

# WebSphere Application Server V7 OSGi, JPA, and Modern Batch Feature Packs

David Follis and Don Bagwell  
IBM Corporation

Wednesday, August 4, 2010



**SHARE** in Boston

# Trademarks

**The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.**

CICS*	Parallel Sysplex*
DB2*	RACF*
GDPS*	System z9
Geographically Dispersed Parallel Sysplex	WebSphere*
HiperSockets	z/OS
IBM*	zSeries*
IBM eServer	
IBM logo*	
IMS	
On Demand Business logo	

\* Registered trademarks of IBM Corporation

**The following are trademarks or registered trademarks of other companies.**

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.

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

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

\* All other products may be trademarks or registered trademarks of their respective companies.

## Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

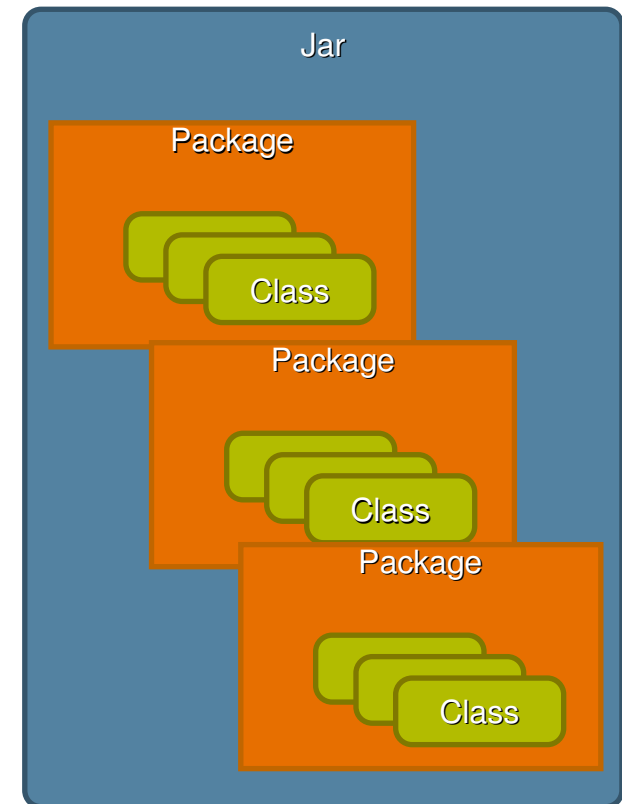
# Disclaimer

- The information contained in this documentation is provided for informational purposes only. While efforts were many to verify the completeness and accuracy of the information contained in this document, it is provided “as is” without warranty of any kind, express or implied.
- This information is based on IBM’s current product plans and strategy, which are subject to change without notice. IBM will not be responsible for any damages arising out of the use of, or otherwise related to, this documentation or any other documentation.
- Nothing contained in this documentation 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 the applicable license agreement governing the use of the IBM software.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
- All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

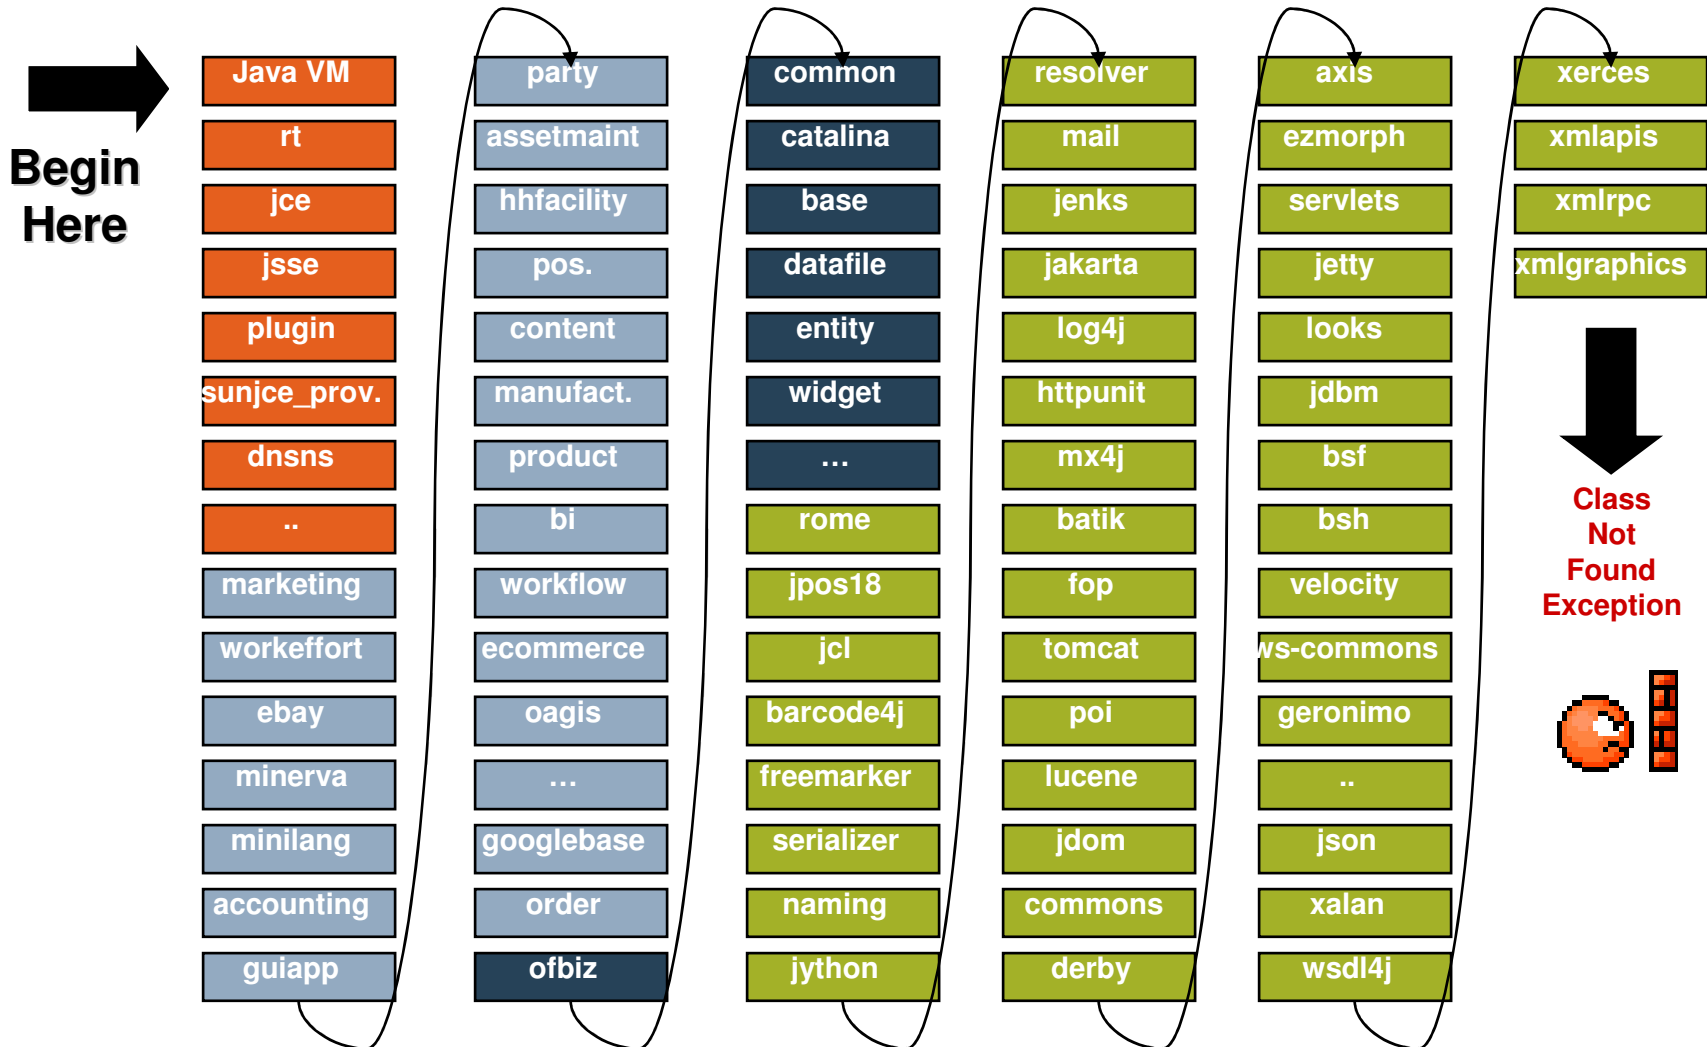
# OSGi Feature Pack

# Modularization in Java – Problems with Jars

- Java Platform Modularity
  - Classes encapsulate data
  - Packages contain classes
  - Jars contain packages
- Class visibility:
  - private, package private, protected, public
- No “jar scoped” access modifiers.
- No means for a jar to declare its dependencies.
- No versioning.
- Jars have no modularization characteristics
  - At runtime there is just a collection of classes on a global classpath




