

IBM WebSphere Application Server

Batch applications overview

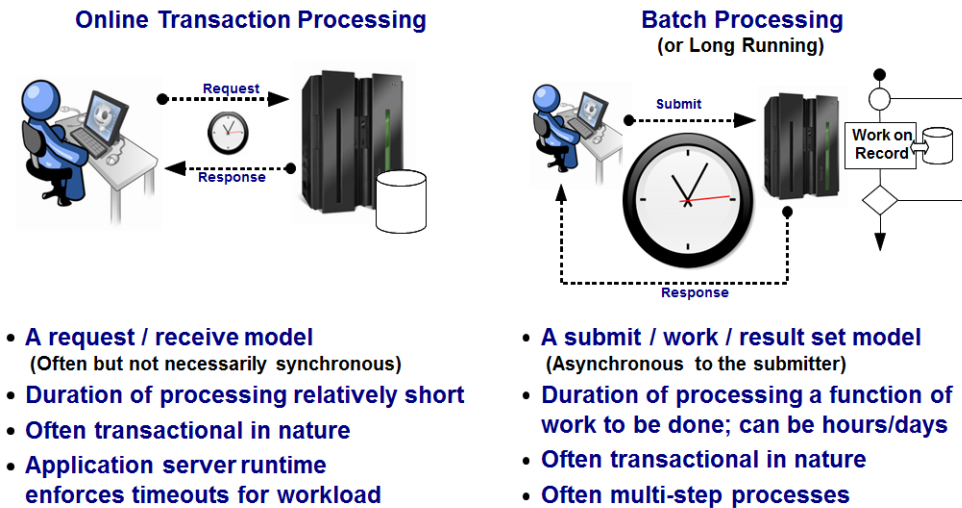


©2011 IBM Corporation

This module provides an overview of the batch applications feature of IBM WebSphere Application Server.

Batch processing

Comparison of online transaction processing and batch processing



2

Overview

© 2011 IBM Corporation

Online transaction processing is a request/receive model where the duration of the processing is relatively short, and the tasks are typically transactional in nature. In this model, the application server runtime enforces timeouts for the workload.

Batch processing is a submit/work/result set model where the duration of the processing is a function of the tasks to be completed. In some cases, the tasks can require hours or even days to complete. In this model, the work tasks are typically transactional in nature and typically involve multi-step processes.

Batch modernization

Providing greater flexibility / control of batch window processing

- Moving to a continuously available batch model intermixed with OLTP
- Using goal-oriented workload management to enable job completion within deadlines

SOA transformation - Running batch as a service

- Hosting batch processes within SOA-capable application server environments
- Optimally reusing services through efficient co-location

Leveraging Java skills for all applications

- Utilizing Java for batch processing and OLTP

Reducing cost of running batch on z/OS

- Moving batch processing to Java to take advantage of zAAP specialty engines
- Utilizing processor resources available when OLTP workloads are light


Modernization does not mean wholesale replacement of existing batch function that serves the business needs. Business imperatives drive technical solutions

Batch modernization involves these techniques:

- Providing greater flexibility and control of batch window processing by moving to a continuously available batch model intermixed with online transaction processing and using goal-oriented workload management to enable job completion within deadlines.
- Hosting batch processes within service-oriented architecture (SOA) capable WebSphere Application Server environments and optimally reusing services through efficient co-location.
- Leveraging Java skills for batch and online transaction processing tasks.
- On the z/OS platform, reducing costs by moving batch processing to Java to take advantage of System z Application Assist Processor (zAAP) specialty engines and utilizing processor resources available when online transaction processing workloads are light.

Use JVM launcher

Many attempt to use a JVM launcher to run Java batch programs



This approach works, but has limitations:

- High overhead associated with repeated cycling of JVMs
- Application must provide most batch processing services

4

Overview

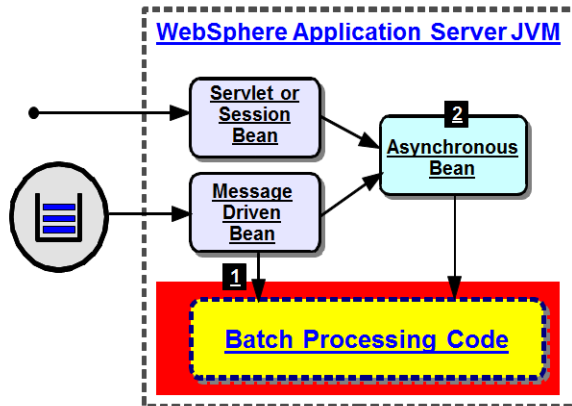
© 2011 IBM Corporation

Java Virtual Machine (JVM) launcher technologies are easy to use and free, and the programming interfaces are very useful; however, a new JVM is created for every step, and the resulting overhead is significant. In addition, JVM launcher technologies do not provide a transaction manager. Ultimately, the application developer must author a transaction manager within their application.

