


---

# IBM WebSphere Extended Deployment Compute Grid V8

## Overview



© 2011 IBM Corporation

This module provides an overview of the enhancements included in IBM WebSphere Extended Deployment (XD) Compute Grid V8.

## Agenda

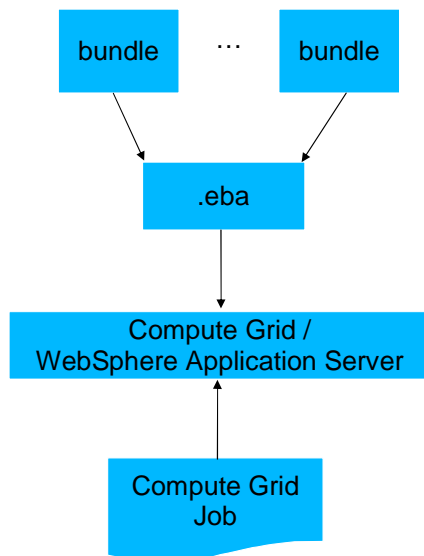
- Programming model enhancements
- Job definition enhancements
- Operational enhancements

This presentation covers the programming, job definition, and operational enhancements included in WebSphere XD Compute Grid V8.

## ***Programming model enhancements***

This section describes the programming model enhancements included in WebSphere XD Compute Grid V8.

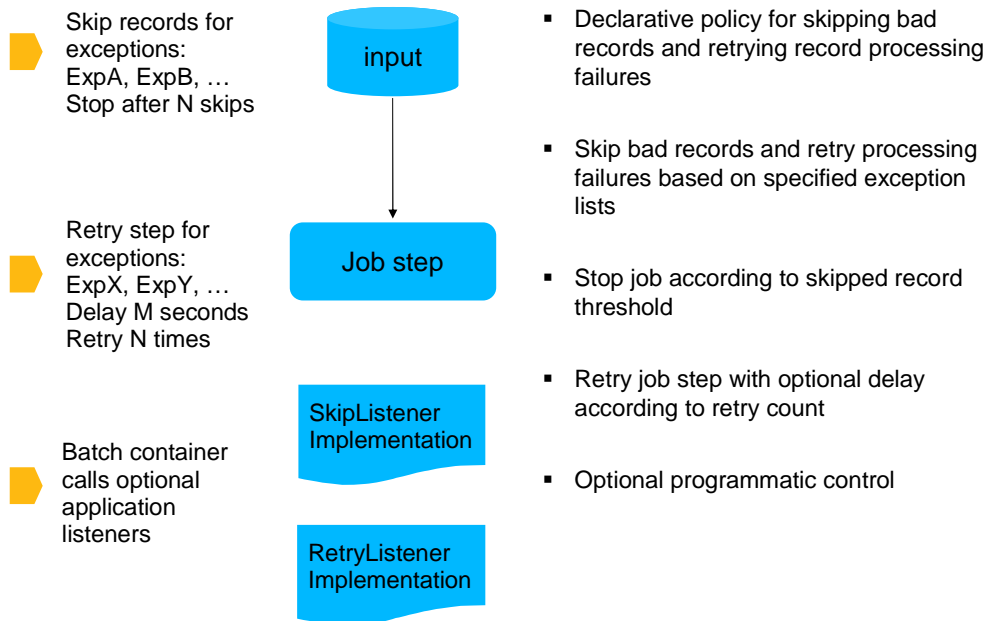
## OSGi batch applications



- Enables use of OSGi for batch application development
- Full batch programming model available to OSGi framework
- Supports standard and blueprint bundles
- Enterprise Bundle Archive deployment

WebSphere XD Compute Grid V8 enables you to re-use your existing blueprint OSGi applications to run batch workloads. Batch artifacts are exported as blueprint services which are invoked by the Compute Grid batch container during the job lifecycle.

