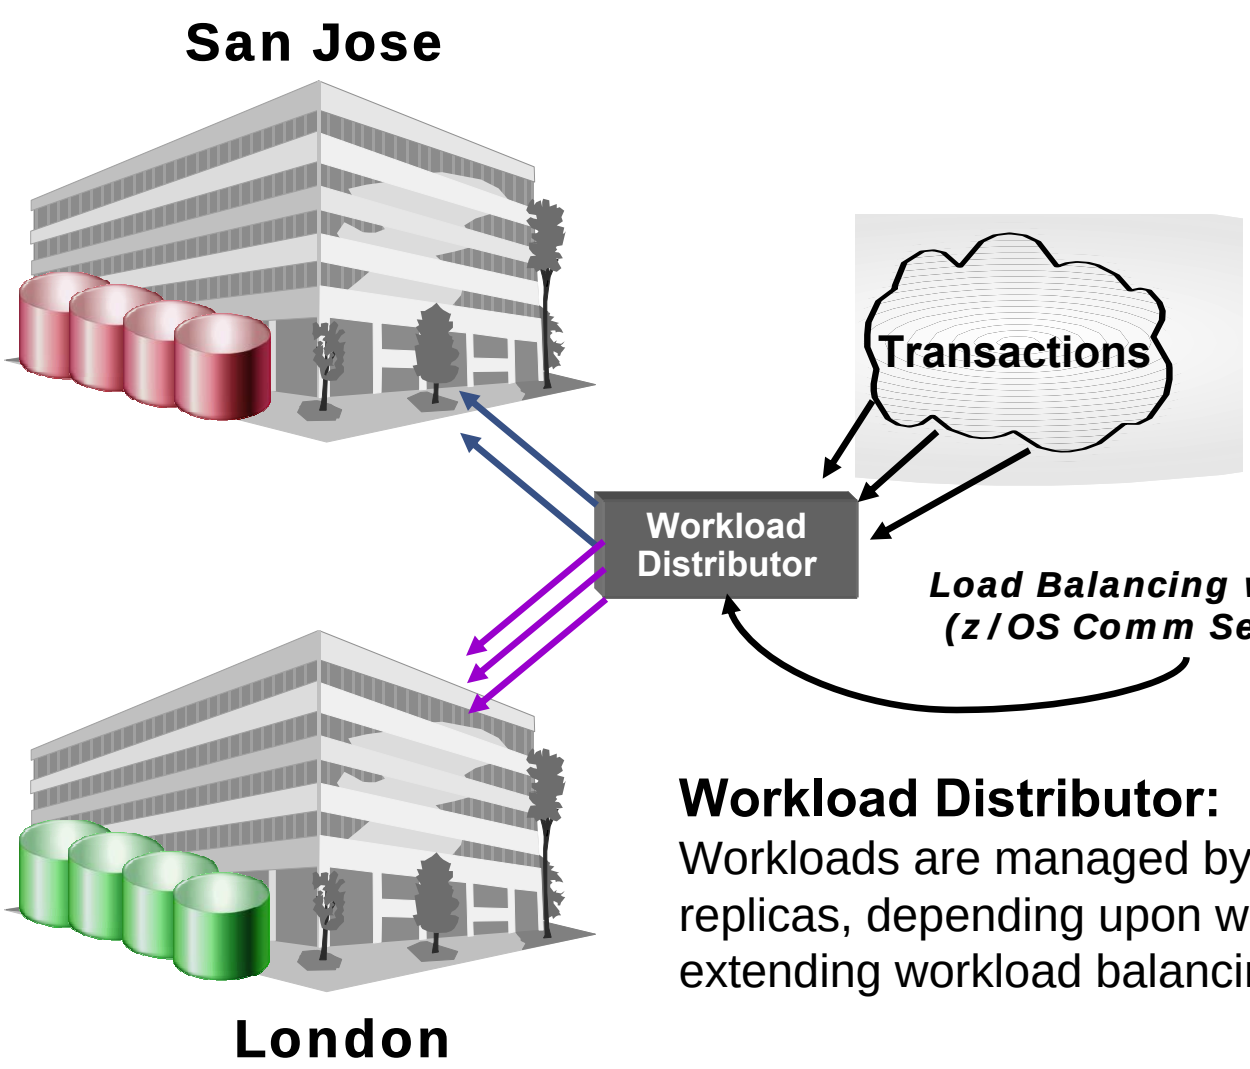

Two or more sites, separated by ***unlimited*** distances, running the same applications & having the same data to provide:

- Cross-site Workload Balancing
- ***Continuous*** Availability
- Disaster Recovery

Workload Distributor:

Workloads are managed by a client and routed to one of many replicas, depending upon workload weight and latency constraints, extending workload balancing to SYSPLEXs across multiple sites!