On the z/OS platform, WebSphere Application Server and WebSphere XD Compute Grid complement the Java Batch Toolkit for z/OS (JZOS). The JZOS launcher can be used when only a few Java batch steps must be executed. Once the size of the Java batch infrastructure grows to the point where the overhead of creating and terminating JVMs is an issue, WebSphere Application Server and WebSphere XD Compute Grid can be used. The second component of JZOS, its API's for accessing z/OS resources, can be used from WebSphere Application Server and WebSphere XD Compute Grid applications.

Use application server runtime

A second approach uses the application server runtime and authors the batch application as a message driven bean or an asynchronous bean.



1. When a message driven bean is used, the application must manage the timeout on the message driven bean input
2. The preferred technique (used by WebSphere Application Server and WebSphere XD Compute Grid) is to use an asynchronous bean

In both cases, the application developer must author custom middleware to support the batch processing

5

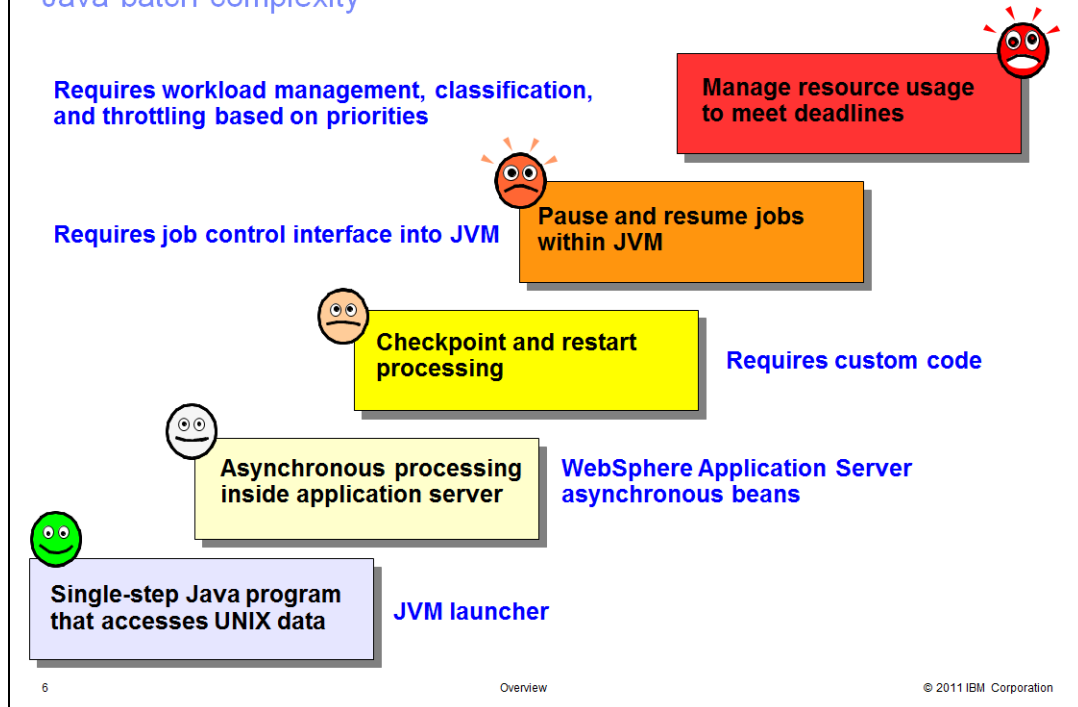
Overview

© 2011 IBM Corporation

A second approach uses the application server runtime and authors the batch application as message driven bean or an asynchronous bean. When a message driven bean is used, the application must manage the timeout on the message driven bean input; therefore, using an asynchronous bean is the preferred technique. However, in both cases, the application developer must author custom middleware to support the batch processing, given that the solution lacks these features:

- A batch programming model, job definition language, and supporting development tools
- Interweaving of online transaction processing tasks and batch tasks
- Checkpointing and restarting of jobs
- Integration with enterprise schedulers
- Graphical and command line interfaces for job monitoring and management
- And Job usage accounting and activity logging

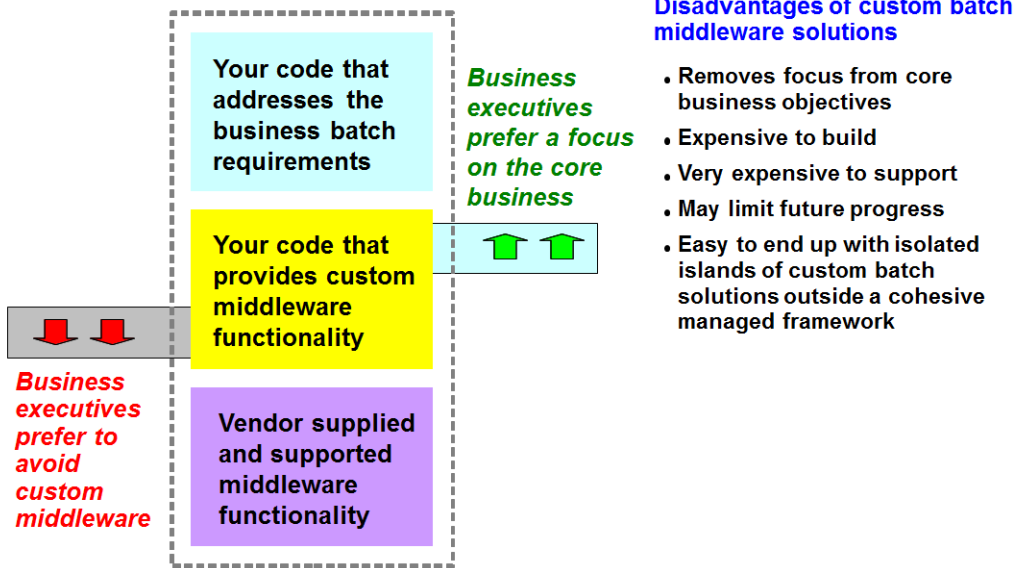
Java batch complexity



Initial Java batch programming projects typically begin with a relatively simple Java Virtual Machine (JVM) launcher design, but it grows more complex over time as additional management features are required. These projects frequently progress in this manner:

- An organization uses the Java command line interface or the Java Batch Toolkit for z/OS (JZOS) to run a single-step Java program that access UNIX data.
- After realizing that transaction management capabilities are required, the organization moves the application to a Java Enterprise Edition (Java EE) application server.
- As the organization begins to process large numbers of records as a single batch job, it finds that it is no longer practical to restart entire jobs when failures occur in the middle of a job. The organization then introduces checkpointing and restart capability.
- A large number of jobs in the batch environment compete for the same resources, and administrators need a way to manage the jobs and handle resource conflicts. The organization then builds a job control interface into the JVM that processes the batch applications.
- As batch loads increase even further, it becomes impossible to manually handle resource conflicts; as a result, the organization now requires a batch job management system that can optimally distribute workloads based on priorities and deadlines.