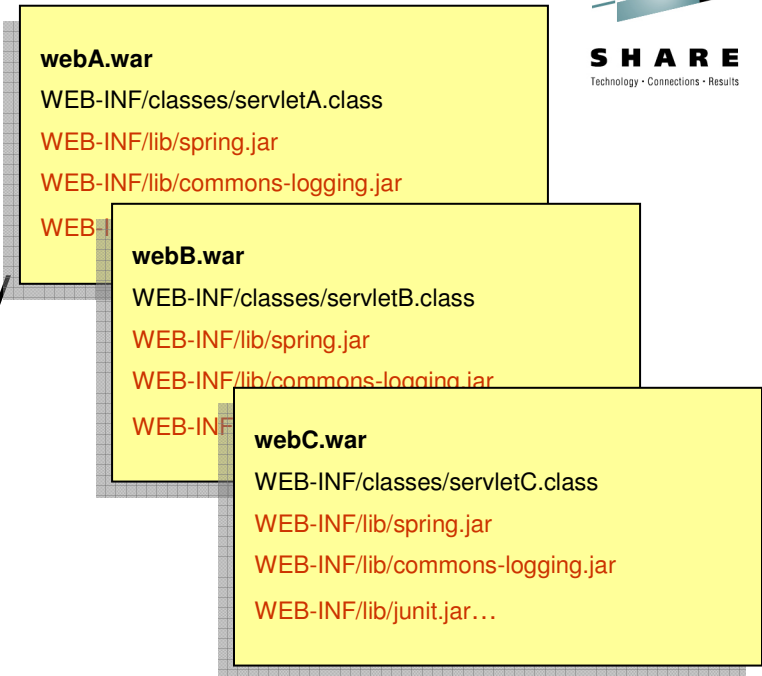
# Problems with Global Java ClassPath



# Problems with EARs/WARs

Enterprise Apps have isolated classpaths but...

- *Across applications* - each archive typically contains all the libraries required by the application
  - Common libraries/frameworks get installed with each application
  - Multiple copies of libraries in memory
- *Within applications* – third party libraries consume other third party libraries leading to version conflicts

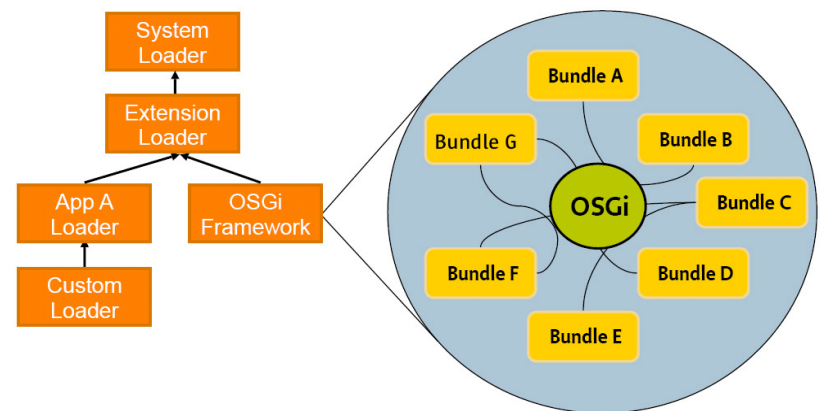


# OSGi Bundles and Class Loading

- OSGi Bundle – A jar containing:
  - Classes and resources.
  - OSGi Bundle manifest.
- What's in the manifest:
  - Bundle-Version: Multiple versions of bundles can live concurrently.
  - Import-Package: What packages from other bundles does this bundle depend upon?
  - Export-Package: What packages from this bundle are visible and reusable outside of the bundle?
- Class Loading
  - Each bundle has its own loader.
  - No flat or monolithic classpath.
  - Class sharing and visibility decided by declarative dependencies, not by class loader hierarchies.
  - OSGi framework works out the dependencies including versions.

```

Manifest-Version: 1.0
Bundle-ManifestVersion: 2
Bundle-Name: MyService bundle
Bundle-SymbolicName: com.sample.myservice
Bundle-Version: 1.0.0
Bundle-Activator: com.sample.myservice.Activator
Import-Package: com.something.i.need;version="1.1.2"
Export-Package: com.myservice.api;version="1.0.0"
  
```





# OSGi Enterprise Specification

- Released 22 March 2010
  - The product of the OSGi Enterprise Expert Group (EEG)
- Brings Enterprise technologies and OSGi together
- Using existing Java SE/EE specifications:
  - JTA, JPA, JNDI, JMX, WebApps...
- Adds Spring-derived *Blueprint* component model and DI container
- **Java EE provides the core enterprise application programming model**
- **Deploying modules as OSGi bundles simplifies reuse between applications, provides versioning, encourages (and enforces) modular design and enables dynamic module updates.**



# Enterprise OSGi in Open Source



- Apache “Aries” created as a new Apache incubator project in Sep 2009:
  - to provide enterprise OSGi spec implementations  
<http://incubator.apache.org/aries/>
  - to provide an environment to collaborate and experiment with new technologies to inform further EEG standardization.
    - In particular the programming model aspects of OSGi applications in an enterprise environment such as the Blueprint container and multi-bundle composites.
  - to build a broad development community to encourage implementation and adoption of EEG specs
- Aries componentry supporting an enterprise OSGi programming model are being integrated into both Geronimo and WebSphere Application Server.
  - As well as Apache Felix Karaf, JBossOSGi and others

# Application exploitation of OSGi in WebSphere

- OSGi has been used internally in WebSphere Application Server since V6.1 and in Eclipse since R3.
- Application-level exploitation is introduced in the **WebSphere Application Server Feature Pack for OSGi Applications and Java Persistence API (JPA) 2.0**
  - <http://www-01.ibm.com/software/webservers/appserv/was/featurepacks/>
  - Generally available May 2010
- Early Program available since Nov 2009
  - <https://www14.software.ibm.com/iwm/web/cc/earlyprograms/websphere/wasfposgiajp>
  - More downloads in a shorter period of time than any previous WebSphere Application Server v7 feature pack open beta
- Two installable features:
  - OSGi Application feature simplifies the development, assembly, and deploy of enterprise applications
  - JPA 2.0 feature introduces Java EE 6 JPA 2.0 enhancements to object-relational persistence to simplify data access and optimize performance



# New: Bundle Repository Configuration in WebSphere Application Server



Integrated Solutions Console Welcome Help | Logout Close page

Cell=irobinsNode01Cell, Profile=AppSrv01

**Internal bundle repository**

**Internal bundle repository > com.ibm.json.java**

The internal bundle repository is used to store bundles that are referenced by OSGi applications running in WebSphere Application Server. When an OSGi application is imported as an asset, the provisioner attempts to satisfy all its dependencies by using the contents of the asset, the contents of the internal bundle repository, and the contents of any available external bundle repositories.

Configuration

**General Properties**

Bundle symbolic name

Bundle version

Bundle name

Bundle description

Imported packages

Exported packages

# JPA Overview

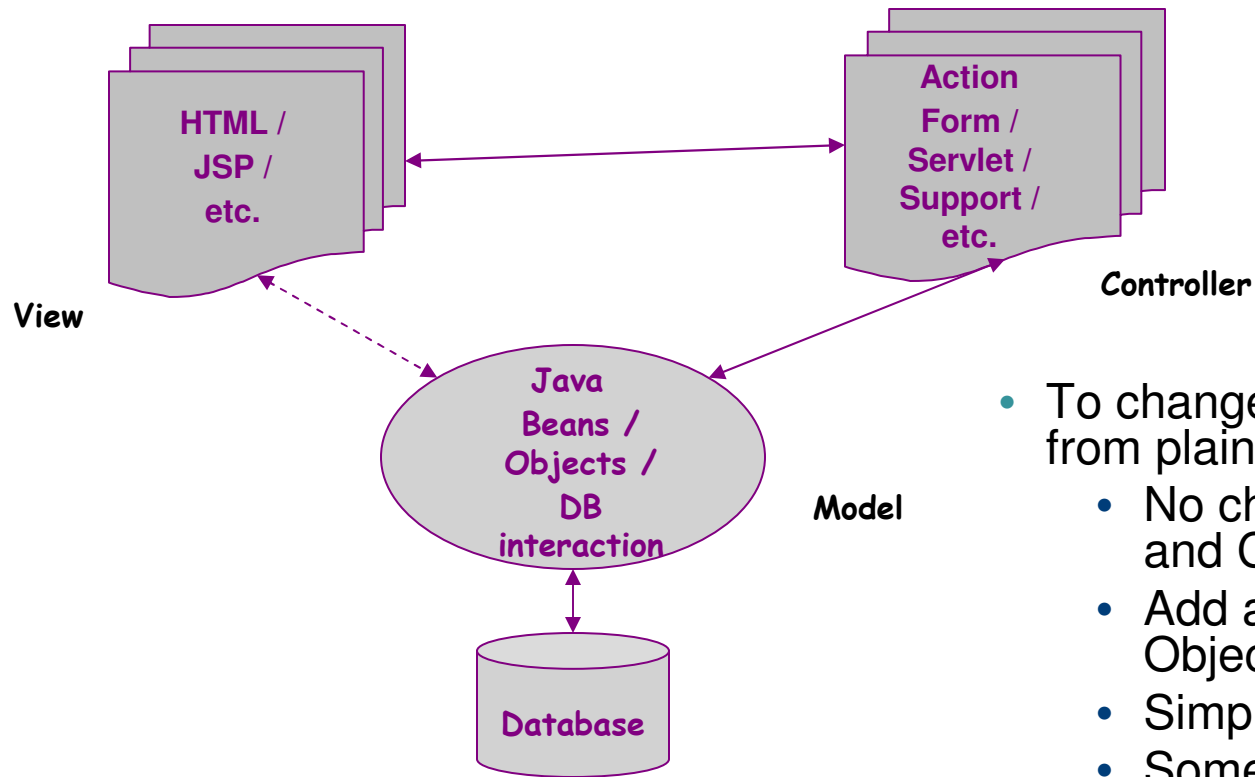
## What is JPA?