## Record processing policy



5

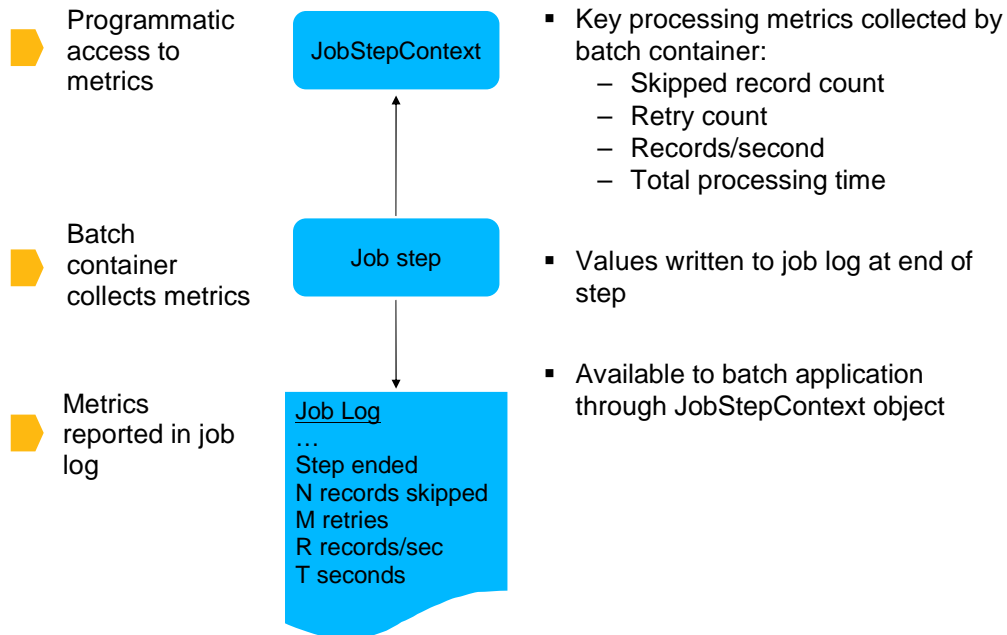
Overview

© 2011 IBM Corporation

WebSphere XD Compute Grid V8 includes two different declarative record processing policies, skip records and step retry. The skip records policy covers the input stream, and the step retry policy covers record processing and the output stream.

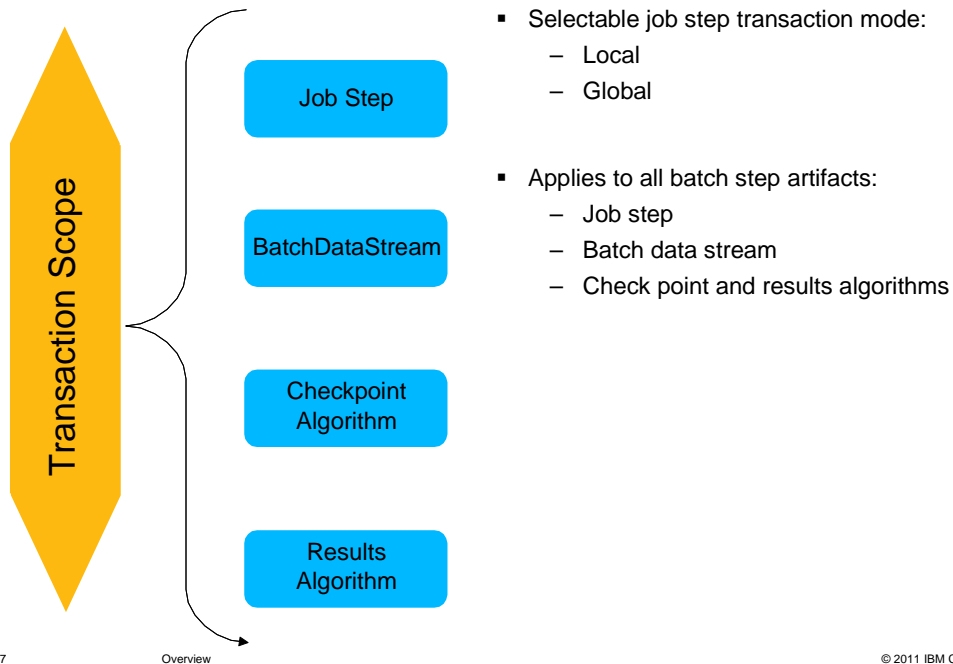
The step retry policy is NOT related to skipped records. It will automatically retry a failed step since its last checkpoint.

## Record metrics



The record metrics feature provides convenient statistics about the records processed by each job step. Key statistics include the skipped record count, retry count, records processed per second, and total processing time. The values are written to the job log at the end of a step, and are available to a batch application through the JobStepContext object.

