Enterprise Java Batch for z/OS with WebSphere XD Compute Grid

Snehal S. Antani, antani@us.ibm.com
WebSphere XD Technical Lead
SOA Technology Practice
IBM Software Services

ON DEMAND BUSINESS™

Presentation Outline (Sections)

- **Batch Computing**
  - **What is it? What are batch jobs?**
- Challenges Related to Batch Computing
  - How does WebSphere XD address?
- Batch and SOA
- What is WebSphere XD?
- Positioning XD Compute Grid in the world of Batch
- XD Compute Grid Use-cases
  - Batch Modernization
  - Highly parallel batch jobs
  - Dynamic OLTP and Batch infrastructure
  - Batch as a service
  - Replacing existing java batch frameworks
  - Sharing business logic across OLTP and Batch

What are “Batch Jobs”?

- Process millions of business critical transactions “behind the scenes”
  - Without user interaction, asynchronous…
- **IBM has a strong heritage and significant capabilities in batch!**

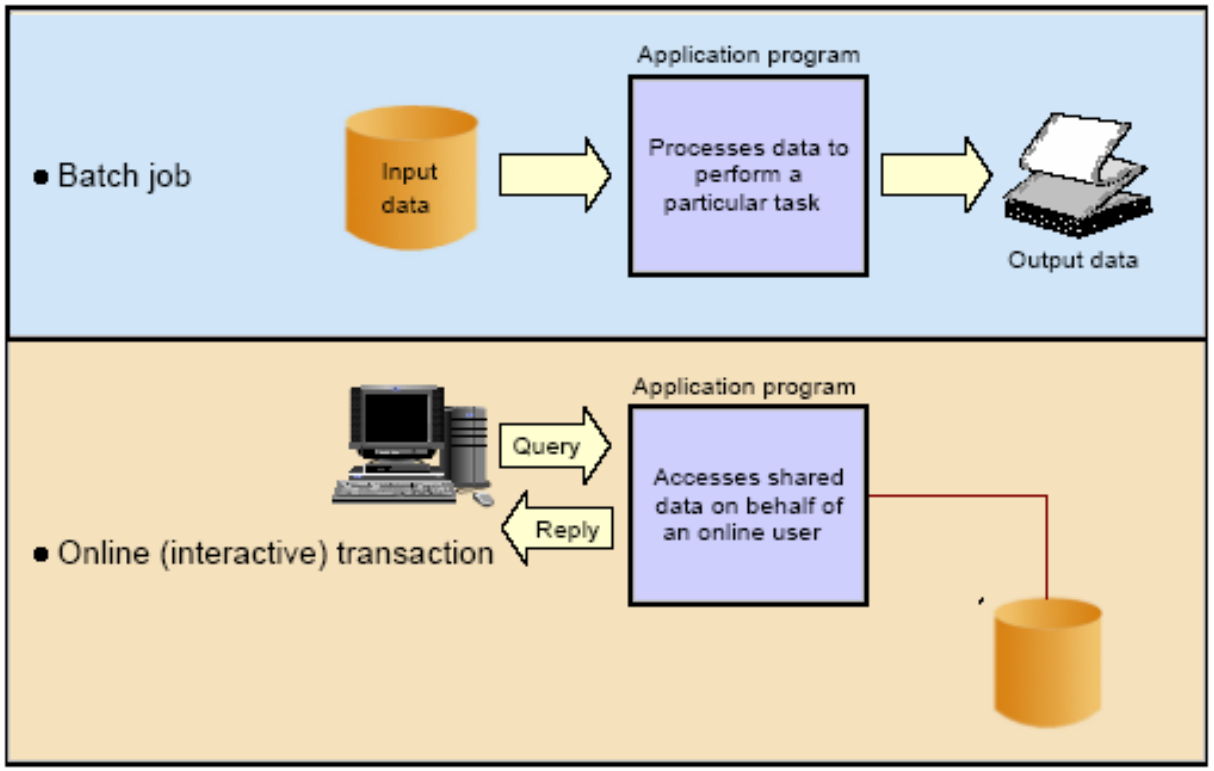

Examples

- Payment Processing
- Shipment Processing
- Report Generation
- Claims Processing
- Inventory Management
- End of Day/Month/Quarter/year processing

Presentation Outline (Sections)

- Batch Computing
  - What is it? What are batch jobs?
- **Challenges Related to Batch Computing**
  - **How does WebSphere XD address?**
- Batch and SOA
- What is WebSphere XD?
- Positioning XD Compute Grid in the world of Batch
- XD Compute Grid Use-cases
  - Batch Modernization
  - Highly parallel batch jobs
  - Dynamic OLTP and Batch infrastructure
  - Batch as a service
  - Replacing existing java batch frameworks
  - Sharing business logic across OLTP and Batch

IT Departments are challenged to…

- Absorb tremendous growth rates while executing with constant IT budgets
- Reduce Execution Costs
- Reduce Maintenance Costs
- Meet Aggressive Performance Demands

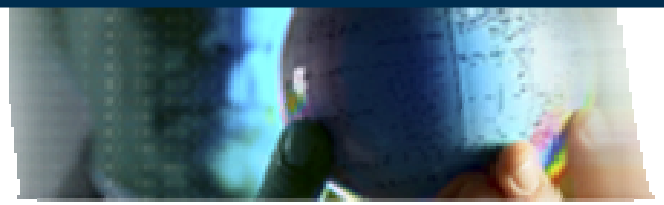

Challenge: Reduce Execution Costs

- Traditional batch processing does not leverage zAAP processors on z/OS
- Customers mitigate this by
  - Building homegrown Java batch solutions ***(code maintenance costs)***
  - Shifting workload to distributed platforms (***performance issues, risky all-or-nothing approach***)

Solution: Lower Costs with WebSphere XD Compute Grid

- Delivers a batch runtime and platform (Infrastructure Services for an SOA)
- Builds on the Qualities of Service delivered by WAS z/OS (high availability, scalability, etc)
- zAAP-eligible

Challenge: Reduce Maintenance Costs

- Disparate business logic (separate business logic for OLTP and Batch)
- Diminishing COBOL skills in the market place
- Batch applications are not portable across platforms
- Customers mitigate this with hybrid execution models
  - Collocating COBOL (native) programming languages with Java ***(only exists in basic form today)***
  - Bridging batch to WebSphere z/OS via WebSphere MQ ***(high performance and MIPS costs)***

➢ Solution: Lower Costs with WebSphere XD Compute Grid

- Seamlessly share business across OLTP and Batch
- Batch applications are written in Java (abundance of skills and tooling)
- Batch applications are portable across platforms
- Advanced hybrid execution models and System-Z integration
  - Collocating COBOL (native) programming languages with Java
  - Can integrate Java batch applications with traditional z/OS services (DFSort, etc)

Challenge: Meet Aggressive Performance Demands

- Narrower batch windows due to 24x7 datacenter operations
- Ability to do continuous batch processing
- Utilize hardware resources efficiently
- Close proximity of batch applications and data **is essential for performance**

➢ **Solution: Exceed performance demands with WebSphere XD Compute Grid**

- Highly-parallel and stream-lined batch jobs for performance
- Uses simple java objects (POJO), not heavy-weight services
- Facilitates close proximity of batch applications and data
- Benefits from Java 5 performance on z/OS

The Effect of Co-locating Data with Batch Applications

If your data is on z/OS, keeping your batch applications on z/OS makes the most sense…