- Standard persistence technology
  - JPA 1.0 introduced in Java EE 5
  - JPA 2.0 is part of Java EE 6 standards
  - Object/Relational Mapping
  - Designed for highly distributed applications
    - Especially web-enabled applications
  - Life cycle manageable by application server to increase quality of service
- Simplifies development
  - Provides framework/ tools around providing automatic persistence
  - Objects are mapped to Tables using Metadata
  - Metadata is used to transform data from one representation to another

## What is JPA? (continued)

- JPA specifications made very significant simplifications:
  - Standardizes O/R mapping metadata (not the case for EJB 2.x)
  - Java SE 5.0 Annotations can be used instead of XML deployment descriptor
  - No deploy code implementing abstract classes— Entities are POJOs
  - Application Server is not required
    - The Java Persistence API is available outside of Containers
    - Unit testing greatly simplified
  - Removal of checked exceptions (for instance, remote exceptions)
  - No bean (or business) interface required

# Web Application Example – MVC Framework



- To change existing applications from plain JDBC to use JPA
  - No change required to View and Controller
  - Add annotations to Java Objects
  - Simplify DB interaction code
  - Some configuration changes (for example, persistence.xml)



# Annotate Entity Object Class

Import the javax.persistence packages

```
import javax.persistence.*;
```

Declare the class as an entity

```
@Entity
```

Identify a property as a primary key

```
@Id
```

```
@GeneratedValue(strategy=GenerationType.AUTO)
```

```
private int ID;
```

```
private String name;
```

```
private String industry;
```

```
private String headquarterLocation;
```

```
private String cepClient;
```

```
private String betaClient;
```

```
private String productsInUse;
```

```
private String wasVersions;
```

```
public Client() {
```

```
    ...
```

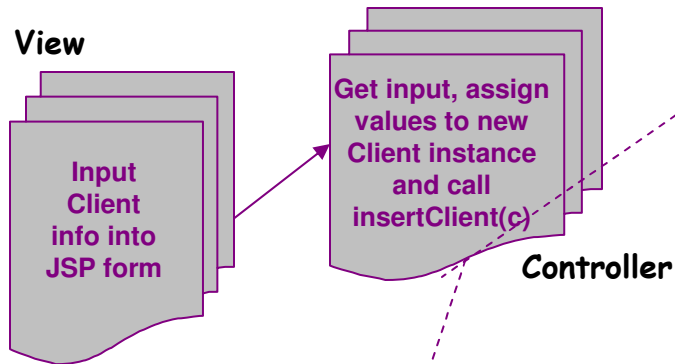
```
}
```

```
//all necessary getter and setter methods
```

```
    ...
```

```
}
```

# JPA compared to Plain JDBC – Insert



**JPA**

```
@PersistenceContext(unitName="test")
EntityManager em;

public void insertClient(Client client){
    doInsert(Client)
} //end insertClient

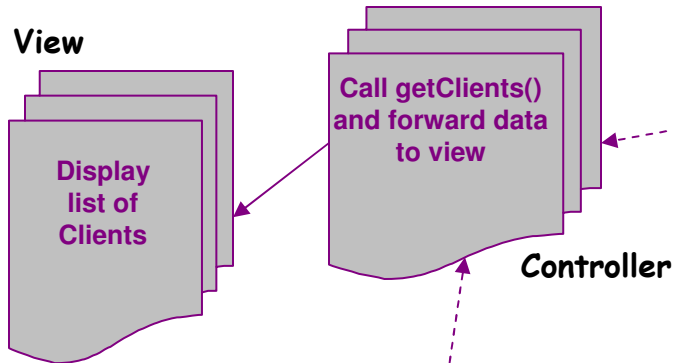
public void doInsert(Object newObject){
    em.persist(newObject);
} //end doInsert
```



```
public void insertClient(Client client){
    String insert = "INSERT INTO clientInfo ("name, industry,"+
        "headquarterLocation, cepClient, betaClient,"+
        "productsInUse, wasVersions) VALUES(' " +
        client.getName()+"', '"+client.getIndustry()+"', '"+
        client.getHeadquarterLocation()+"', '"+
        client.getCepClient()+"', '"+ client.getBetaClient()+"', '"+
        client.getProductsInUse()+"', '"+client.getWasVersions()+"')";
    doInsert(insert);
} //end insertClient

public void doInsert(String insert){
    Statement stmt;
    Connection conn;
    CONSTANTS c = new CONSTANTS();
    try {
        java.lang.Class.forName(c.DRIVER).newInstance();
    } catch (Exception E) {
        System.err.println("Unable to load driver.");
        E.printStackTrace();
    }
    try{
        conn = DriverManager.getConnection(c.DB_CONNECTION);
        stmt = conn.createStatement();
        stmt.executeUpdate(insert);
        stmt.close();
        conn.close();
    } catch (SQLException E) {
        System.out.println("SQLException: " + E.getMessage());
    }
} //end doInsert
```

# JPA compared to Plain JDBC – Retrieve



```

@PersistenceContext (unitName="test")
EntityManager em;

public Client [] getClients(){
    String query = "SELECT c FROM Client c";
    Query q = em.createQuery(query, Client.class);
    Client {} cArr = q.getResultList().toArray();
    return cArr;
} //end getClients
    
```

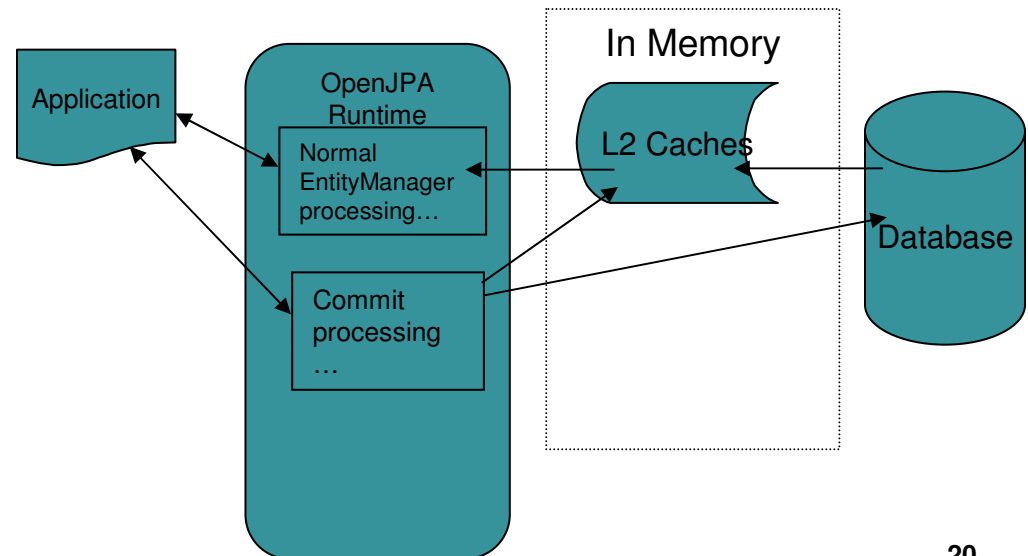


```

public Client [] getClients(){
    String query = "SELECT * FROM clientInfo;";
    LinkedList list = doQuery(query);
    int size = list.size();
    Client [] cArr = new Client[size];
    for(int i=0; i<size; i++){
        cArr[i] = new Client();
        String[] item = (String[]) list.get(i);
        cArr[i].setID(item[0]);
        cArr[i].setName(item[1]);
        cArr[i].setIndustry(item[2]);
        cArr[i].setHeadquarterLocation(item[3]);
        cArr[i].setCepClient(item[4]);
        cArr[i].setBetaClient(item[5]);
        cArr[i].setProductsInUse(item[6]);
        cArr[i].setWasVersions(item[7]);
    } //end for i
    return cArr;
} //end getClients
    
```