## Configurable transaction model



WebSphere XD Compute Grid V8 adds support for a local transaction mode of operation. In this mode, no global transaction is started; the connection object is shared between the container and the user code via the JobStepContext object. Local transaction mode significantly improves performance by reducing the overhead required to manage global transactions. This mode is especially useful if your data is located on the same database system as the Compute Grid tables. In addition, large cursors are less likely to be closed when transactions are committed.

## Batch data stream timeouts

xJCL:

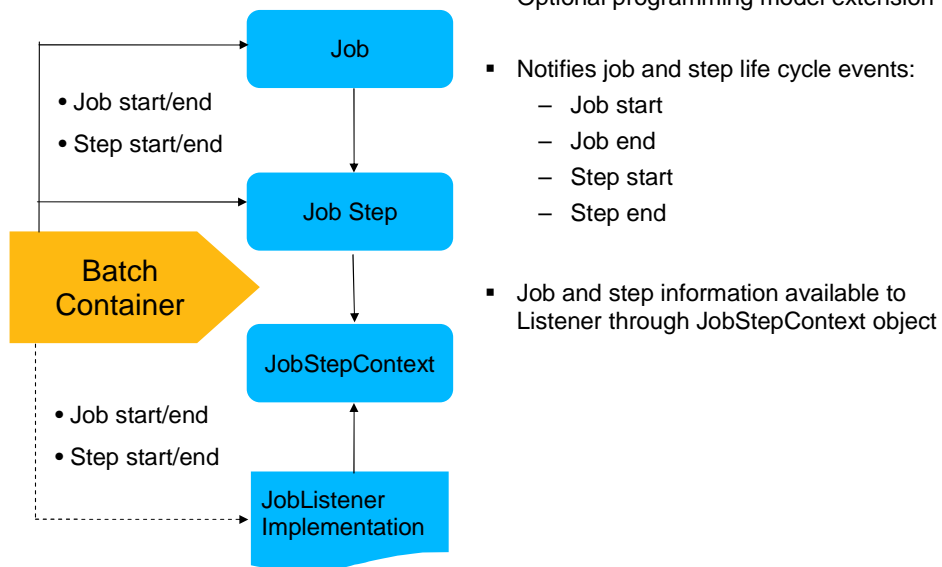
```
<job name=... >
  <step name=...>
    <batch-data-streams timeout=N>
      <bds>...</bds>
      <bds>...</bds>
      <bds>...</bds>
    </batch-data-streams>
  </step>
</job>
```

- Problem
  - Job steps need short transaction timeouts
  - Batch data streams (large results sets) need long timeouts
- Solution
  - Configurable batch data stream timeout
  - Can be specified per job step

WebSphere XD Compute Grid V8 allows you to configure a batch data stream timeout for each job step; as a result, you can now define short timeouts for transactions, and long timeouts for BatchDataStreams.