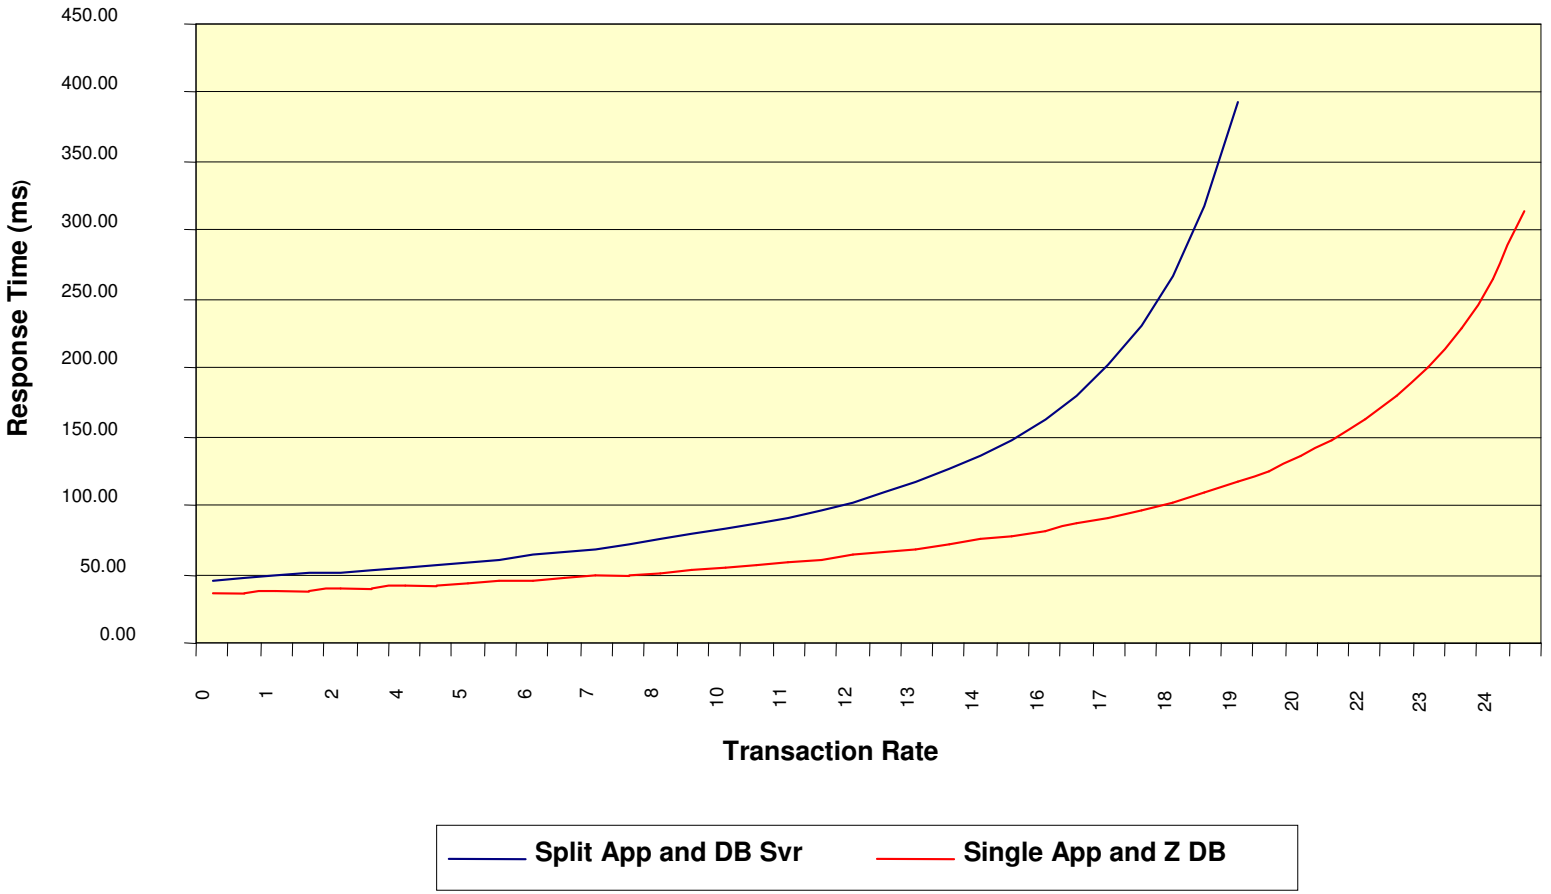

The Effect of Co-locating Data with Batch Applications

Split Application Server and Database Server

- Physical separation of the application layers is achieved by using network connections (DRDA over TCP/IP)
- Adds latency to the overall transaction response time as a consequence of queuing, protocol processing, and network transition
- At an individual transaction level the impact is slight
- But queuing delays increase exponentially as the number of database requests and CPU resources consumed increase
- Latency has implications on the overall design and implementation
  - Overall transaction throughput of the system is reduced
  - Increased memory required to hold the in-flight transactions
    - Potential bottleneck and extra cost

Split Batch and Database Server

- Batch will continue to be a major element of processing
- Batch continues to be efficient means of processing a wide range of tasks
- Common need to run both Batch and OLTP simultaneously
- Database server must be capable of dynamically supporting mixed workloads running simultaneously
- Queuing effects of physical separation of Batch from Database Server has more pronounced effect
- Application design to exploit DRDA block fetch & SQLJ/JDBC batching would help offset, some not all, of penalty
- Elapsed time could easily increase by 50%

Presentation Outline (Sections)

- Batch Computing
  - What is it? What are batch jobs?
- Challenges Related to Batch Computing
  - How does WebSphere XD address?
- **Batch and SOA**
- What is WebSphere XD?
- Positioning XD Compute Grid in the world of Batch
- XD Compute Grid Use-cases
  - Batch Modernization
  - Highly parallel batch jobs
  - Dynamic OLTP and Batch infrastructure
  - Batch as a service
  - Replacing existing java batch frameworks
  - Sharing business logic across OLTP and Batch

Batch and SOA

"… business function used in online transactions may be the same business function used in batch processes, so organizations should think about their IT modernization strategy and consider SOA as a standardized application integration mechanism" - Gartner Research

Reusing business services is a fundamental principle of SOA

Batch workloads are an integral part of any IT infrastructure

How do you integrate your batch & OLTP environments with a common services infrastructure?

Presentation Outline (Sections)

- Batch Computing
  - What is it? What are batch jobs?
- Challenges Related to Batch Computing
  - How does WebSphere XD address?
- Batch and SOA
- **What is WebSphere XD?**
- Positioning XD Compute Grid in the world of Batch
- XD Compute Grid Use-cases
  - Batch Modernization
  - Highly parallel batch jobs
  - Dynamic OLTP and Batch infrastructure
  - Batch as a service
  - Replacing existing java batch frameworks
  - Sharing business logic across OLTP and Batch

WebSphere Extended Deployment (XD)

XD contains 3 components on z/OS available as a single, integrated package or 3 individual components

Compute Grid	Operations Optimization	Data Grid
- *Transactional Batch* - *Compute Intensive Tasks* - *Manage non-Java workloads* - *z/OS Integration Patterns*	- *On-Demand Router* - *Extended Manageability* - *Application Editions* - *Health Management* - *Runtime Visualization* - *Virtualization*	- *Distributed Caching* - *Partitioning Facility* - *In-memory Databases*

XD Design Principles

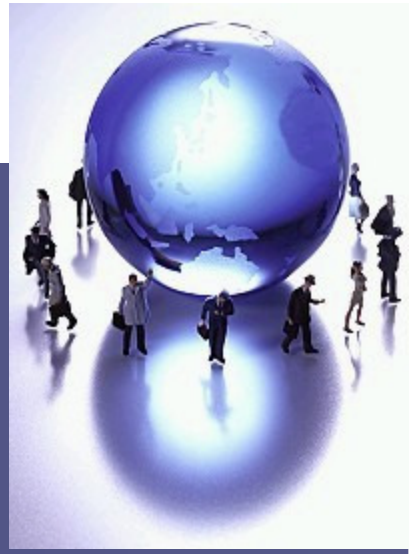

An platform "**extender**" …. pre-reqs WebSphere Application Server

Installation is a simple delta to an existing environment

- ✓ Does not require "migration" or restructuring of current installation

Completely integrated into the WebSphere Application Server environment

- ✓ Extends the WAS Admin Console
- ✓ Extends the wsadmin scripting environment

Meaningful without implementing all new concepts and features

- ✓ Manual and Supervised modes allow autonomics to be adopted gradually
- ✓ Goals-directed WLM can be implemented without autonomic features
- ✓ Long running (J2EE batch) and partitioned programming models can be introduced incrementally
- ✓ Extended Manageability features such as application editions, repository checkpoints, visualization and health monitoring have broad appeal

WebSphere XD z/OS- Compute Grid summary

- The XD Batch Container supports J2EE-based Batch Applications.
- These applications allow for batch access to enterprise applications hosted in WebSphere. They have available to them WebSphere resources:
  - **Transactions**
  - **Security**
  - **high availability**
  - **Leverages the inherent WAS QoS**
- The XD Batch Container provides services such as
  - **Check Pointing**- the ability to resume batch work at some selected interval
  - **Batch Data Stream Management**- the ability to handle reading, positioning, and repositioning data streams to files, relational databases, native z/OS datasets, and many other input and output sources.
  - **Result Processing**- the ability to intercept step and job return codes and subsequently process them using any J2EE facility (Web Service, JMS message, and so on)