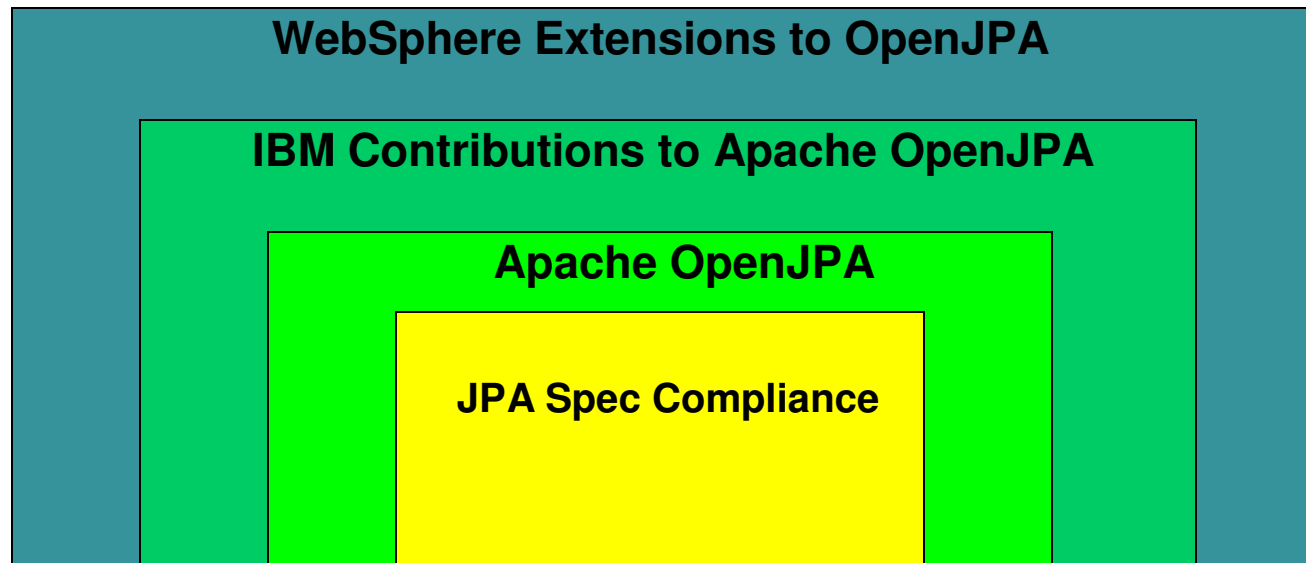
# Advantages to using OpenJPA

- Object level programming model
  - More natural for programmers
- Container “free”
  - No requirement on an application server
  - JEE, JSE, and now OSGi
- Common programming model across databases
  - No database-specific SQL or JDBC
- Ability to Cache database records in memory
  - Fewer trips to the database



# WebSphere's JPA Architecture

- WebSphere's JPA Solution
  - Spec Compliant
  - Feature Rich
  - Extensible
  - Competitive with Hibernate and EclipseLink
  - Production Ready



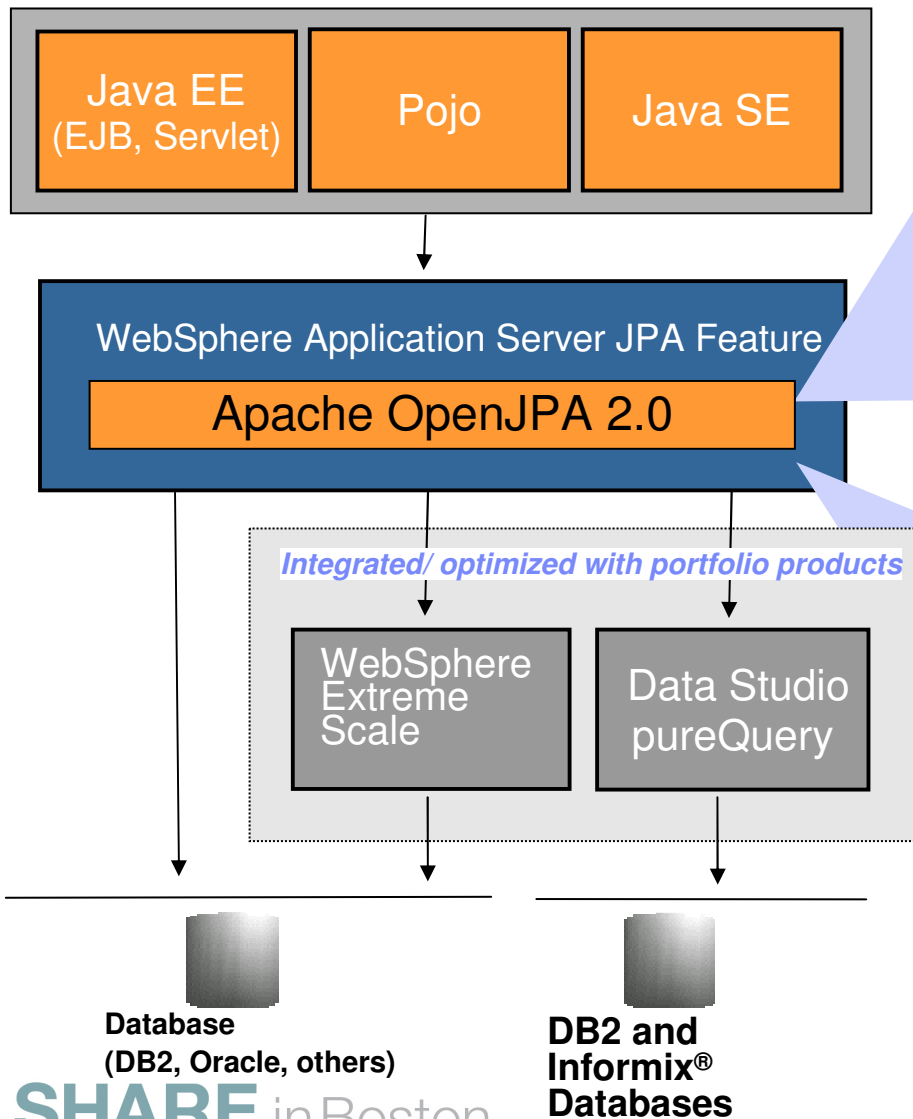

# JPA 2.0 Feature Pack

# The WebSphere Application Server V7 Feature Pack for OSGi Applications and Java Persistence API 2.0



- Based on Apache OpenJPA, a leading open source Java persistence framework.
- Provides the Apache OpenJPA 2.0 implementation with IBM enhancements to benefit integration with WebSphere Application Server.
- The Apache OpenJPA 2.0 implementation includes improvements and benefits over previous releases and even beyond the JPA 2.0 specification.

# WebSphere Application Server JPA 2.0 feature Architecture



Apache OpenJPA 2.0 provides:

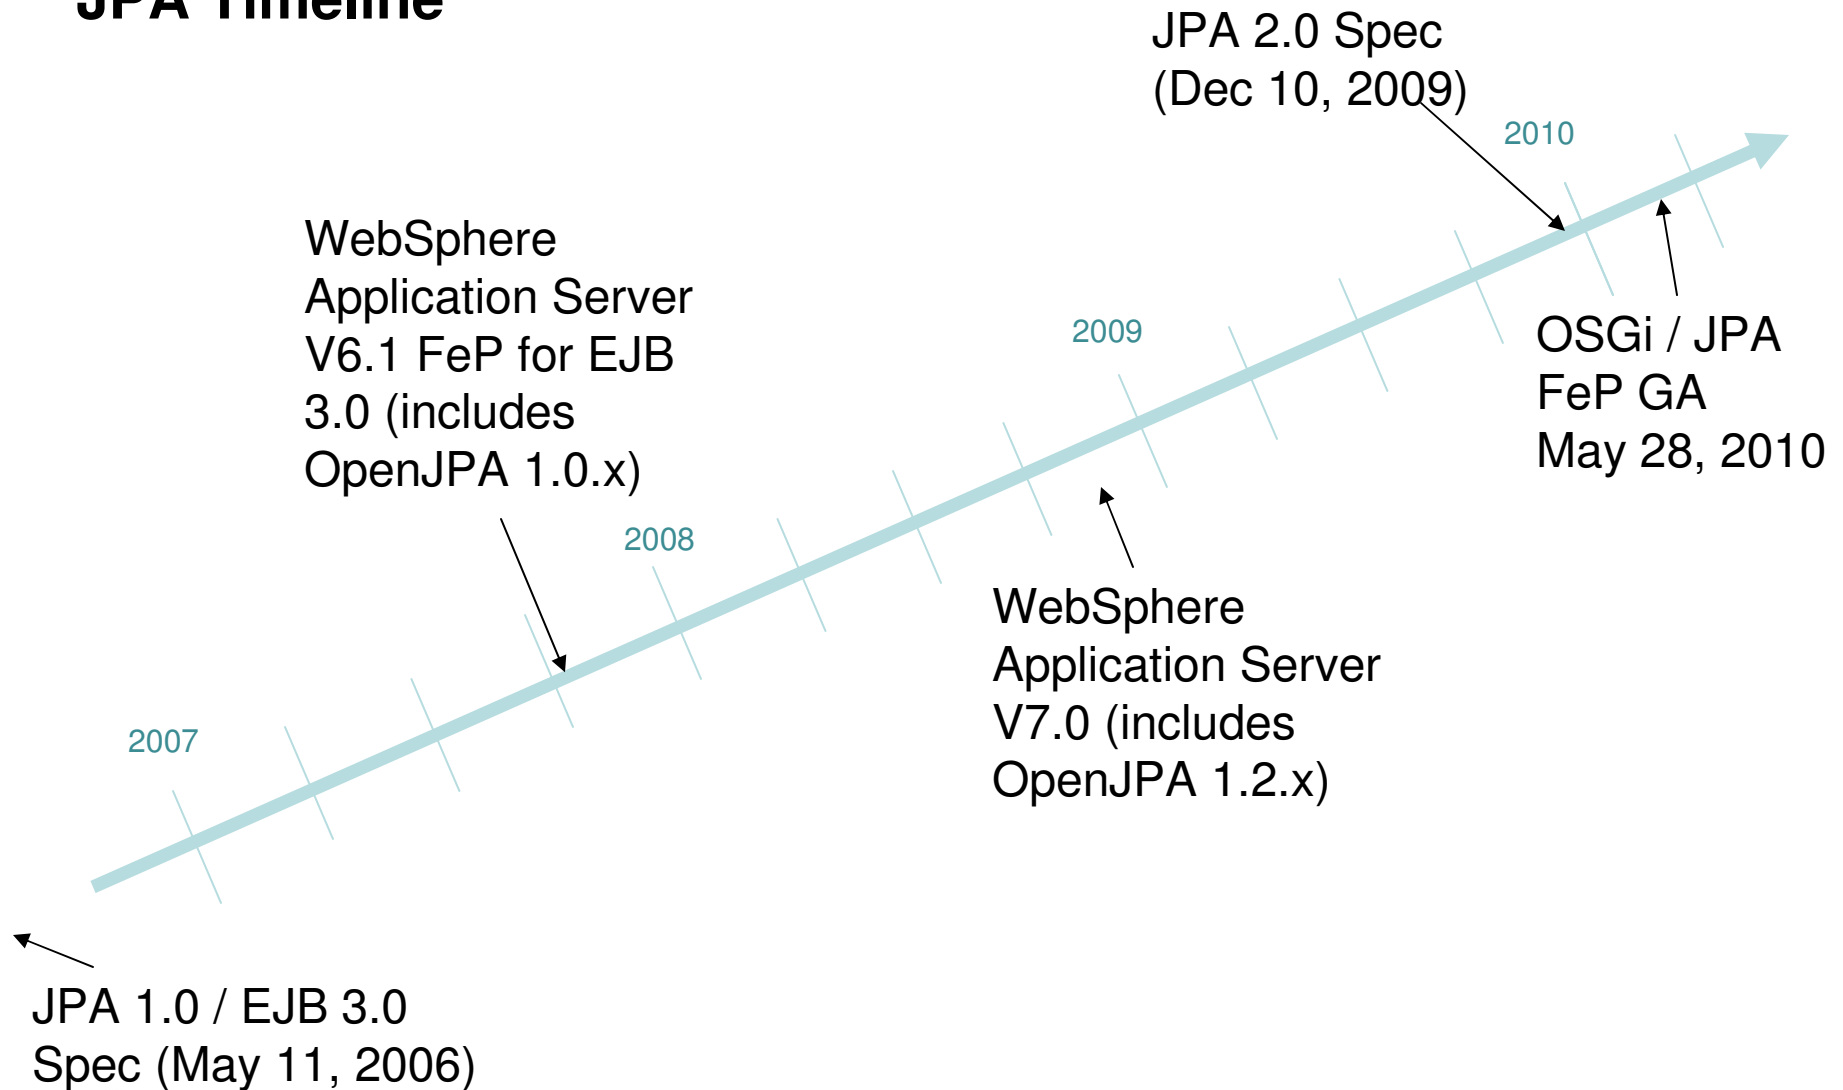
- JPA 2.0 standards currency: **JPA 2.0 (JSR 317)** includes important improvements on the data persistence and relational database management capabilities. Early delivery of part of the JEE 6 standard.
- WebSphere Application Server Integration: **Enhanced management, debugging, monitoring, and security** for WAS.
- WebSphere Extreme Scale Integration Enhancements: **Features that enable JPA to use WebSphere eXtreme Scale as a level2 cache in order to improve data access performance.**

*WebSphere Application Server Feature Pack for OSGi Applications and Java Persistence API 2.0 JPA feature adds:*

- Additional WebSphere Application Server Integration: **Install, Extended Debugging**
- Data Studio pureQuery Runtime optimization: **Performance and monitoring improvements for database access**



# WebSphere Application Server V7 Feature Pack for JPA Timeline



# JPA 2.0 Highlights

- O/R Mapping and Domain Modeling
  - Embeddables
  - Access Types
  - Enhanced Map Collections
  - Derived Identities
  - JPQL Updates
- Runtime API Updates
  - EntityManagerFactory API, EntityManager API, Query API, Query Result API
- Metamodel and Criteria API
- Bean Validation
  - Provides entity validation based on JSR-303 semantics



WebSphere Application Server

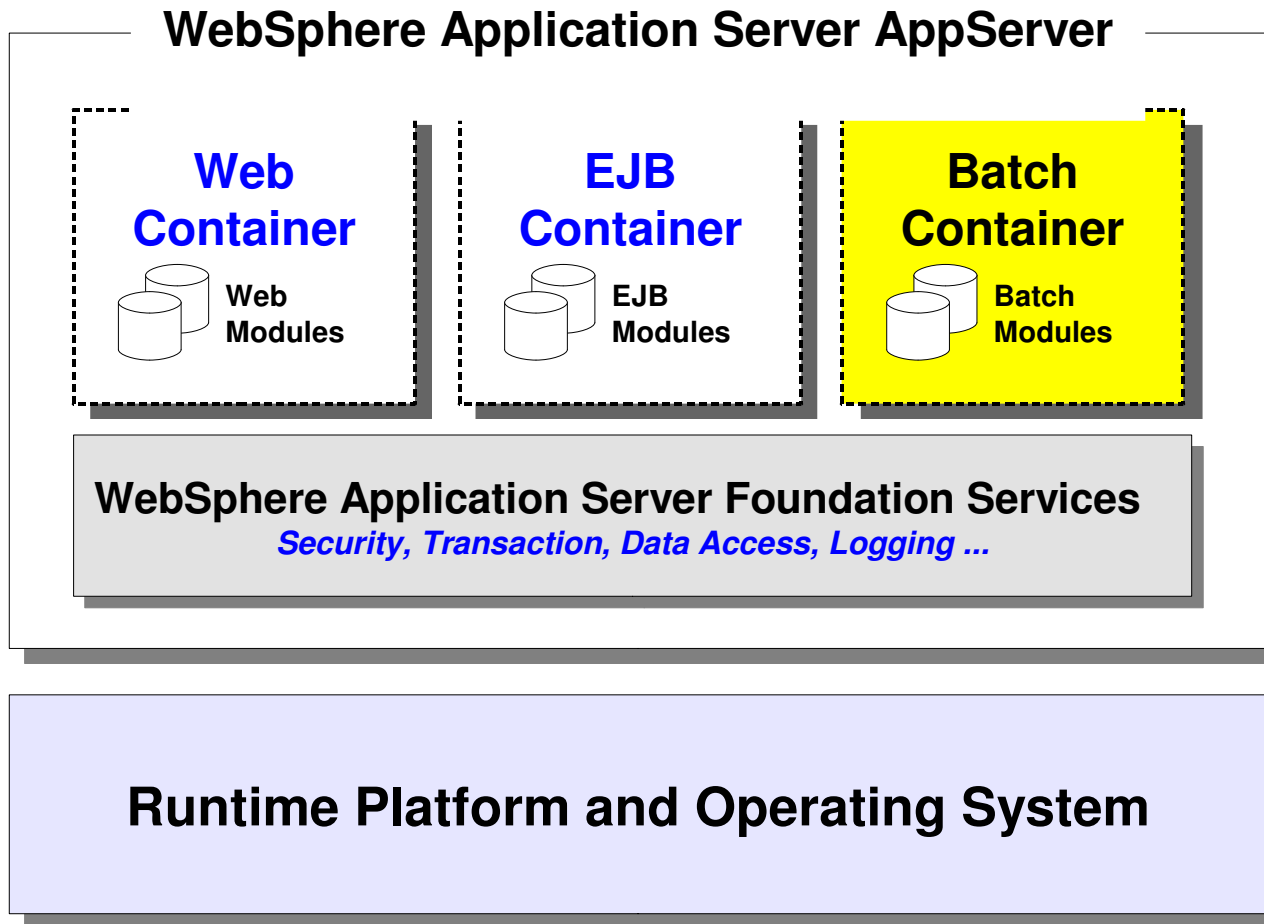
# Feature Pack for Modern Batch

Currently in Open Beta

# What is the "Feature Pack for Modern Batch?"



It is a batch execution framework within the WebSphere Application Server runtime platform:



Feature Pack adds function to an existing WebSphere Application Server runtime environment