Active/Active concepts

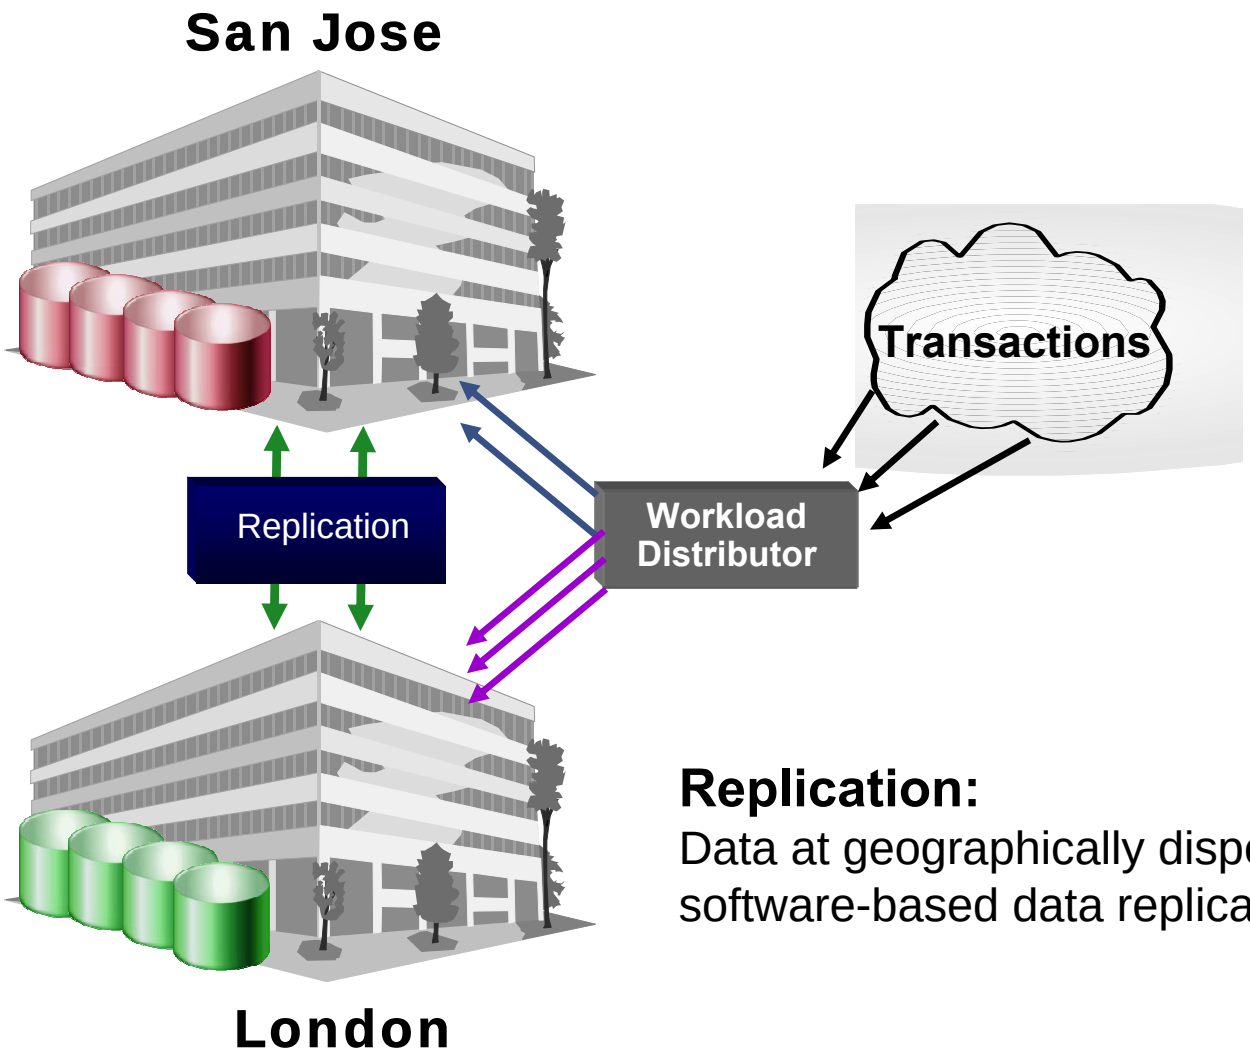

Two or more sites, separated by ***unlimited*** distances, running the same applications & having the same data to provide:

- Cross-site Workload Balancing
- ***Continuous*** Availability
- Disaster Recovery

Replication:
Data at geographically dispersed sites are kept in sync via software-based data replication