XD Compute Grid Topology

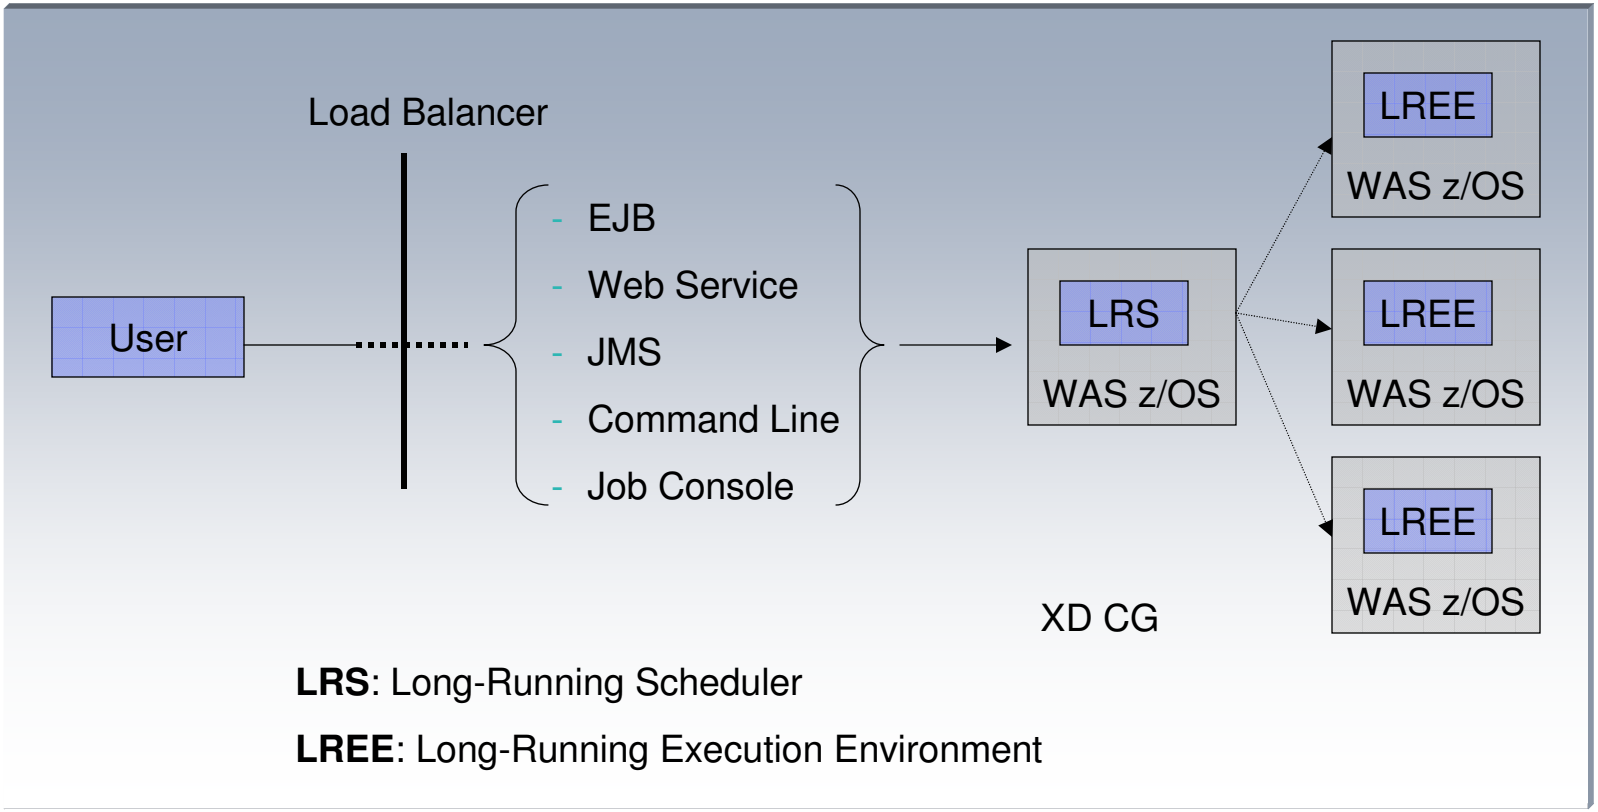

XD Compute Grid and z/OS Integration

XD Compute Grid leverages and integrates with native z/OS workload management (WLM) to enhance job execution & management

SMF accounting records for J2EE batch jobs

- SMF 120 (J2EE) records tailored to jobs
- Record includes: job id, user, CPU time

Dynamic Servants for J2EE batch job dispatch

- Exploits WLM to start new servants to execute J2EE batch jobs on demand

Service policy classification and delegation

- Leverage XD classification to select z/OS service class by propagating transaction class from scheduler to z/OS app server for job registration with WLM

Presentation Outline (Sections)

- Batch Computing
  - What is it? What are batch jobs?
- Challenges Related to Batch Computing
  - How does WebSphere XD address?
- Batch and SOA
- What is WebSphere XD?
- **Positioning XD Compute Grid in the world of Batch**
- XD Compute Grid Use-cases
  - Batch Modernization
  - Highly parallel batch jobs
  - Dynamic OLTP and Batch infrastructure
  - Batch as a service
  - Replacing existing java batch frameworks
  - Sharing business logic across OLTP and Batch

Positioning XD Compute Grid

Role of Enterprise Schedulers

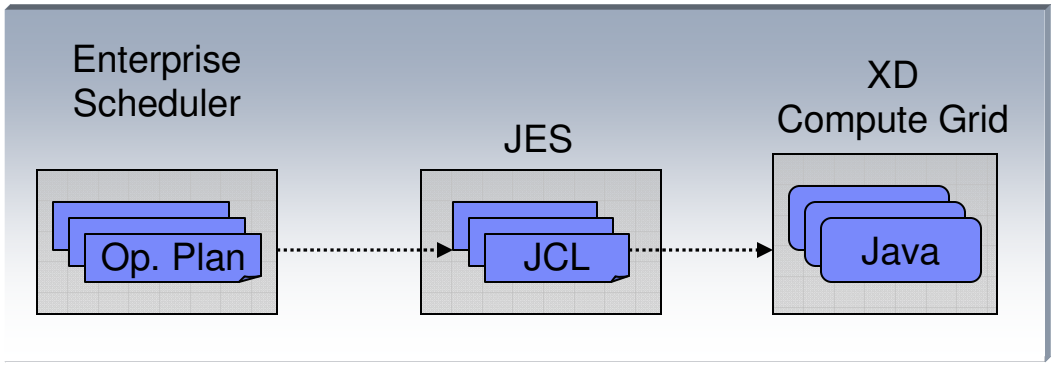

- Tivoli Workload Scheduler, CA7, Control-M, Zeke, and others are ***Enterprise Schedulers.*** These schedulers manage Operational plans, job dependencies, resource definitions, job recovery, and so on

- XD Compute Grid is a ***Batch Execution Environment***, it is **not** in the scheduling business. XD CG executes jobs with expected QoS.

- XD Compute Grid ***complements*** Enterprise Schedulers, integrates with them, works alongside them.

Positioning XD Compute Grid- Role of JZOS

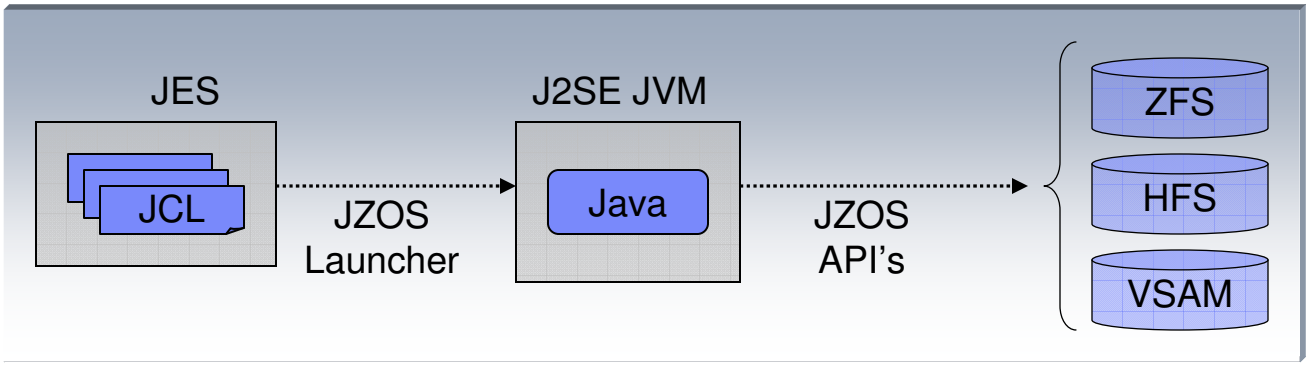

- JZos delivers 2 technologies:
  1. ***JZOS Launcher-*** *s*eamless way to initialize a ***J2SE*** runtime from JCL
  2. ***JZOS API's-*** set of library functions for accessing traditional z/OS resources (MVS datasets, VSAM files, etc) from Java

- ***JZOS launcher*** is excellent for simple java applications, not for 1000's of enterprise batch jobs to be run within a batch window
  - J2SE JVM has no:
    - security, transaction, or connection management
    - checkpoint or restart facility for batch jobs
    - inherent high availability, or other WAS z/OS qualities of service
  - JVM is not persistent or reusable.