Custom middleware versus core business



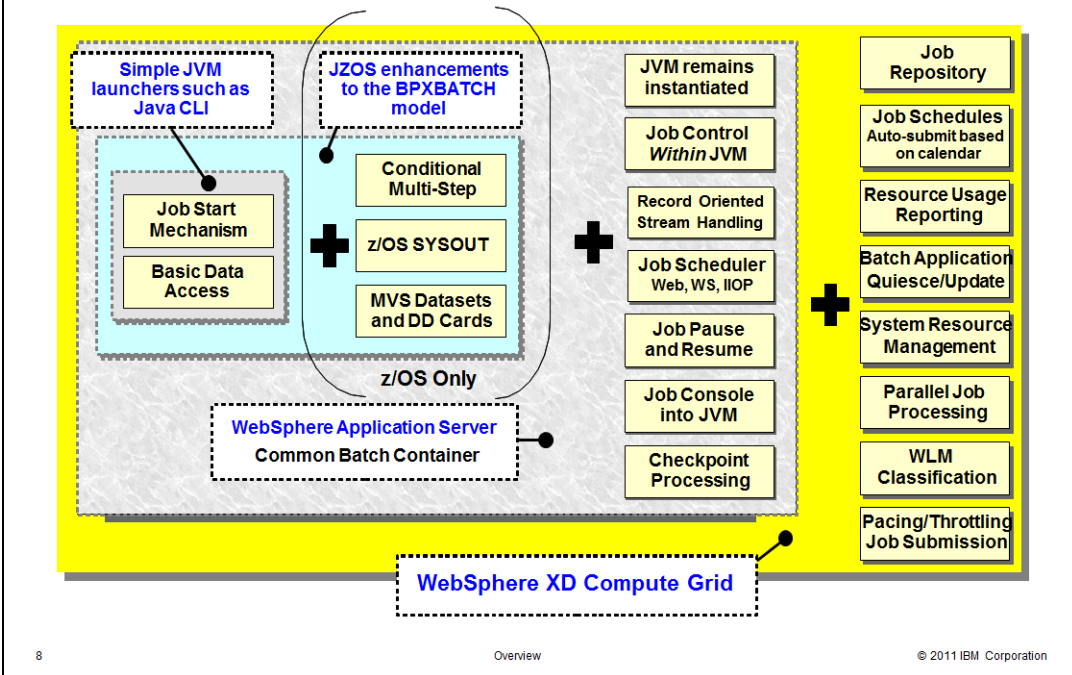
7

Overview

© 2011 IBM Corporation

Organizations are frequently tempted to solve tactical issues with custom middleware code; however, the middleware code requires extensive time to develop and maintain, which ultimately decreases an organization's ability to develop custom code that addresses core business objectives.

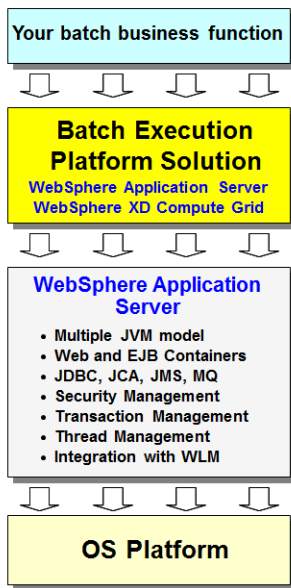
IBM Java batch offerings



Java batch tasks can be managed in several different ways:

- Simple Java Virtual Machine (JVM) launchers can manage single step batch jobs that require basic data access.
- On the z/OS platform, the Java Batch Toolkit for z/OS (JZOS) enhances the BPXBATCH model by supporting conditional multi-step batch jobs with access to MVS datasets and use of Data Definition (DD) cards.
- WebSphere Application Server provides multi-step job support, a managed container for execution of batch jobs, a job control interface, job checkpoint and restart capability, and a batch application development framework.
- Finally, WebSphere XD Compute Grid adds support for job repository and schedules, workload management, job usage reporting, batch application quiesce and update, parallel job support, and pacing and throttling of jobs.

Java batch foundation



Key benefits

- Custom middleware code is no longer required, so developers can stay focused on core business imperatives
- Online Transaction Programming (OLTP) and batch loads can be mixed while maintaining respective service level agreements

The basic IBM Java batch foundation includes WebSphere Application Server. Customers demanding advanced workload management, accounting, and job scheduler integration capabilities can use WebSphere Application Server and WebSphere XD Compute Grid.

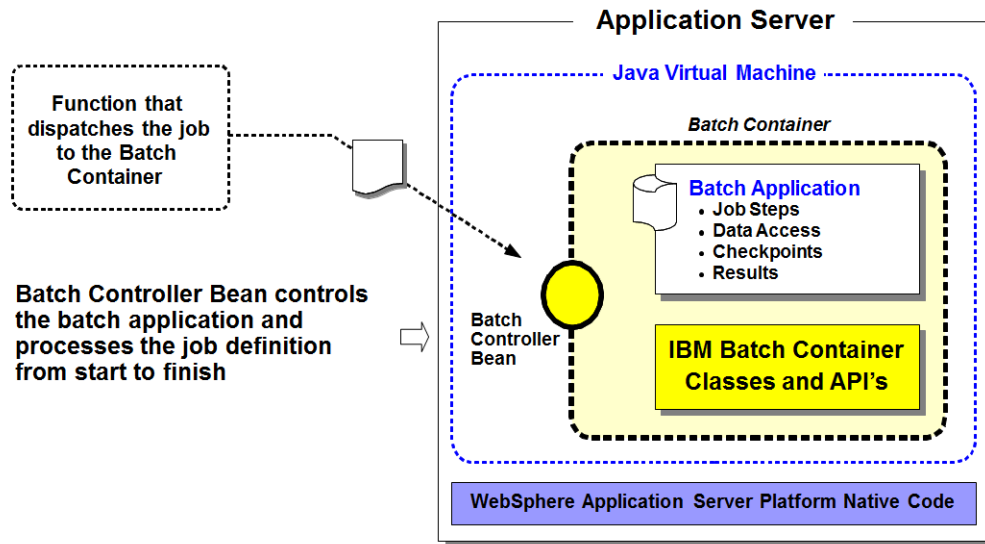
Both solutions provide these benefits:

-Custom middleware code is no longer required, so developers can stay focused on core business imperatives.

-Online Transaction Programming (OLTP) and batch loads can be mixed while maintaining the service level agreements for each type of traffic.

Batch Container

The Batch Container is the heart of the WebSphere Application Server and WebSphere XD Compute Grid products



10

Overview

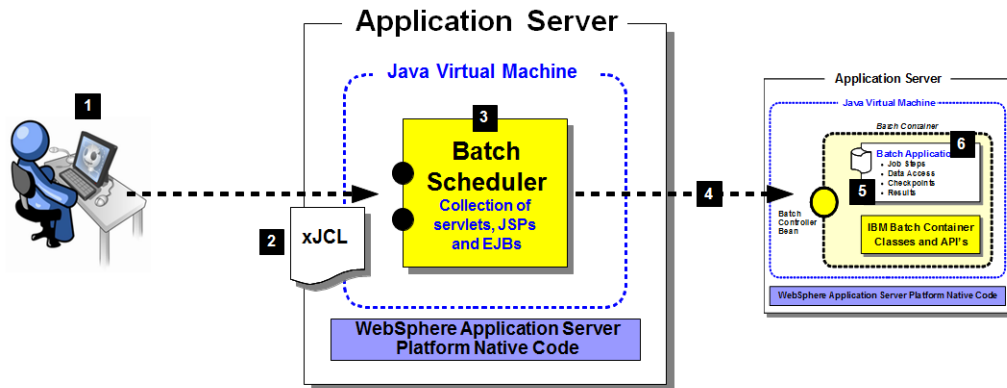
© 2011 IBM Corporation

The Batch Container is the heart of the batch application support provided in the WebSphere Application Server and WebSphere XD Compute Grid offerings. It runs a batch job under the control of an asynchronous bean, which can be thought of as a container-managed thread. The batch container ultimately processes a job definition and carries out the life cycle of a job.

The Batch Container provides these services:

- Checkpointing, which involves resuming batch work from a selected position.
- Result processing, which involves intercepting and processing step and job return codes.
- And Batch data stream management, which involves reading, positioning, and repositioning data streams to files, relational databases, native z/OS datasets, and many other different types of input and output resources.

Batch components and workflow



- | | |
|--|--|
| 1. Job is submitted | 4. Job is dispatched to batch endpoint |
| 2. Job Control Definition is specified | 5. Batch endpoint begins execution |
| 3. Scheduler analyzes the request | 6. Batch application is invoked |

11

Overview

© 2011 IBM Corporation

-Batch jobs are submitted to the system using the Job Management Console or programmatically by way of Enterprise Java Beans (EJB), Java Message Service (JMS), or web services.