Available for all platforms

Currently in Open Beta

# What's the Objective?

In a picture and a few words:



**Allows you to maintain a focus on your core business objectives while leveraging Java as a batch execution language**



**Helps you avoid writing expensive and hard-to-maintain custom batch middleware functionality**



**Makes you a hero in the eyes of management, who's focused on the business imperatives**

**Funny pictures ... *serious point*. Countless hours and dollars have been spent by companies pursuing custom middleware solutions. It's costly, time consuming and may lead to disparate islands of Java batch processes**

# What Does it Provide?

Here's a summary of the key features the Feature Pack provides:

- **Batch container environment**  
Provides the structure and support function your Java batch application uses
- **Job scheduler and dispatcher function**  
Provides a key separate of application code from batch job, and controls the dispatching of the submitted job to a batch container hosting the application
- **Declarative job control file (xJCL)**  
Very much like regular JCL except it uses JCL.
- **Development class libraries**  
So batch applications may be written that run within this batch container
- **Batch Data Stream (BDS)**  
Function that abstracts data into easily accessible record format, taking the burden off you with respect to stream handling
- **Conditional multi-step job support**  
A combination of declaration in xJCL with step execution control
- **Checkpoint processing leveraging WAS transaction manager**  
Record or time, specified at the job step level.
- **Results and return code coordination**  
A consistent results framework within the batch container for each job

# Relationship to WebSphere Compute Grid

It's a matter of degree of functional sophistication. The Feature Pack is a subset of the much more powerful WebSphere Compute Grid product:

## Previous Chart

- **Batch container environment**  
Provides the structure and support function your Java batch application uses
- **Job scheduler and dispatcher function**  
Provides a key separate of application code from batch job, and controls the dispatching of the submitted job to a batch container hosting the application
- **Declarative job control file (xJCL)**  
Very much like regular JCL except it uses JCL.
- **Development class libraries**  
So batch applications may be written that run within this batch container
- **Batch Data Stream (BDS)**  
Function that abstracts data into easily accessible record format, taking the burden off you with respect to stream handling
- **Conditional multi-step job support**  
A combination of declaration in xJCL with step execution control
- **Checkpoint processing leveraging WAS transaction manager**  
Record or time, specified at the job step level.
- **Results and return code coordination**  
A consistent results framework within the batch container for each job

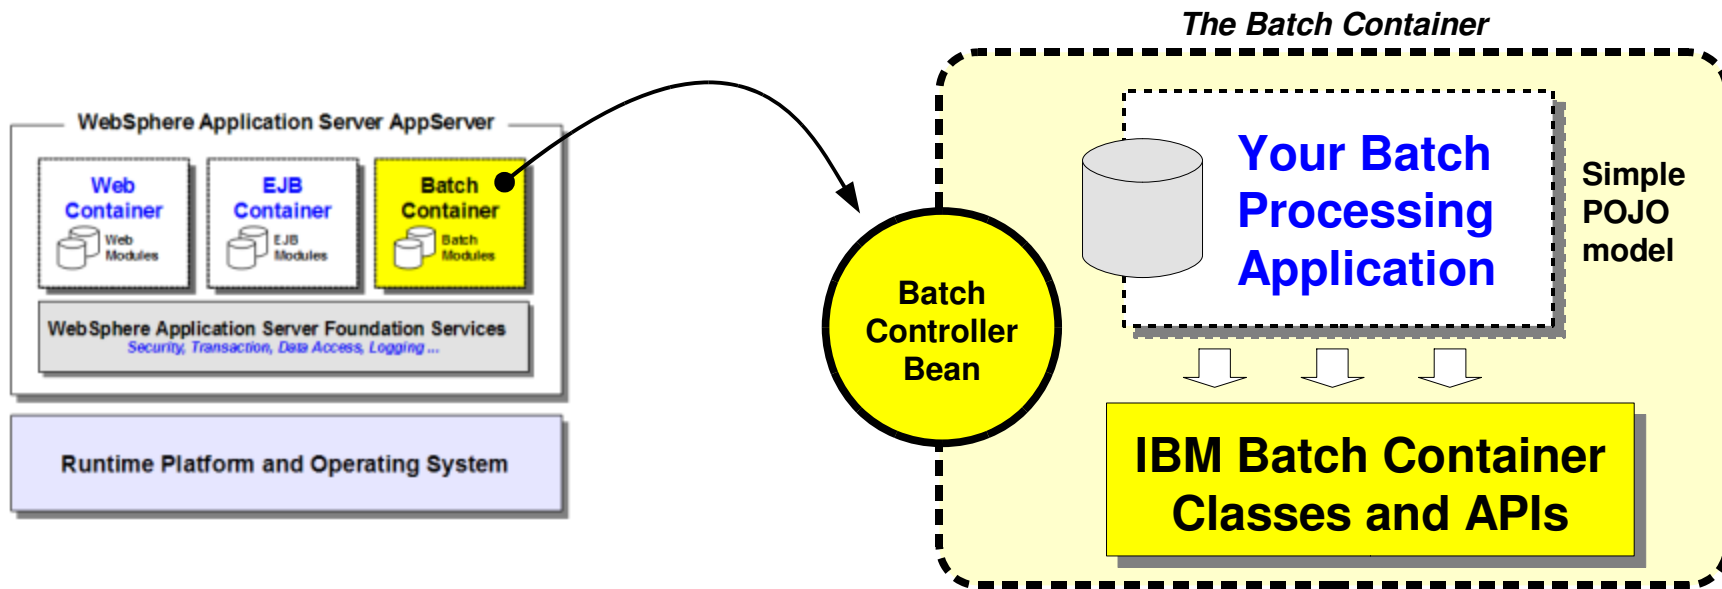


- **Calendar and clock scheduling of jobs from repository of xJCL**
- **Out of the box integration with enterprise schedulers**
- **Usage reporting with SMF 120.20 records (z/OS only)**
- **WLM transaction classification *by job* (z/OS only)**
- **Application quiesce and update**
- **Job submission pacing and job execution throttling**
- **Parallel job management and dispatching**

**Feature Pack is designed to introduce you to WAS container processing of Java batch  
WebSphere Compute Grid completes the picture with enterprise functionality**

# How Does it Provide This Function?

Batch by its nature has an long running execution profile. Therefore it can't be run under the traditional request / response model. So it uses asynchronous beans:



**This is an asynchronous bean**

**Your batch application runs under the control of this bean**

**You can think of this as a container-managed thread**

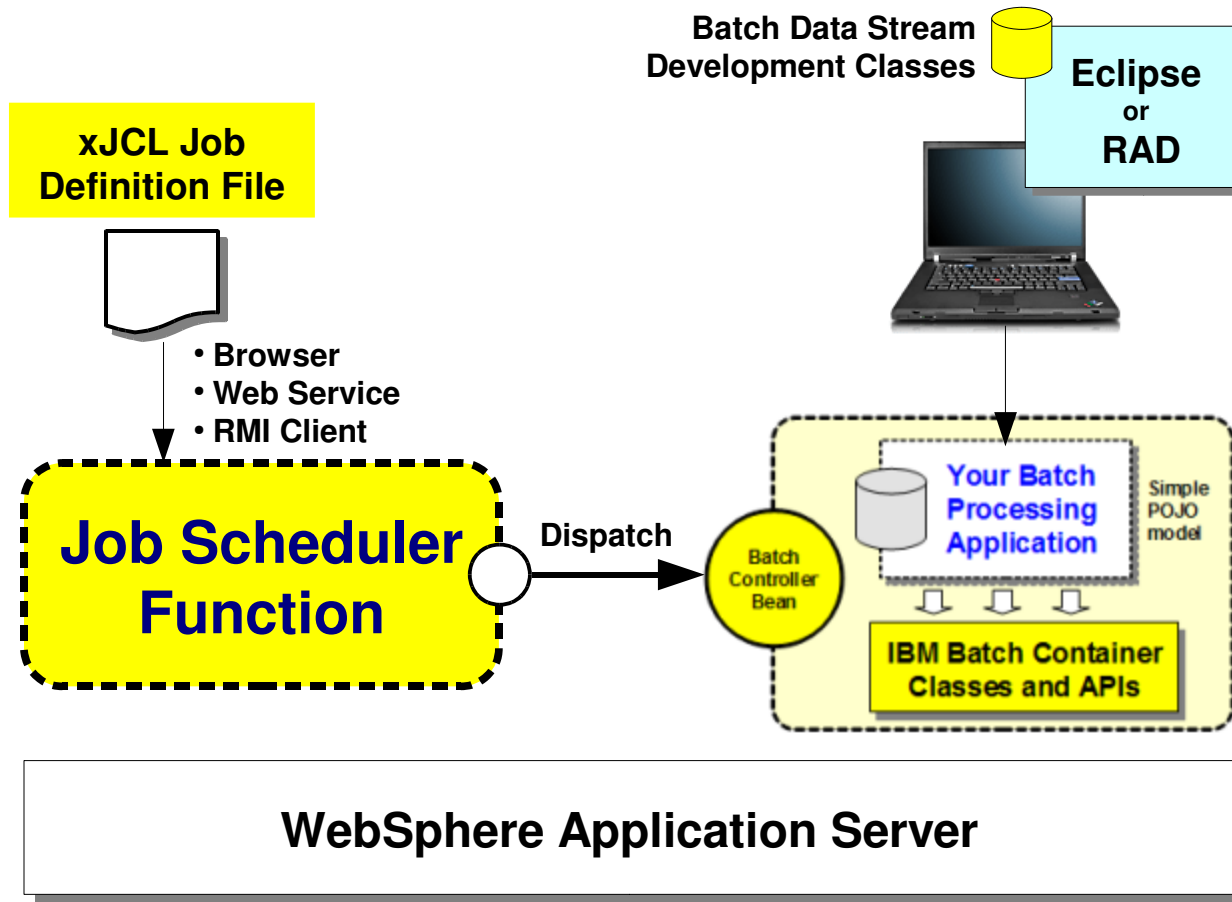
**It processes the job definition and carries it from start to finish**