Positioning XD Compute Grid- Role of JZOS

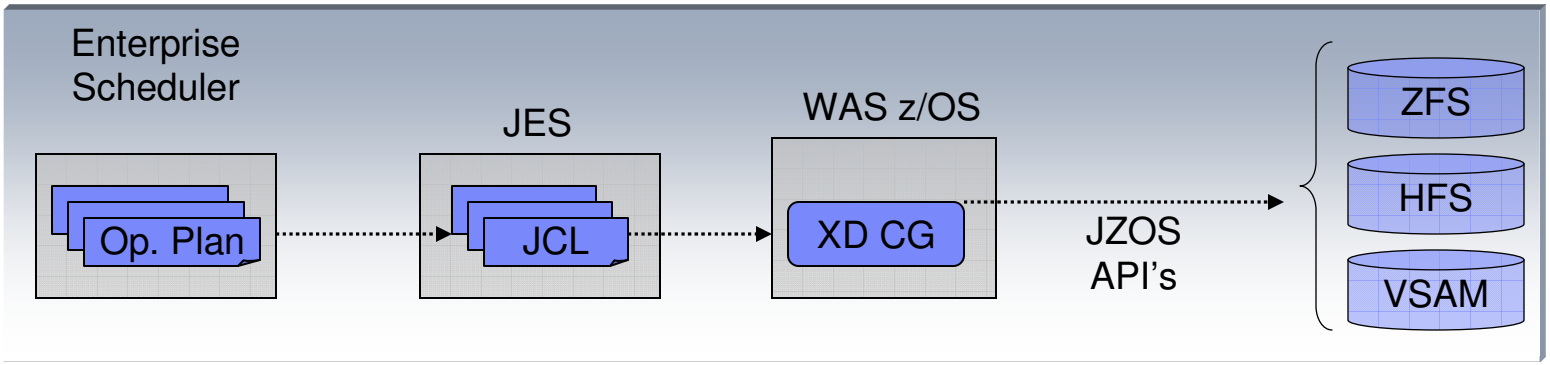

- XD Compute Grid is ***built on WebSphere z/OS***
  - leverages QoS and services provided by the WAS z/OS runtime *(security, transaction, connection management; thread pooling; HA, etc)*
  - Runs within a persistent, reusable JVM and Execution Container
- JZOS Api's can be leveraged from XD CG applications
- JZOS Api's provide ***a strong integration point*** for Java and traditional z/OS

Presentation Outline (Sections)

- Batch Computing
  - What is it? What are batch jobs?
- Challenges Related to Batch Computing
  - How does WebSphere XD address?
- Batch and SOA
- What is WebSphere XD?
- Positioning XD Compute Grid in the world of Batch
- XD Compute Grid Use-cases
  - **Batch Modernization**
  - Highly parallel batch jobs
  - Dynamic OLTP and Batch infrastructure
  - Batch as a service
  - Replacing existing java batch frameworks
  - Sharing business logic across OLTP and Batch

Batch Modernization Use-case

- **Motivations for modernization**
  - IT departments are challenged to absorb tremendous growth rates while executing with a constant IT budget

- **Primary strategic goals for batch modernization**
  1. No loss of performance.
     Can be achieved with: JIT compilers in Java, parallelization, caching, etc.
  2. No loss of qualities of service such as job restart, availability, security
  3. Reduced operations costs. Primarily delivered through zAAP processors

- **Secondary strategic goals**
  1. A more agile runtime infrastructure that can better tolerate future changes
  2. Common development, testing, deployment, security, and production management processes and tooling across OLTP and Batch

Batch Modernization with XD Compute Grid

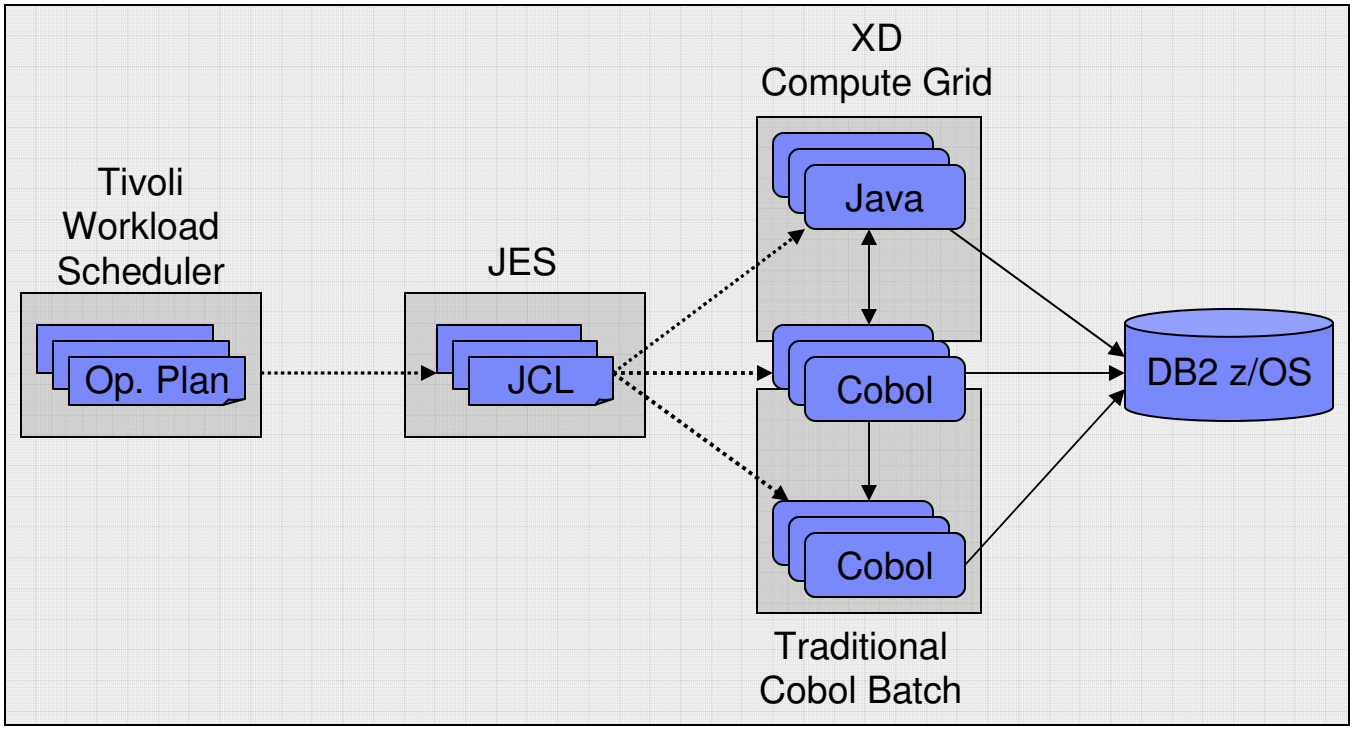

Today: Executing tradition batch with Cobol

Phase 1: Implement new business logic in java with XD Compute Grid

Phase 2: Share existing Cobol modules across both batch domains

Phase 3: Incrementally migrate remaining Cobol Modules to Java with XD Compute Grid

Completion: All Cobol batch modules are replaced with java and are running in XD Compute Grid

Presentation Outline (Sections)

- Batch Computing
  - What is it? What are batch jobs?
- Challenges Related to Batch Computing
  - How does WebSphere XD address?
- Batch and SOA
- What is WebSphere XD?
- Positioning XD Compute Grid in the world of Batch
- XD Compute Grid Use-cases
  - Batch Modernization
  - **Highly parallel batch jobs**
  - Dynamic OLTP and Batch infrastructure
  - Batch as a service
  - Replacing existing java batch frameworks
  - Sharing business logic across OLTP and Batch