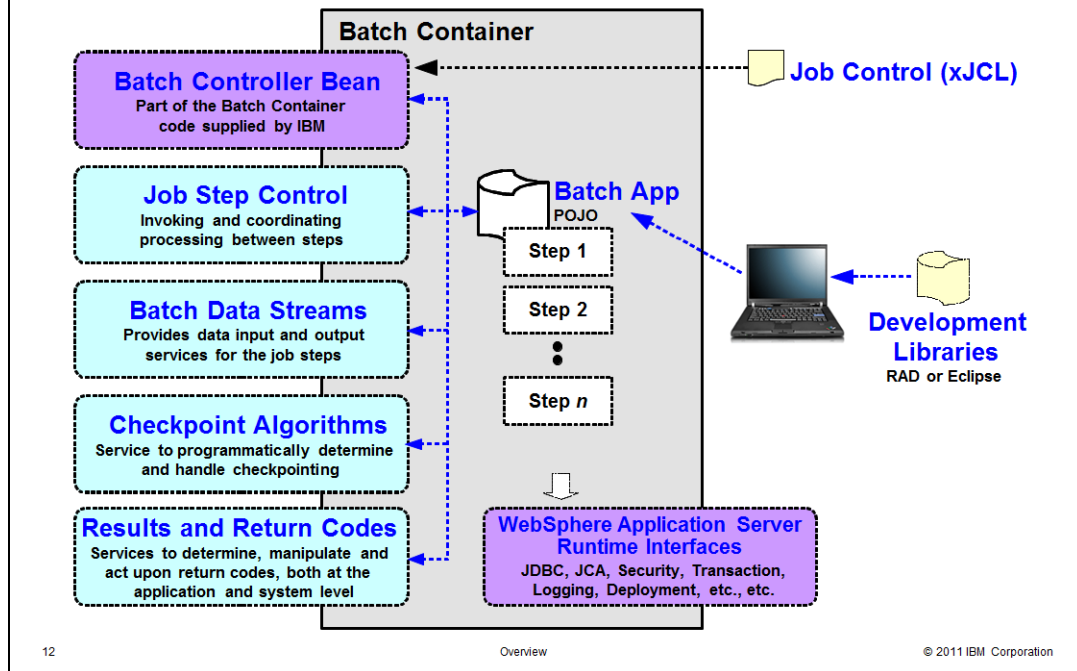
-Each job is submitted in the form of an XML Job Control Language (xJCL) document.

-The Job Dispatcher selects the best endpoint for job execution based on several different metrics.

-The endpoint sets up the jobs in the Batch Container and executes the batch steps based on the definitions in the xJCL.

-The Job Dispatcher aggregates job logs and provides life-cycle management functions like start, stop, cancel, and so on.

Batch programming model



The batch programming model consists of four principal interfaces, two of which are essential to building a batch application, and two which are optional and intended for advanced scenarios:

- The batch job step defines the interaction between the batch container and the batch application.
- The batch data stream abstracts a particular input source or output destination for a batch application and defines the interaction between the batch container and a concrete BatchDataStream implementation.
- A checkpoint policy algorithm defines the interaction between the batch container and a custom checkpoint policy implementation. A checkpoint policy is used to determine when the batch container will checkpoint a running batch job to enable a restart to occur after a planned or unplanned interruption. WebSphere Application Server and WebSphere XD Compute Grid include two ready-to-use checkpoint policies.
- A results algorithm defines the interaction between the batch container and a custom results algorithm. The purpose of the results algorithm is to provide the overall return code for a job. The algorithm has visibility to the return codes from each of the job steps. WebSphere Application Server and WebSphere XD Compute Grid include one ready-to-use results algorithm.

Job control definition - xJCL

```
<?xml version="1.0" encoding="UTF-8" ?>
```

```
<job name="name" ... >
  <jndi-name>batch_controller_bean_jndi</jndi-name>
  <substitution-props>
    <prop name="property_name" value="value" />
  </substitution-props>
```

**Roughly analogous
to the JOB card**

```
<job-step name="name">
  <classname>package.class</classname>
  <checkpoint-algorithm-ref name="chkpt"/>
  <results-ref name="jobsum"/>
  <batch-data-streams>
    <bds>
      <logical-name>input_stream</logical-name>
      <props>
        <prop name="name" value="value" />
      </props>
    </bds>
  </batch-data-streams>
</job-step>
```