Active/Active concepts

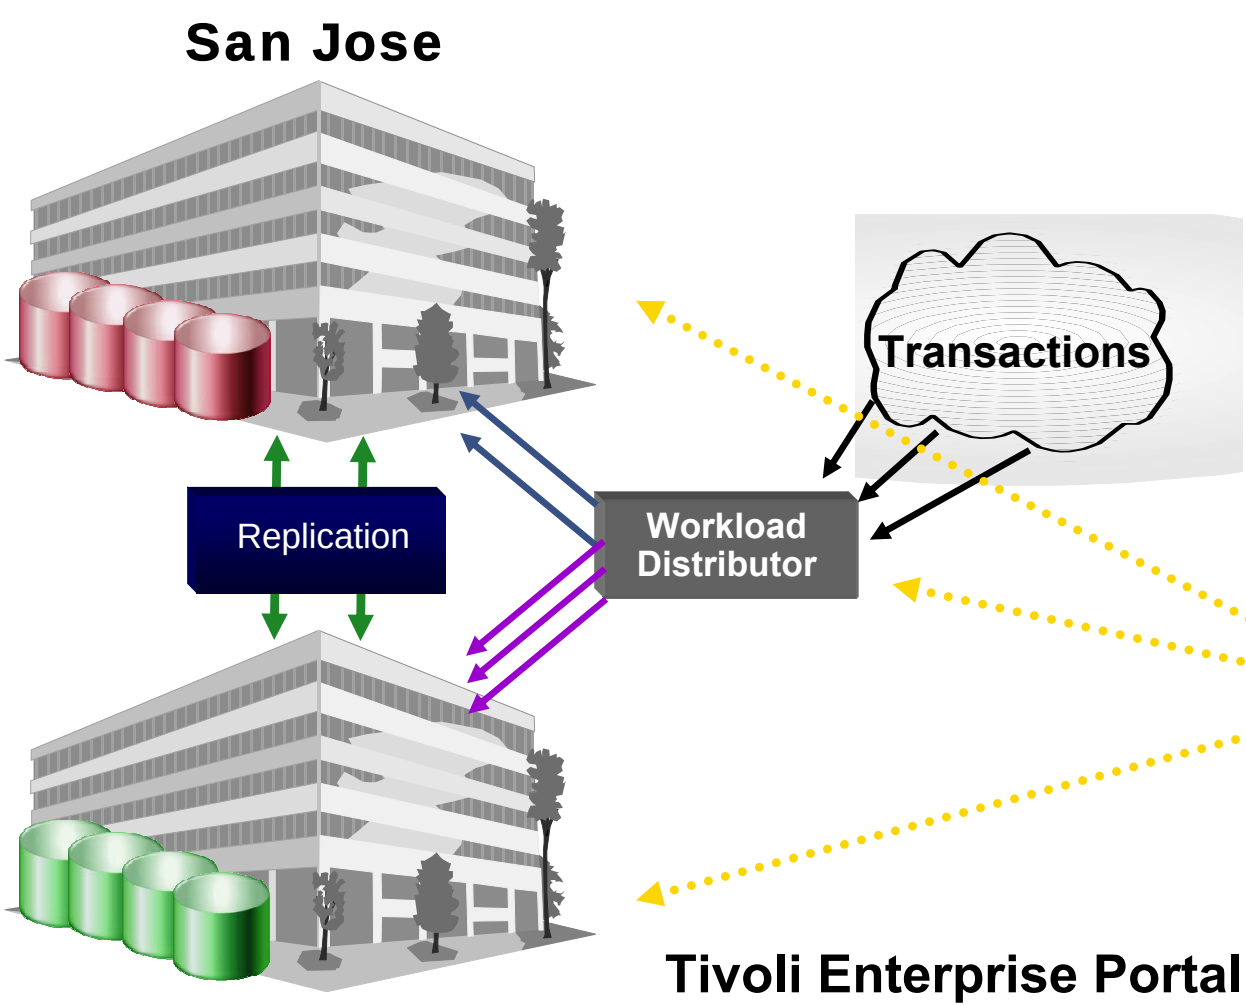

Two or more sites, separated by ***unlimited*** distances, running the same applications & having the same data to provide:

- Cross-site Workload Balancing
- ***Continuous*** Availability
- Disaster Recovery

Tivoli Enterprise Portal

Tivoli Enterprise Portal:
Monitoring spans the sites and now becomes an essential element of the solution for site health checks, performance tuning, etc.