Highly Parallel Batch Jobs

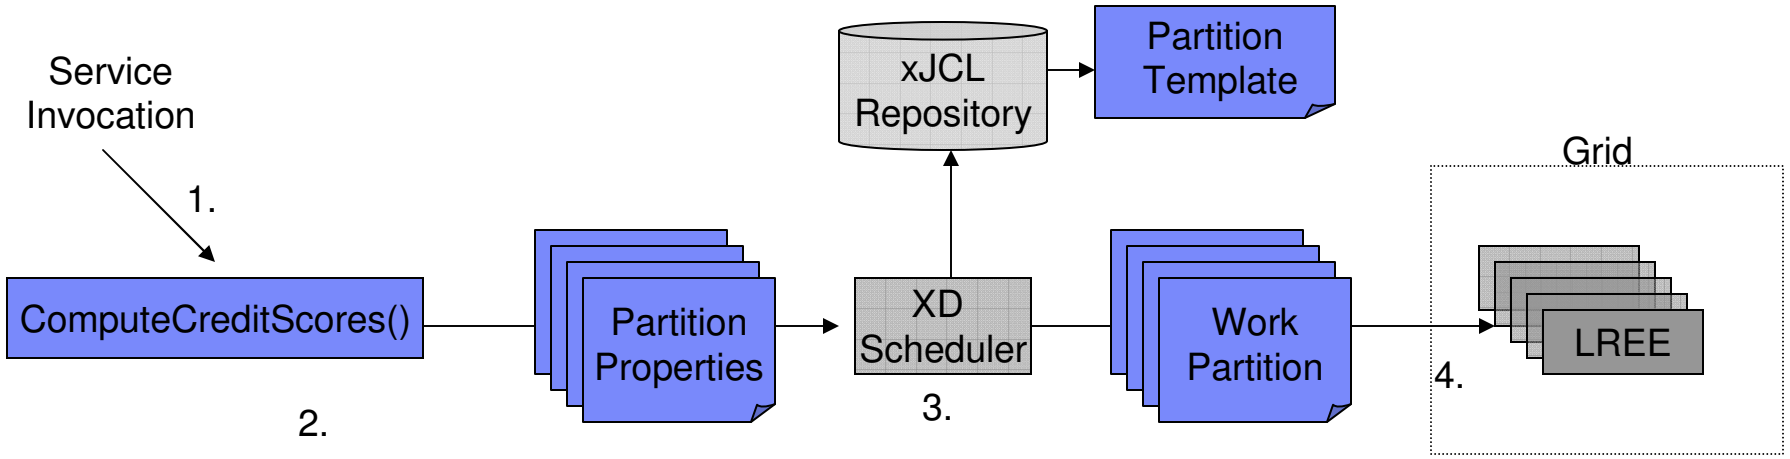

1. The exposed service ComputeCreditScores() is invoked
2. This service partitions a single, large batch job into many smaller jobs and submits them to the XD Scheduler
3. The XD Scheduler receives each job, pre-processes them, and sub-dispatches to the grid of execution endpoints (LREE)
4. Grid of LREE's execute the partitions as independent units of long-running work in parallel, ensuring that qualities of service such as transactionality, restartability, and so on are maintained.

Presentation Outline (Sections)

- Batch Computing
  - What is it? What are batch jobs?
- Challenges Related to Batch Computing
  - How does WebSphere XD address?
- Batch and SOA
- What is WebSphere XD?
- Positioning XD Compute Grid in the world of Batch
- XD Compute Grid Use-cases
  - Batch Modernization
  - Highly parallel batch jobs
  - **Dynamic OLTP and Batch infrastructure**
  - Batch as a service
  - Replacing existing java batch frameworks
  - Sharing business logic across OLTP and Batch