A job step

**Like the EXEC PGM=
statement in JCL**

**Similar to DD
statements**

```
</job-step>
</job>
```

13

Overview

© 2011 IBM Corporation

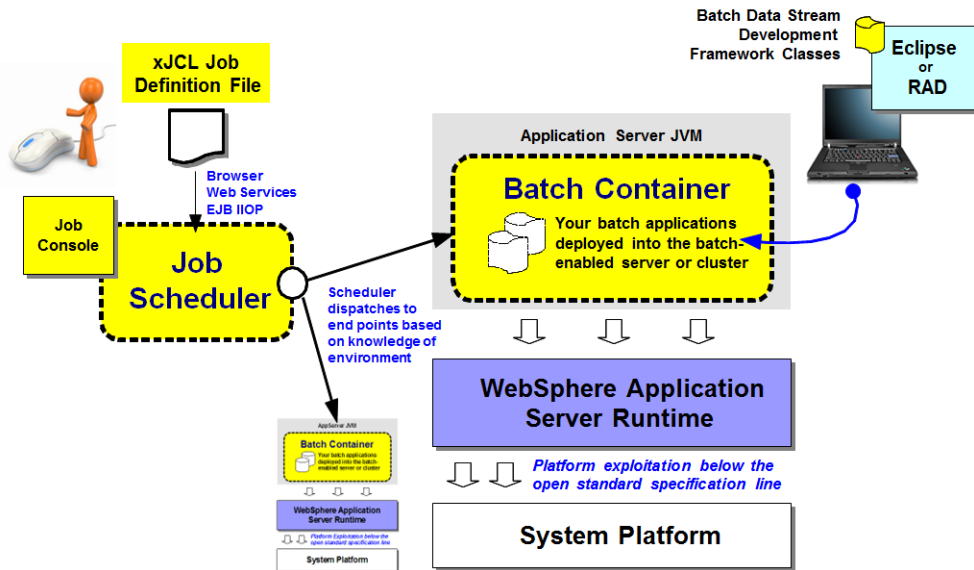
Jobs are expressed using an Extensible Markup Language (XML) dialect called XML Job Control Language, or xJCL.

The xJCL definition is very similar to the traditional JCL.

The xJCL definition of a job is not part of the batch application, but is constructed separately and submitted to the job scheduler to run. The job definition identifies which batch application to run, and its inputs and outputs. It also identifies which checkpoint algorithms and results algorithms to use.

The Job Scheduler uses information in the xJCL to determine where and when the job should be run.

Batch environment review



14

Overview

© 2011 IBM Corporation

The Job Scheduler provides the job management functions such as submit, cancel, restart, and so on. It maintains a history of all job activity, including waiting jobs, and running jobs, and completed jobs. The job scheduler is hosted in a WebSphere Application Server or a server cluster.

Jobs are described using job control language called XML Job Control Language (xJCL), which identifies the batch application to run, its inputs and outputs.

The Batch Container provides the execution environment for batch jobs. There can be multiple Batch Containers in a WebSphere cell.

Batch applications are regular WebSphere Java Enterprise Edition (Java EE) applications, deployed as Enterprise Archive (EAR) files.

Deployment options

	WebSphere Application Server	WebSphere Application Server Network Deployment or z/OS	WebSphere Application Server Network Deployment with Compute Grid
Common batch container, development tools, and operational commands to manage batch job life cycle	■	■	■
Container managed checkpoint/restart capabilities	■	■	■
Job management console	■	■	■
Application execution platform	■	■	■
Basic scheduler/job dispatcher	■	■	■
System-managed job logs	■	■	■
High-availability and clustering of Job Scheduler/Dispatcher		■	■
Multi-site disaster recovery for batch platform		■	■
Integration with workload management on z/OS		■	■
Interoperability between Java and COBOL on z/OS			■
Non-disruptive batch application update/endpoint quiesce			■
Job usage accounting, including SMF integration on z/OS			■
Job classes and workload classification			■
Integrated Parallel Job Manager for job parallelization across multiple JVM's			■
Enterprise scheduler connectors			■
Enterprise monitoring capabilities			■
Disaster recovery with operational state transfer			■
Integration with WebSphere Virtual Enterprise for goal-oriented job placement			■

15

Overview

© 2011 IBM Corporation

IBM offers several different Java batch environments:

- WebSphere Application Server Base Edition
- WebSphere Application Server Network Deployment Edition
- WebSphere Application Server for z/OS

And WebSphere XD Compute Grid

WebSphere Application Server Base Edition provides these features:

- A managed batch container
- A batch application framework
- A job management console
- A job dispatcher
- And System-managed job logs

WebSphere Application Server Network Deployment Edition and WebSphere Application Server for z/OS provides these additional features:

- High-availability and clustering of the job dispatcher
- Multi-site disaster recovery support
- And Integration with z/OS workload management

WebSphere XD Compute Grid provides these features beyond those provided by WebSphere Application Server:

- Job classes and workload classification
- Job parallelization across multiple Java Virtual Machines
- Job usage accounting
- Connectors for integration with enterprise schedulers
- Enterprise monitoring capabilities
- Interoperability between Java and COBOL on z/OS
- Non-disruptive batch application update support
- And Integration with WebSphere Virtual Enterprise for goal oriented job placement

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, System z, WebSphere, and z/OS are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java, and all Java-based trademarks and logos are trademarks of Oracle and/or its affiliates.

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

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 OR SOFTWARE.

© Copyright International Business Machines Corporation 2011. All rights reserved.