Conceptual view

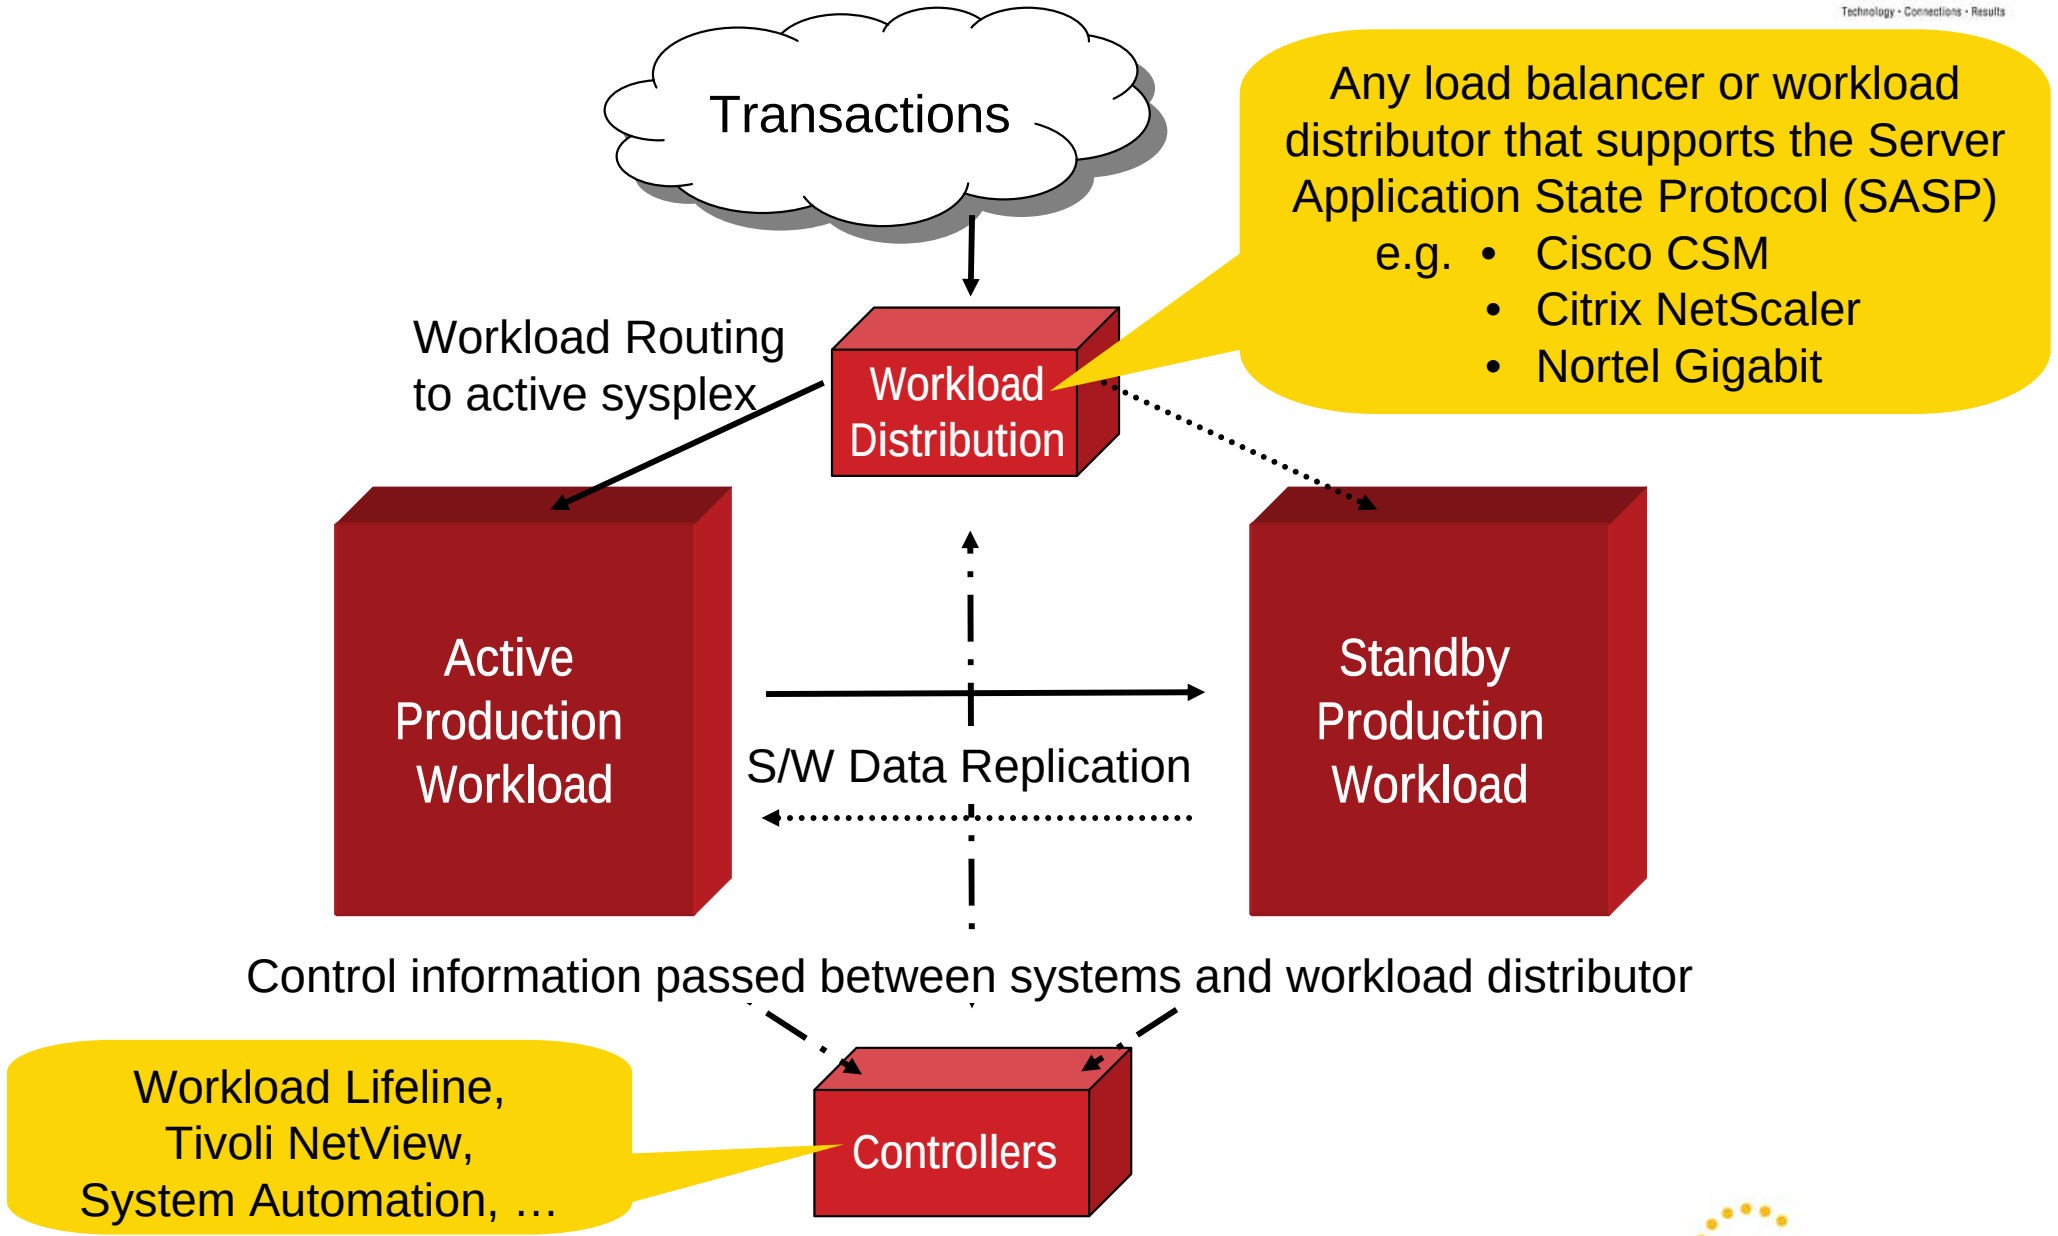