Dynamic OLTP and Enterprise Grid Runtime

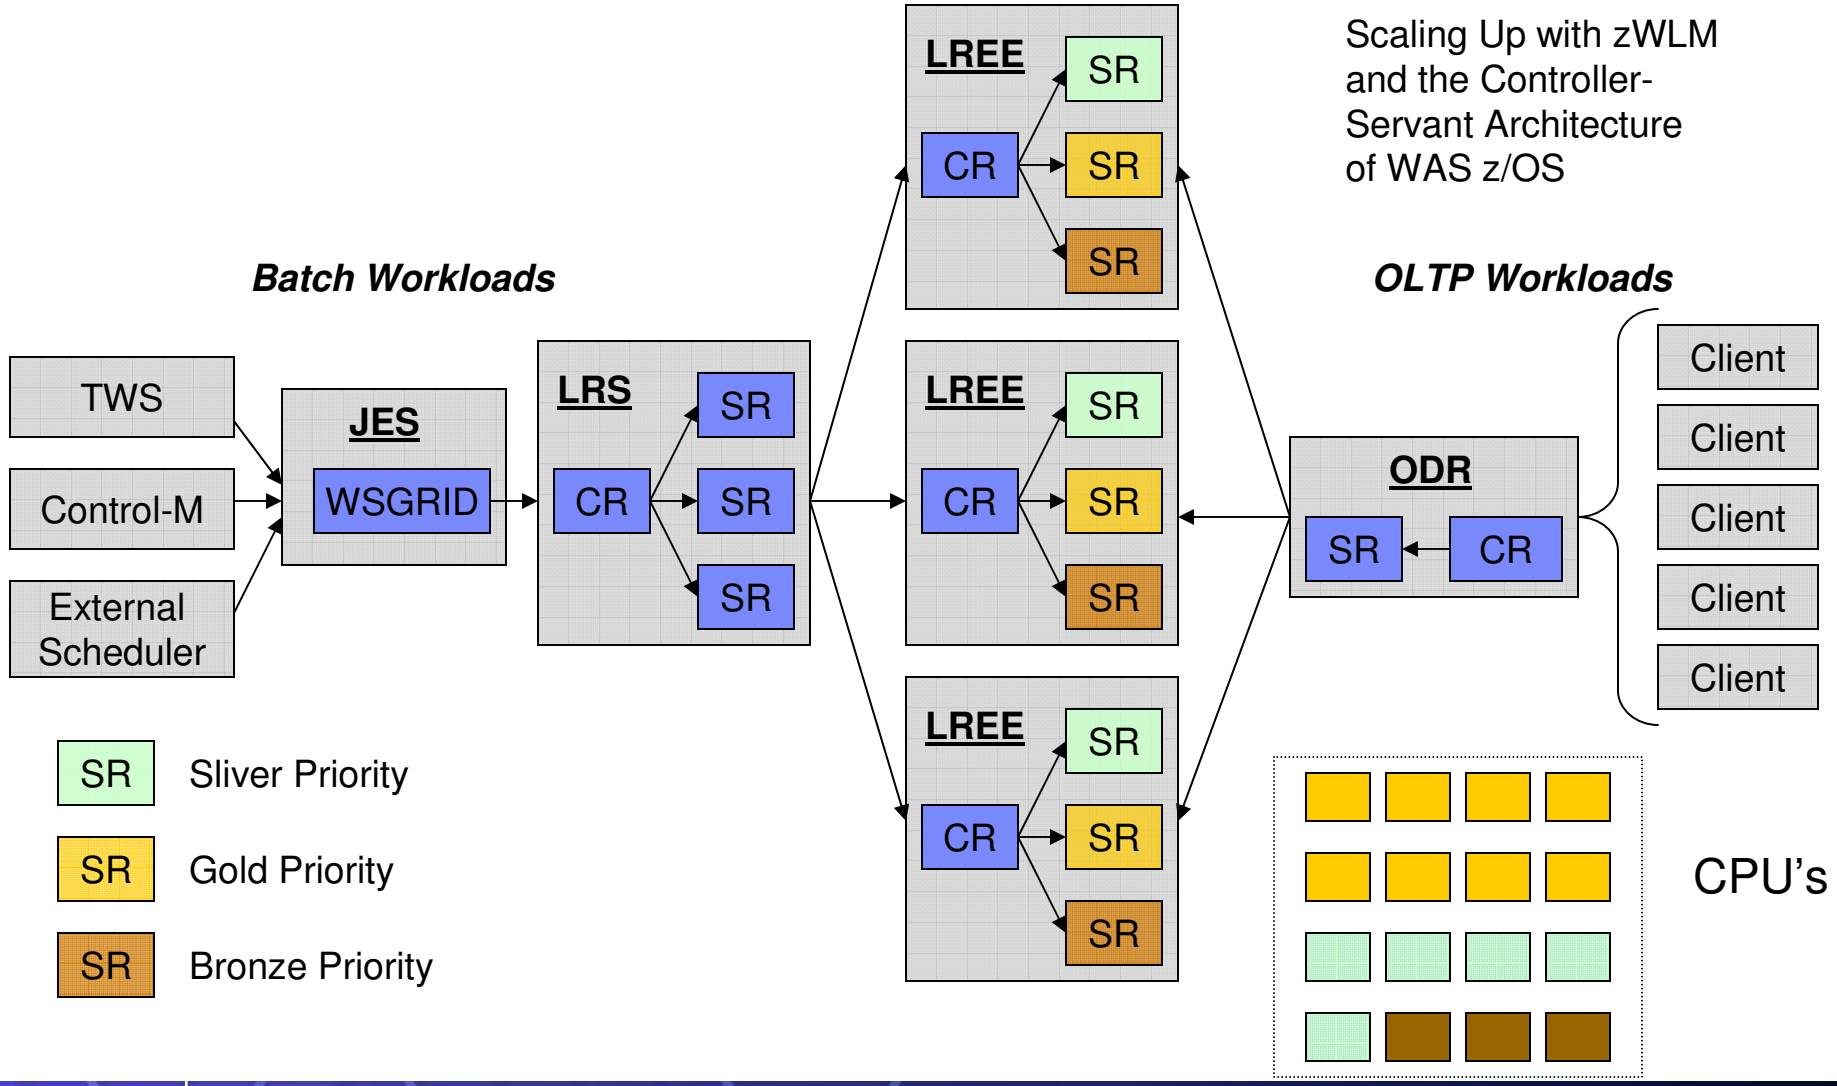

Presentation Outline (Sections)

- Batch Computing
  - What is it? What are batch jobs?
- Challenges Related to Batch Computing
  - How does WebSphere XD address?
- Batch and SOA
- What is WebSphere XD?
- Positioning XD Compute Grid in the world of Batch
- XD Compute Grid Use-cases
  - Batch Modernization
  - Highly parallel batch jobs
  - Dynamic OLTP and Batch infrastructure
  - **Batch as a service**
  - Replacing existing java batch frameworks
  - Sharing business logic across OLTP and Batch

Batch as a service

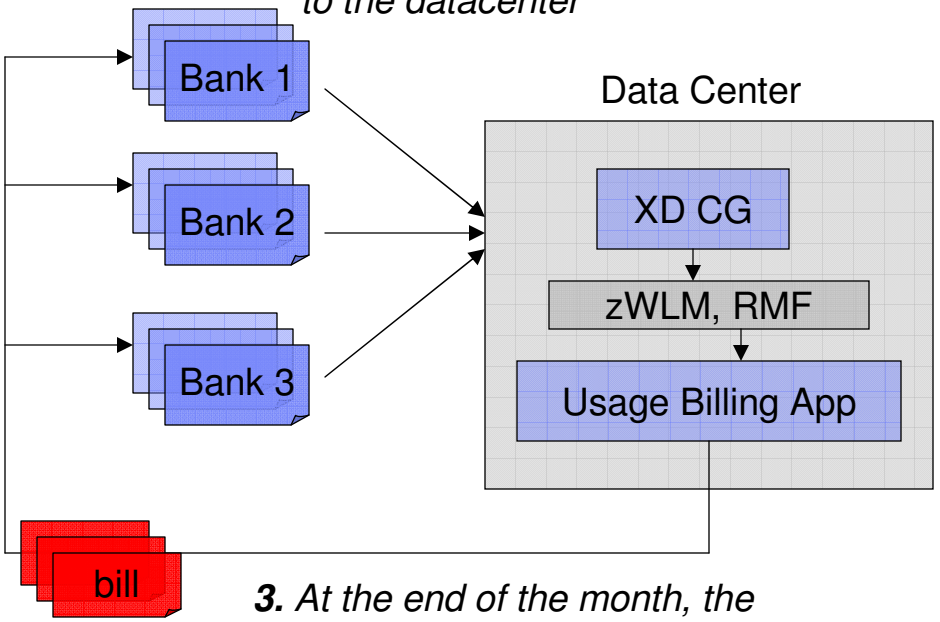

1. *Smaller banks submit job requests to the datacenter*

2. *Datacenter executes workloads for each bank, keep tracking of exactly how many resources each bank's jobs used by leveraging the usage accounting facilities of zWLM, RMF, and other system facilities of z/OS*

3. *At the end of the month, the datacenter sends bills for services rendered, based on the exact CPU seconds consumed, to each bank.*

Presentation Outline (Sections)

- Batch Computing
  - What is it? What are batch jobs?
- Challenges Related to Batch Computing
  - How does WebSphere XD address?
- Batch and SOA
- What is WebSphere XD?
- Positioning XD Compute Grid in the world of Batch
- XD Compute Grid Use-cases
  - Batch Modernization
  - Highly parallel batch jobs
  - Dynamic OLTP and Batch infrastructure
  - Batch as a service
  - **Replacing existing java batch frameworks**
  - Sharing business logic across OLTP and Batch