# Applications, Jobs and Job Control Definition



A key concept to appreciate the differentiation of this over other Java batch models:



- **Application** deployed into AppServer
- **Job** is an invocation of that application
- The job execution characteristics defined in job control file (**xJCL**)
- xJCL submitted to scheduler function
- Scheduler determines location of named application
- Scheduler dispatches job to the batch end point with application
- Application runs according to xJCL definition

# The Job Control Definition File -- "xJCL"

Syntax different than traditional JCL, but the concepts are very similar:

```
<?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>

  <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>
</job>
```

Roughly analogous  
to the JOB card

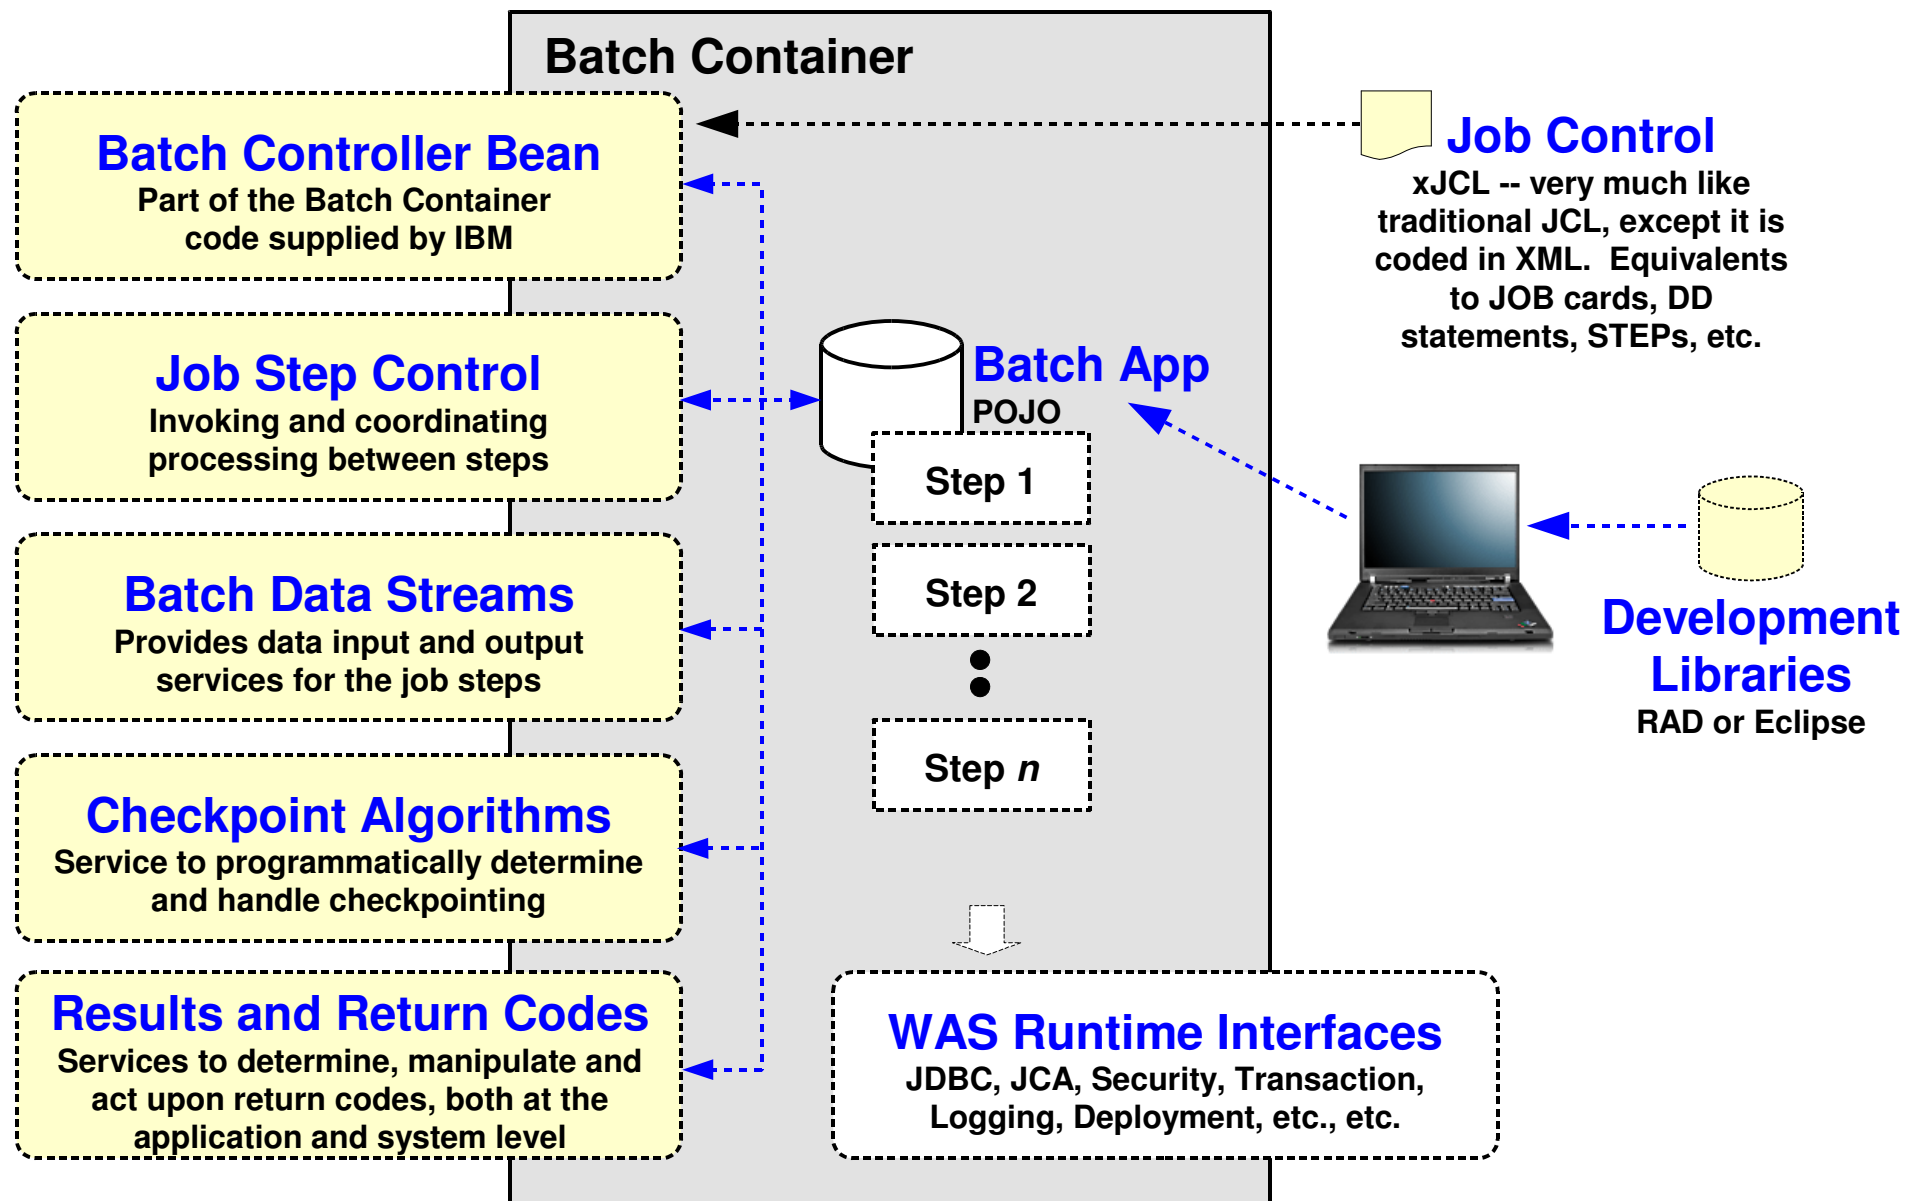
A job step

Like the EXEC PGM=  
statement in JCL

Similar to DD  
statements

# Batch Programming Framework

Provides key functional structures for use by your batch application:

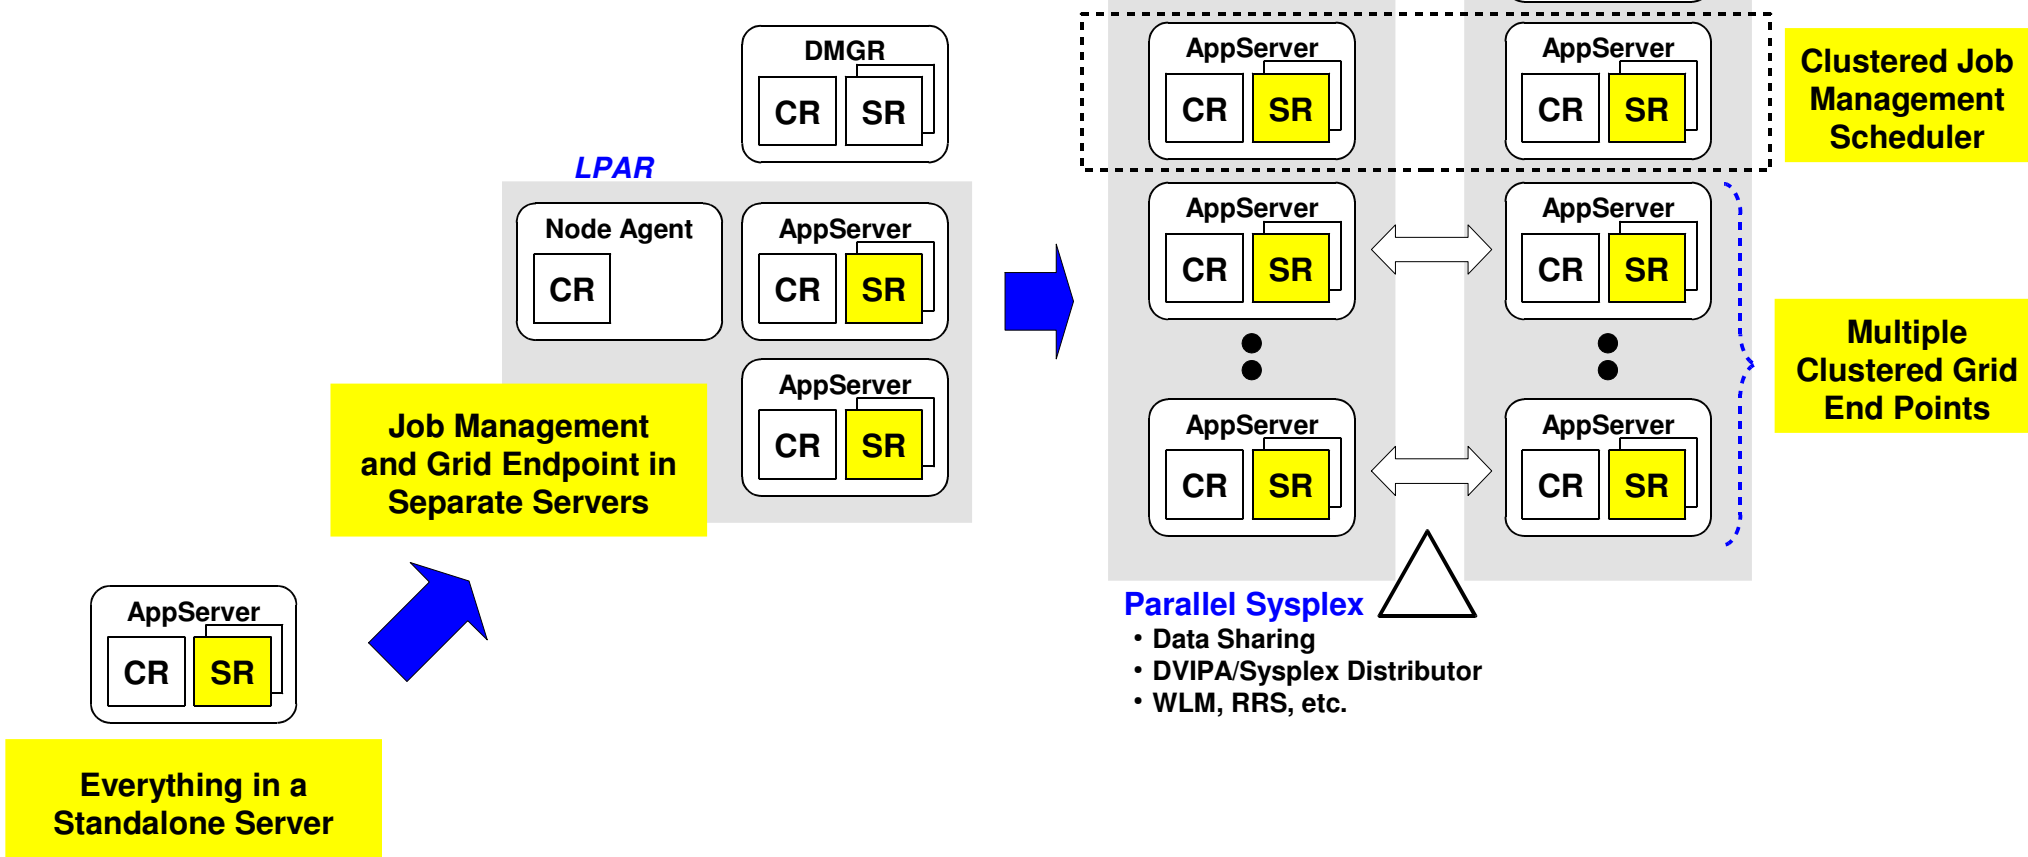


# WAS z/OS Topologies for Batch Feature Pack

As you might expect from WAS z/OS:

Since this is built on top the WAS z/OS design, it's capable of taking on the configuration options WAS z/OS itself has

Similar topologies possible with Distributed WAS ... minus the features of z/OS and Parallel Sysplex



# Implementing the Feature Pack

A relatively simple process:

## Install the Product Files

- Feature Pack for Modern Batch -- typical WAS feature pack, installs under the "Optional Materials" mount point
- WebSphere Compute Grid -- separate FMID, installs under /usr/lpp/zWebSphereXD mount point (default)

## Augment the WAS z/OS Runtime Nodes

- A WebSphere Customization Tools (WCT) extension file is supplied
- Generate augmentation job to indicate what node to augment and where batch installation is located
- Run the job -- one job for the DMGR and one for each node where batch will run

## Create JDBC Provider, Data Sources and Database

- JDBC provider and data source is standard WAS configuration task
- Product comes with DDL to create database in DB2 z/OS, DB2 LUW, Oracle or Derby

## Configure the Job Scheduler and the Grid End Points

- An almost entirely automated process. A few mouse-clicks and it does the rest

## Validation

- "Test Connection" button on the JDBC data sources
- Several sample applications are supplied

**Draft step-by-step Techdoc available on the Beta forum**

# Open Beta Website

<https://www14.software.ibm.com/iwm/web/cc>

</earlyprograms/websphere/iwsasfpmdb/index.shtml>

Software > Early programs > WebSphere >

## IBM WebSphere Application Server V7 Feature Pack for Modern Batch Beta

Overview | Education | **Download** | Library | Support | Live Sessions

IBM Beta Program  
 IBM WebSphere Application Server V7 Feature Pack for Modern Batch Beta  
 The IBM WebSphere Application Server V7 Feature Pack for Modern Batch provides support

**Download Code**

IBM WebSphere Application Server V7 Feature Pack for Modern Batch Beta

Search for  within IBM WebSphere Applicat...

This forum is for technical discussions supporting the IBM WebSphere Application Server V7 Feature Pack for Modern Batch Beta Program. Visit [Beta Program web page](#) for download of the code and various Beta Program documents.

By using this forum, you agree to abide by the [forum etiquette](#) and that you understand the [Terms of Use](#) governing this web site.

[Post New Thread](#) [Watch Forum](#) [Watch List](#) [Forum settings](#) [Help](#)

Messages: 14 - Threads: 8 - Filter: [All Threads](#)

- | Threads  |  |
|--|--|
| <a href="#">Where can I find the sample applications and sample job definitions?</a>     |  |
| <a href="#">DRAPT of WAS z/OS Techdoc for Batch Feature Pack</a>                         |  |
| <a href="#">CWL885240E not authorized for job status when no app security</a>            |  |
| <a href="#">Augmenting this Feature Pack</a>   |  |
| <a href="#">What are the precess/restrictions for the feature pack for modern batch?</a> |  |
| <a href="#">How to install this Feature Pack on Windows/Linux?</a>                       |  |
| <a href="#">Error while trying to download Modern Batch Beta code</a>                    |  |
| <a href="#">Welcome to IBM WebSphere Application Server V7 Feature Pack</a>              |  |

**Support Forum**

[SMP/E install for z/OS](#)  
 This option is for z/OS only.

For this option you will download the product binaries to your system and install using SMP/E.

For installation and configuration instructions, refer to the Getting Started Guide, available from the [Library page](#).