## Job and step listener



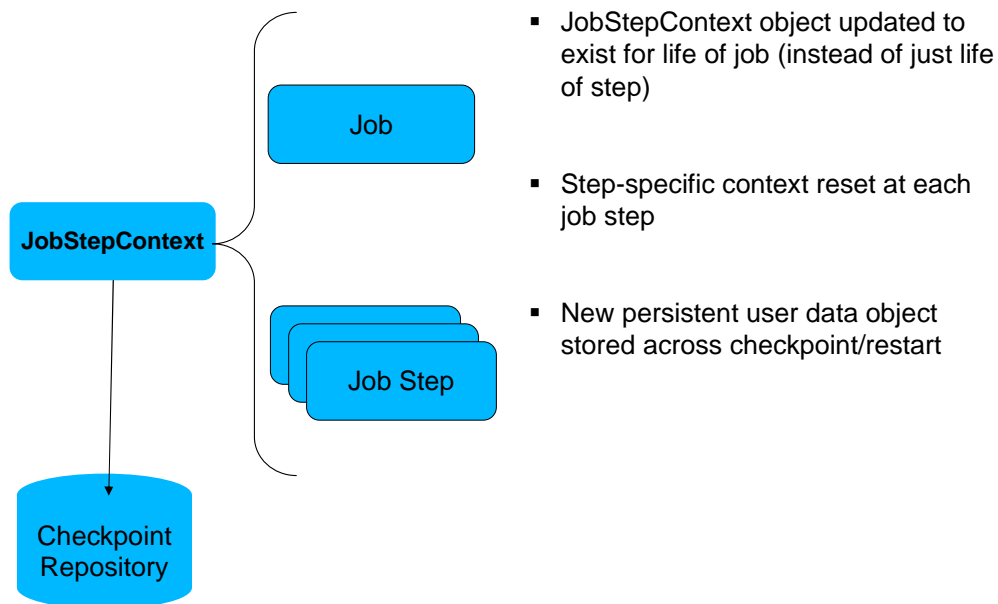
9

Overview

© 2011 IBM Corporation

The job and step listener allows application developers to add cross-cutting initialization and clean up routines to jobs and steps. The listener implementation is specified via a new job-level xJCL element. The batch container invokes the listener at specific events: before job start, after job start, before step start, and before step end. Data can be accessed through the JobStepContext object.

## Persistent job context



10

Overview

© 2011 IBM Corporation

In WebSphere XD Compute Grid V6.1.1, job context information for the top-level job is persisted to a database, and step context information is transient. In Compute Grid V8:

-JobStepContext information is available to all job types.

-The JobStepContext object exists for the life of the job, as opposed to the life of a step.

-Step-specific context information is reset at each job step.


-New step-level user data is persisted to a database at each checkpoint so that it is available at job restart.

## COBOL support

Value

**Reuse existing COBOL in new Java Batch Applications!**

Batch Application



- Call standard COBOL<sup>1</sup> modules from Java on same thread in same process
- Java and COBOL run in same transaction scope
- DB2 connections managed by WebSphere Application Server can be shared with COBOL
- COBOL working storage isolation per job step
- RAD tooling for Java call stub generation

<sup>1</sup> z/OS only

11

Overview

© 2011 IBM Corporation

COBOL has been the prevalent language to run batch-style workloads in the past, and even today there is a large existing base of COBOL code. On the z/OS platform, WebSphere XD Compute Grid V8 adds support to easily call into existing COBOL assets to run modern batch workloads.

## ***Job definition enhancements***


This section describes the job definition enhancements included in WebSphere XD Compute Grid V8.

## Multi-threading

### xJCL:

```
<job name=... >  
  <run instances="multiple"  
    jvm="single" />  
  <step name=... >  
    ...  
  </step>  
</job>
```

### Runtime:



13

Overview

© 2011 IBM Corporation

- Option to run parallel job on multiple threads
- Parallel Job Manager local optimization
- Alternative to running parallel job across multiple JVMs
- Optimizes shorter running subjobs

WebSphere XD Compute Grid V8 adds support to run parallel jobs and their sub-jobs in the same JVM if desired. This capability is especially useful when sub-jobs are typically very short in duration and the overhead of distributing the sub-jobs across servers is much larger compared to the actual work to be completed.

## Parallel steps

### xJCL:

```

<job name=... >
  <step name=... >
    <run instances="multiple"
      jvm="single" />
    ...
  </step>
  <step name=... >
    <run instances="multiple"
      jvm="multiple" />
    ...
  </step>
</job>

```

- Ability to specify parallelization configuration per job step
- Supports any step type
- Step might be multi-threaded or multi-process (i.e. multiple JVMs)

Multi threaded step

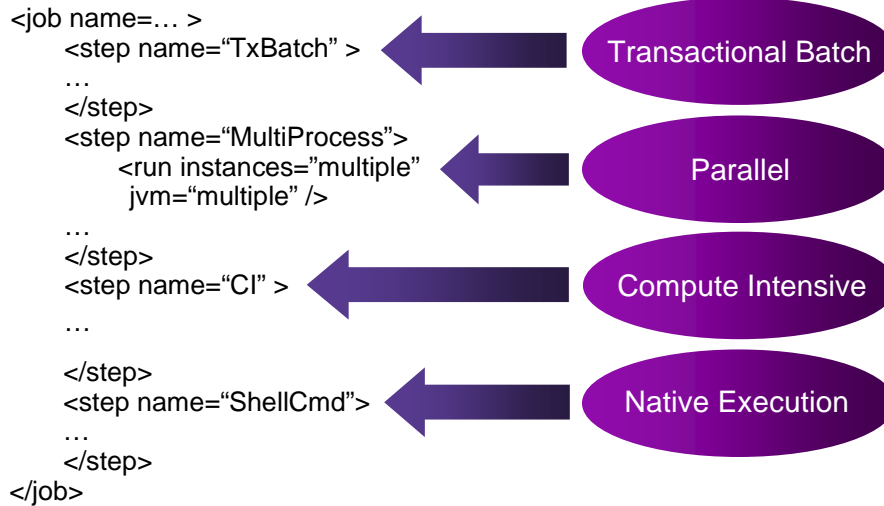
Multi process step

WebSphere XD Compute Grid V8 allows you to run some steps of a job as parallel jobs themselves. This approach is different from running two steps concurrently (a capability that is not currently supported). A parallel step can be thought of as a top level job, with the sub-job xJCL generated using the step's own xJCL snippet.