An Example: Java Batch Pattern (z/OS) - Good

The home-grown solution required many components to be integrated together…

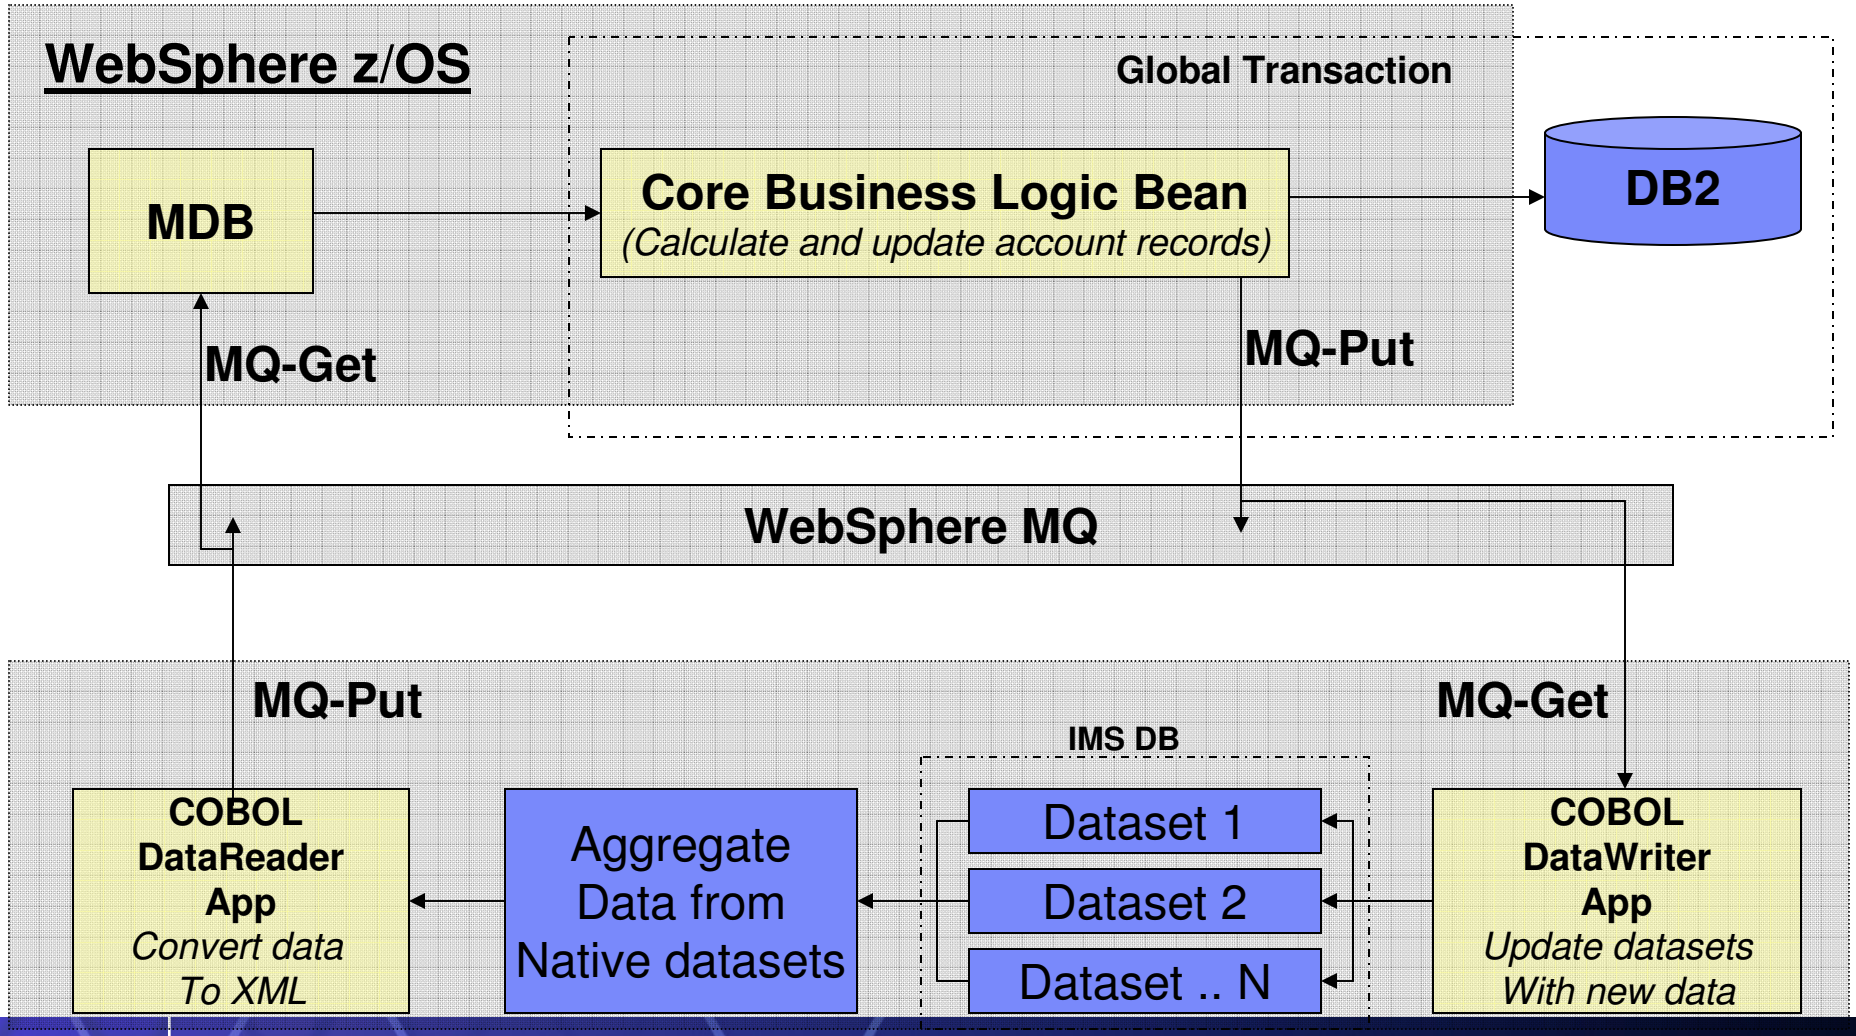

An Example: Java Batch Pattern (z/OS) - Better

WebSphere XD streamlines the Java-centric batch processing on z/OS…

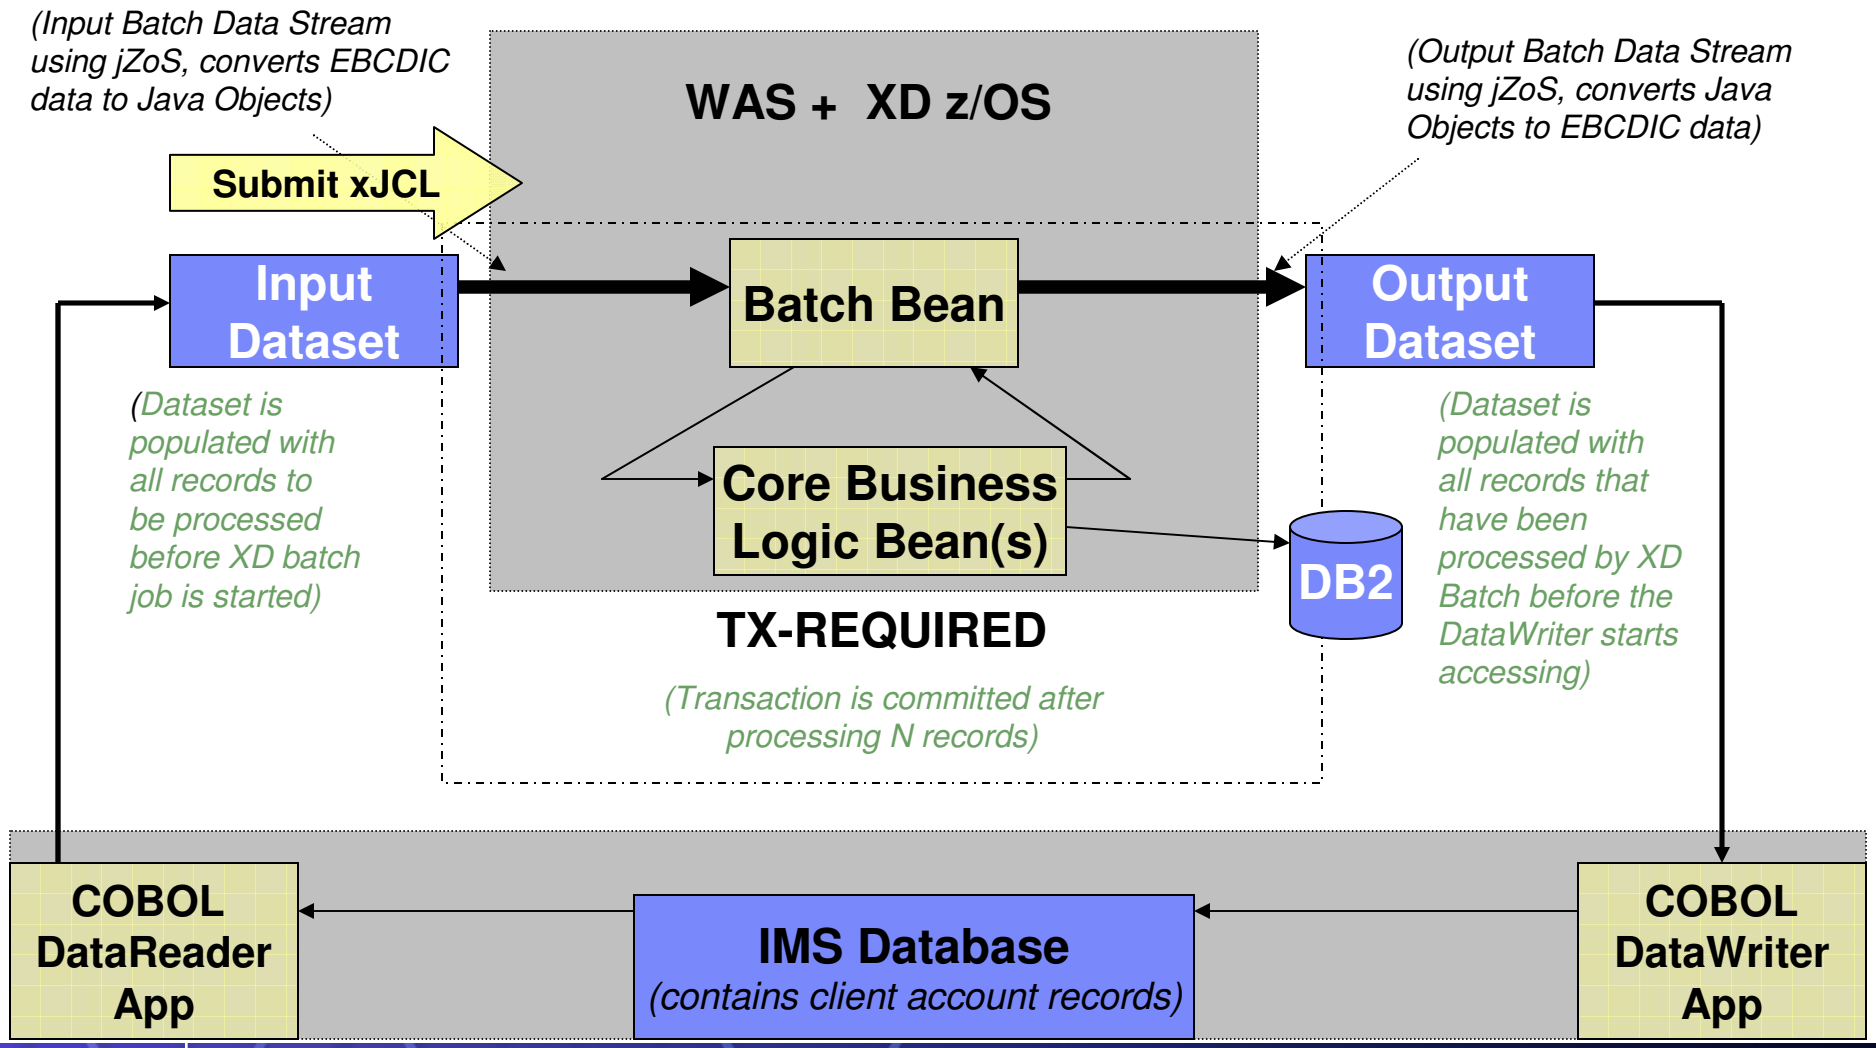

Presentation Outline (Sections)