Active/Active Summary workspace

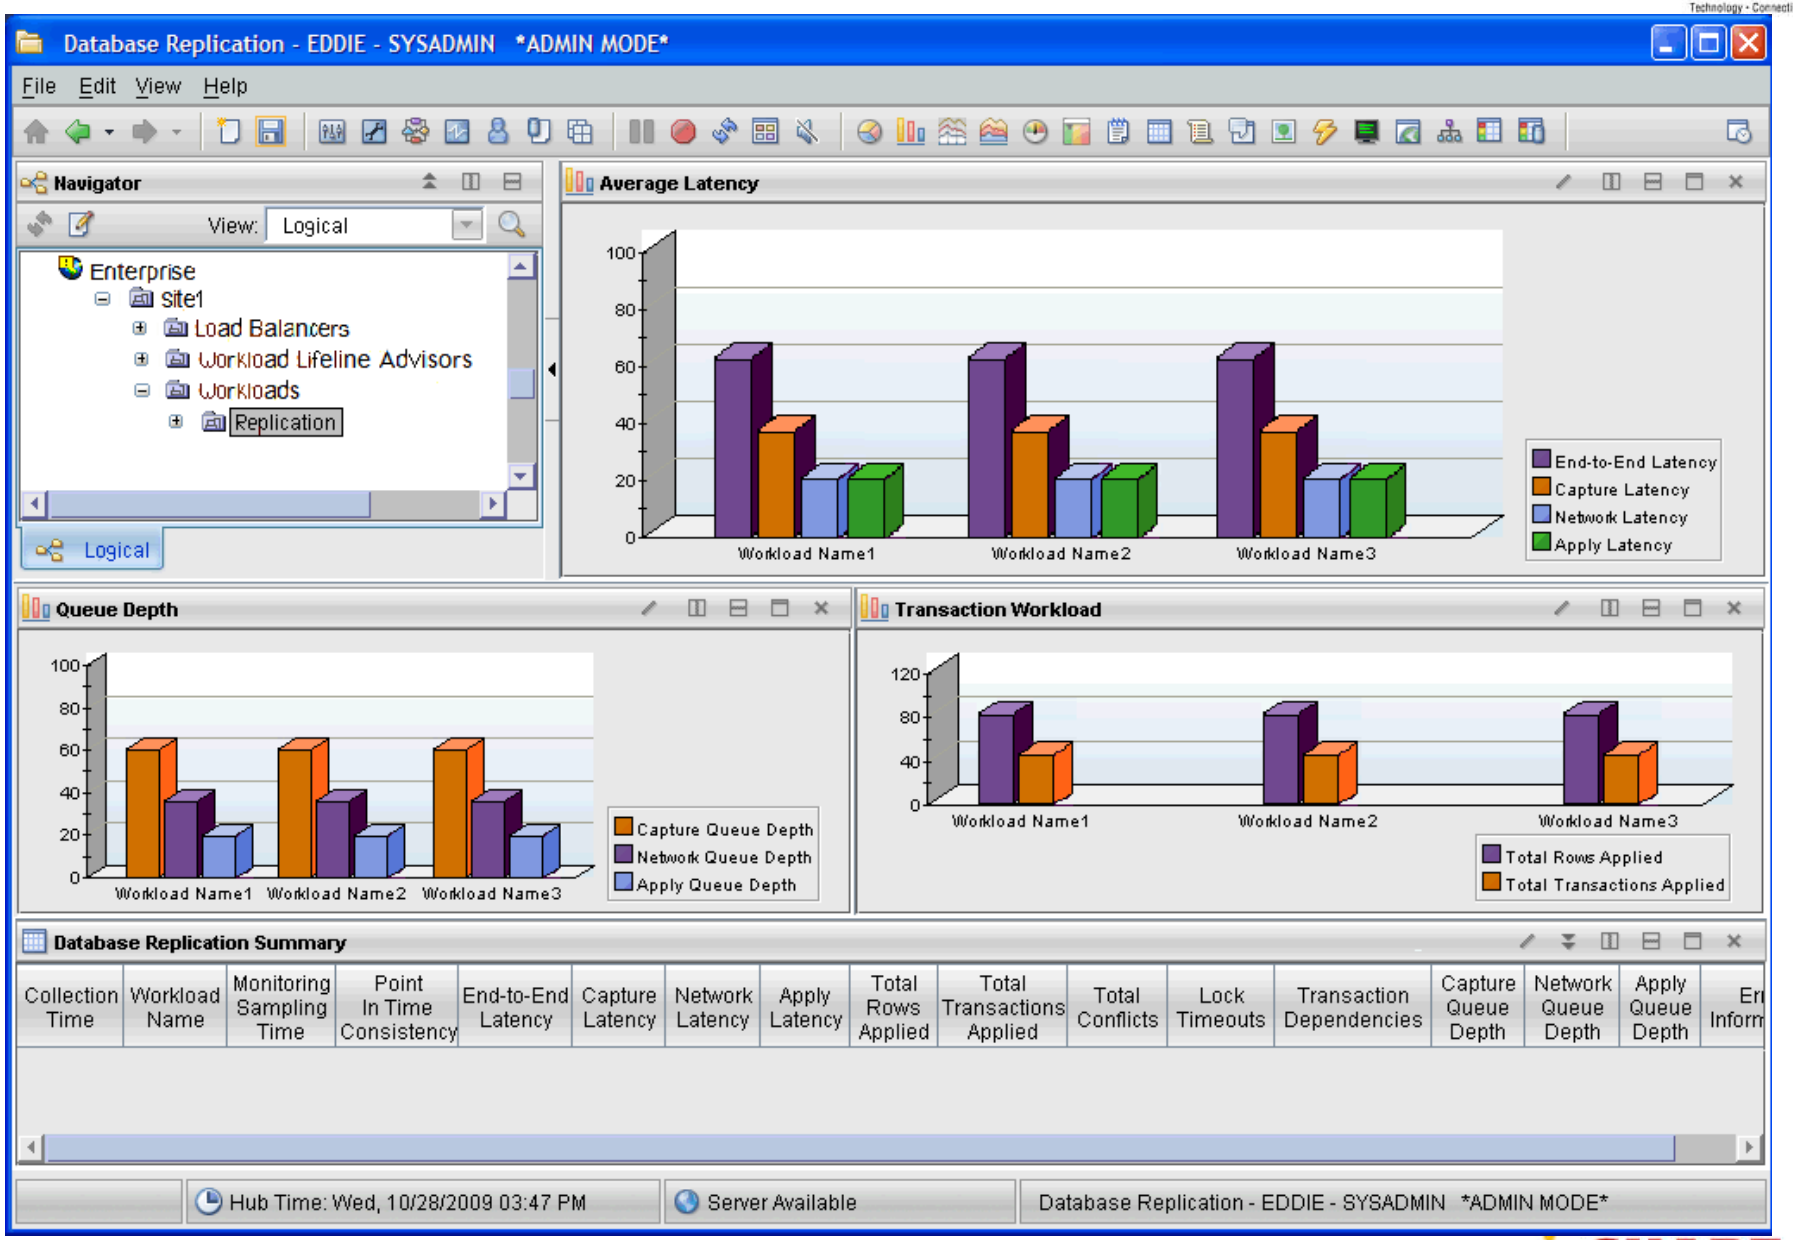


SHARE
Technology · Connections · Results
Thank You
SHARE
in Orlando
2011