## Parallel steps

xJCL:

- Ability to mix various step types in the same job



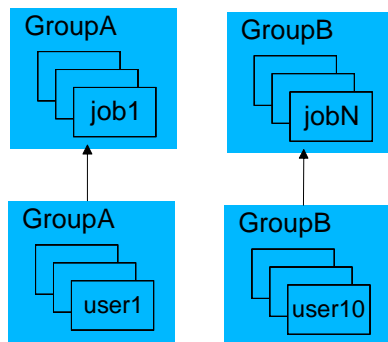
WebSphere XD Compute Grid V8 allows a job to contain different types of steps. For example, a job can include a mixture of batch, compute intensive, and native job steps. This capability adds immense flexibility to the batch programming model, as different parts of a job can be run based their nature, whether transactional or not.

## ***Operational enhancements***

This section describes the operational enhancements included in WebSphere XD Compute Grid V8.



## Group security



- User in GroupA can manage jobs that belong to GroupA
- User in GroupB can manage jobs that belong to GroupB

- Group-level security for jobs
  - Leveraging WebSphere security functions
  - Enable via custom properties
  - Configure to work with or without security roles
- Users have access to jobs based on group membership
- Job management console views are customized based group security settings


In previous versions of WebSphere XD Compute Grid, access to job management tasks and information is controlled via role-based authentication, where each user must be assigned the Iradmin, Irsubmitter, or Irmonitor role. This model is also supported in Compute Grid V8. In an effort to improve flexibility and streamline user management tasks, Compute Grid V8 introduces a new security model that enables a group of users to operate on a common subset of jobs. In the graphic, a user from GroupA can manage jobs that belong to GroupA, and a user from GroupB can manage jobs that belong to GroupB. If a user belongs to both GroupA and GroupB, the user can manage jobs from both groups.

Group security can be configured in combination with role-based security. When this approach is used, a user can perform a job-related action if and only if user and job are members of the same group, and the user's role permits the job action.

The default security model continues to be role-based. To change the security model, you must define several custom properties at the level of the job scheduler.

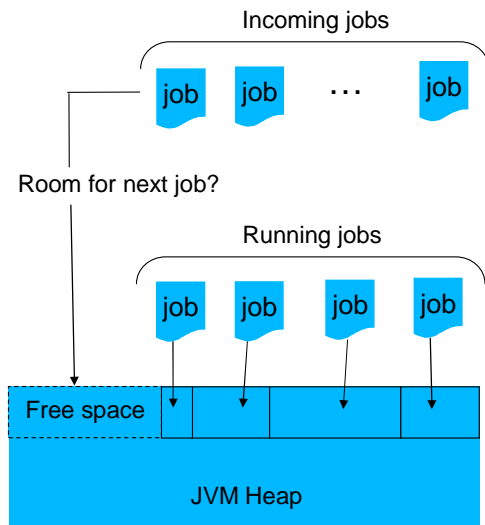
## Group security (continued)

### Job management console

Select	Job ID	Submitter	Last Update	State 	Node	Application Server	Group
<input type="checkbox"/>	<a href="#">Eche:00035</a>	vignola	2010-02-04 11:34:53.953	Ended	VIGNOLA-T60Node01	server1	JSYSDFLT
<input type="checkbox"/>	<a href="#">ParallelJobManager:00083</a>	vignola	2010-05-14 01:06:32.890	Restartable	VIGNOLA-T60Node01	server1	JSYSDFLT
<input type="checkbox"/>	<a href="#">ParallelJobManager:00084</a>	vignola	2010-05-14 01:09:37.593	Restartable	VIGNOLA-T60Node01	server1	JSYSDFLT
<input type="checkbox"/>	<a href="#">ParallelJobManager:00085</a>	vignola	2010-05-14 01:19:15.953	Restartable	VIGNOLA-T60Node01	server1	JSYSDFLT

When group-based security is enabled, the job management console includes the group associated with each job. It is also possible to filter this view by group name.