- Batch Computing
  - What is it? What are batch jobs?
- Challenges Related to Batch Computing
  - How does WebSphere XD address?
- Batch and SOA
- What is WebSphere XD?
- Positioning XD Compute Grid in the world of Batch
- XD Compute Grid Use-cases
  - Batch Modernization
  - Highly parallel batch jobs
  - Dynamic OLTP and Batch infrastructure
  - Batch as a service
  - Replacing existing java batch frameworks
  - **Sharing business logic across OLTP and Batch**

Running Mixed Workloads- OLTP and Batch

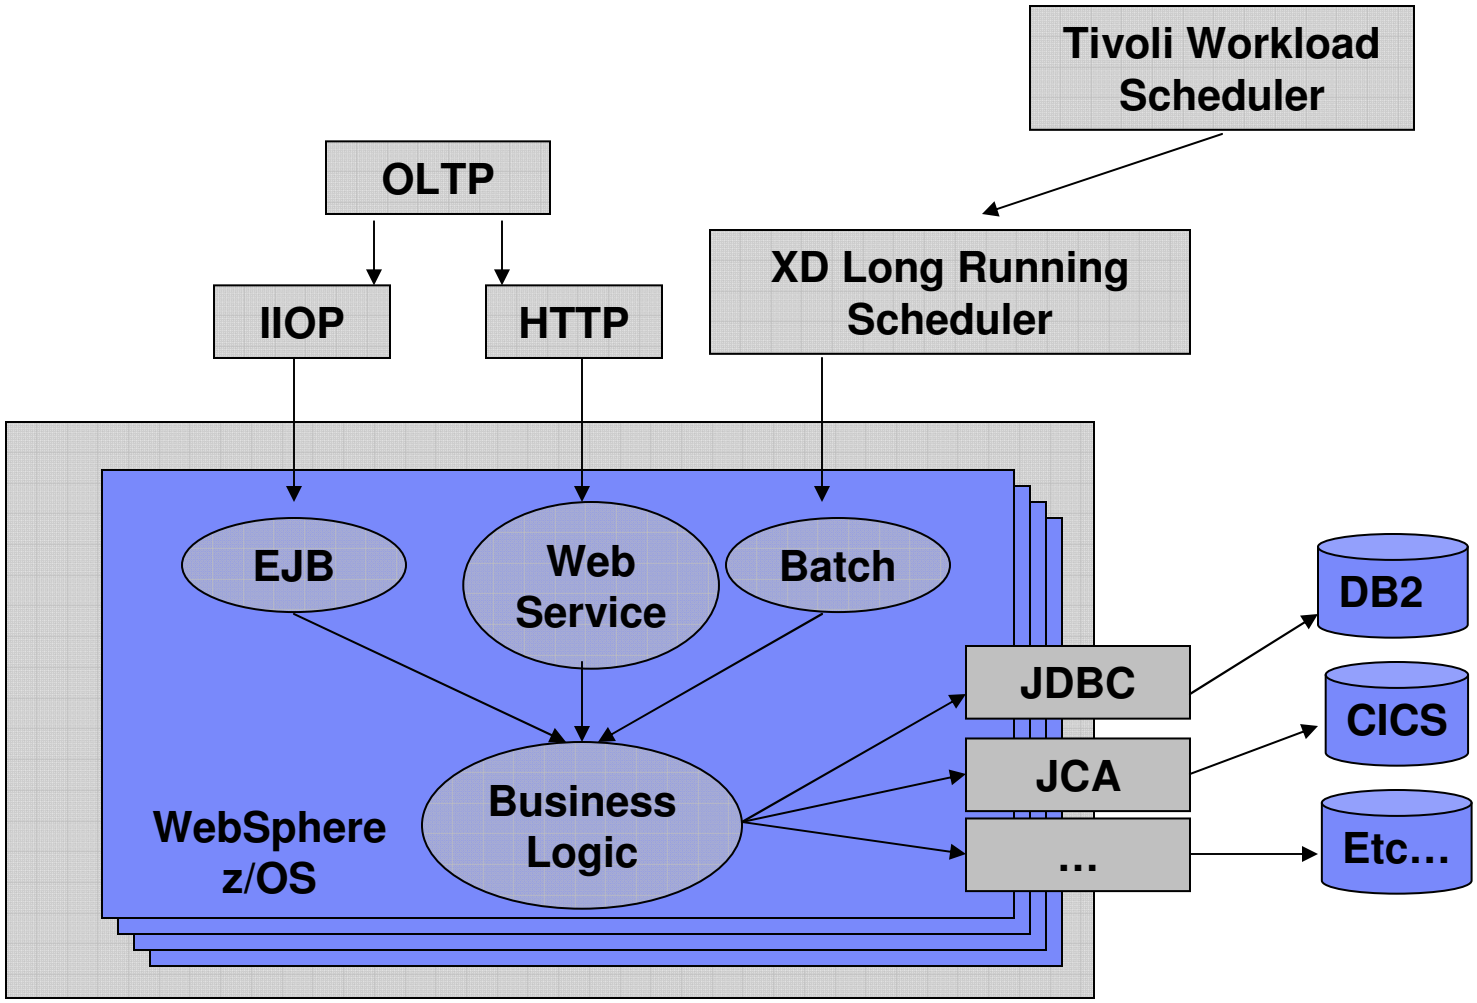

Conclusion

Summary of XD Compute Grid Value Proposition

1. Delivers ***a zAAP-eligible enterprise java batch*** execution environment built on WebSphere for z/OS

2. Enables the ***incremental migration of COBOL to Java*** thereby reducing the risks associated with a batch modernization project

3. ***Integrates with existing enterprise batch schedulers*** such as TWS, CA7, Control-M, Zeke to help deliver a robust, cost-effective, WebSphere-based batch execution environment

4. Enables ***new execution patterns including***: Dynamic OLTP and Batch runtime environment built on WebSphere for z/OS; highly parallel batch jobs; and many others.

5. Integrates with the ***overall SOA strategy of reuse*** by enabling one to share business logic across both the OLTP and Batch worlds

6. Delivers ***high-performance batch processing*** by leveraging the System-z, z/OS, and WAS z/OS performance optimizations gained when executing within close proximity of the data.

WebSphere Extended Deployment: Information Links

Announcement IBM WebSphere Extended Deployment V6.1

IBM Offering Information (United States)

- IBM WebSphere Extended Deployment V6.1 Software Announcement 207-088 PDF File
- IBM WebSphere Extended Deployment for z/OS, V6.1 Software Announcement 207-087 PDF File

IBM WebSphere XD V6.1 Product Information

- IBM WebSphere Extended Deployment Product Overview http://www.ibm.com/software/webservers/appserv/extend/
- **IBM WebSphere XD Features and Benefits** http://www.ibm.com/software/webservers/appserv/extend/features/
- **IBM WebSphere XD Product Documentation** http://www.ibm.com/software/webservers/appserv/extend/library/
- **IBM WebSphere XD InfoCenter** http://publib.boulder.ibm.com/infocenter/wxdinfo/v6r1/index.jsp
- **Redbook: Best Practices for Implementing WebSphere Extended Deployment, SG24-7343** http://www.redbooks.ibm.com/redbooks/pdfs/sg247343.pdf
- **developerWorks: Extended Deployment** http://www.ibm.com/developerworks/websphere/zones/xd/
- **IBM WebSphere Application Server Network Deployment Product Overview** http://www.ibm.com/software/webservers/appserv/was/network/
- **The IBM Solution Assurance Library (Search for Extended Deployment)** http://w3.ibm.com/support/assure/assur30i.nsf/Web/SA

Thank You for Joining Us today!

Go to **www.ibm.com/software/systemz** to:

- Replay this teleconference
- Replay previously broadcast teleconferences
- Register for upcoming events