## Memory overload protection



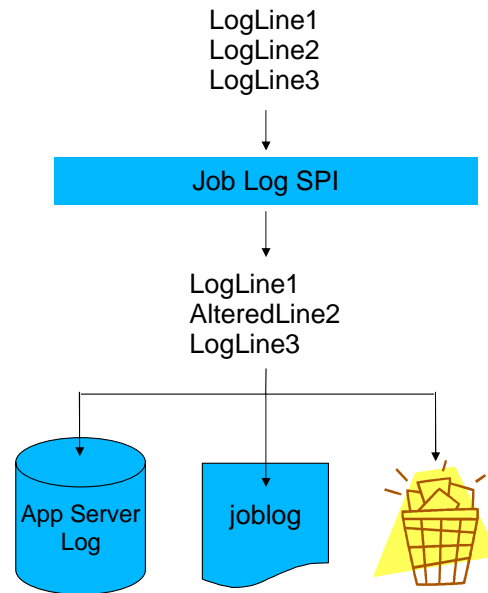
- Protection against over-scheduling jobs to an application server
- Batch container monitors job memory demand against available JVM heap space
- Prevents Java OutOfMemoryError
- Automatic real time job memory estimation with declarative xJCL override

In WebSphere XD Compute Grid V8, the memory overload protection feature prevents the mix of jobs on an endpoint server from causing a Java heap out of memory condition. A new xJCL attribute, `memory`, allows you to declare the amount of memory required by a job.

The batch container monitors job memory requirements against its available heap space. If a job enters the batch container, and the memory requirements of the job exceed the available free space, the job is placed in wait state. As other jobs exit the BatchContainer and resources are released, the held job is again evaluated against the available free space.

This feature does not affect which endpoint server the job scheduler selects to process the job. It simply ensures that the batch container on the selected endpoint does not run into an out of memory condition because an excessive number of memory-intensive jobs are being processed concurrently.

## Job log system programming interface (SPI)



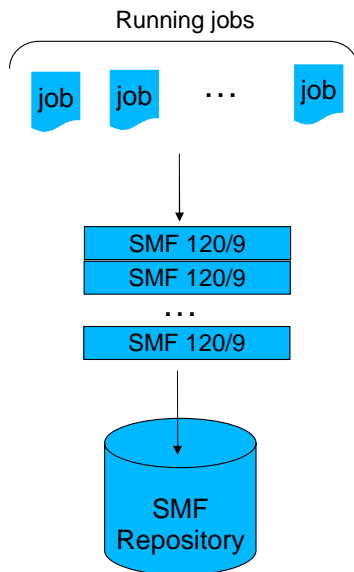
20

Overview

© 2011 IBM Corporation

The job log system programming interface (SPI) allows system-wide customization and control of job log content and destinations. You can use the job log SPI to direct the logging information to only the job log, only the WebSphere Application Server log, both logs, or neither log. The job log SPI can also be used to modify the content of the log line, as shown in the graphic.

## SMF 120 Subtype 9



- Optional Subtype 9 usage record<sup>1</sup> for each job
- Subtype 9 is rich, including:
  - zIIP, zAAP, and general purpose processor time
  - Times given in CPU seconds and service units
- Subtype 9 augmented with additional data for jobs:
  - Accounting string
  - Job ID
  - Submitter

<sup>1</sup> z/OS only

21

Overview

© 2011 IBM Corporation

In WebSphere XD Compute Grid V6.1, job usage accounting information is recorded in SMF 120 Subtype 20 format. Compute Grid V8 continues to support this format. Customers running Compute Grid V8 on WebSphere Application Server V8 can also obtain job usage accounting information in SMF 120 Subtype 9 format. When Compute Grid is running, Subtype 9 is augmented with additional job-related data such as the accounting string, job ID, and submitter.

## Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

[mailto:iea@us.ibm.com?subject=Feedback about XD80 ComputeGrid Overview.ppt](mailto:iea@us.ibm.com?subject=Feedback%20about%20XD80%20ComputeGrid%20Overview.ppt)

This module is also available in PDF format at: [../XD80\\_ComputeGrid\\_Overview.pdf](..../XD80_ComputeGrid_Overview.pdf)

You can help improve the quality of IBM Education Assistant content by providing feedback.



## 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>

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. in the United States, other countries, or both.

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.