IBM Software Group

WebSphere® Process Server V6.0.2
WebSphere Integration Developer V6.0.2

Overview of new functionality

This presentation provides an overview of new functions and enhancements for WebSphere Business Process Integration version 6.0.2.

IBM

Goals

- Introduce WebSphere Business Process Integration Enhancements with V6.0.2

- At the end of the presentation you should be able to:
  - ▶ Identify the new features available in the V6.0.2 release
    - WebSphere Process Server
    - WebSphere Business Process Choreographer

- Prerequisites to understanding this presentation
  - ▶ Knowledge of V6.0.1 WebSphere Business Process Integration Environment

2

The focus of this presentation is on the new features and enhancements to WebSphere Process Server and WebSphere Business Process Choreographer version 6.0.2. This is an overview of the differences between version 6.0.1 and version 6.0.2, therefore an understanding of version 6.0.1 is required. Each slide will list and discuss the V 6.0.1 product, the new features and enhancements in V 6.0.2 and the benefits resulting from the changes.

Agenda

- WebSphere Process Server Features Update
  - ▸ WebSphere Process Server Enhancements
  - ▸ System Administration Improvements
  - ▸ System Integration
  - ▸ Developer Support Ease of Use
  - ▸ Migration Enhancements
  - ▸ Platform Support

- WebSphere Business Process Choreographer Enhancements
  - ▸ Business Flow Manager Enhancements
  - ▸ Human Task Manager Enhancements
  - ▸ Business Process Choreographer Client Enhancements

- Summary

The new features for the WebSphere Process Server are grouped into functional themes. Each of these functional groups contains several new and enhanced features. The Business Process Choreographer enhancements are organized by logical groups. There will be a detailed description of each feature at an overview level in this presentation. Additional details will be described in specific presentations for those features.

Section

WebSphere Process Server features update

4

This section covers WebSphere Process Server features update.

WebSphere Process Server enhancements

Existing 6.0.1	▪Cannot Federate stand-alone profile
New 6.0.2	▪Support for federation of stand-alone profile ▪Promoting servers to a cluster ▪Configuration updates
Benefits	▪Migrate toward WebSphere Process Server distributed environment

5

Many enhancements are provided in WebSphere Process Server V6.0.2 and some of the most important enhancements are covered in the next two slides.

Previously you could not federate a server as the first member of a cell or move the server and its applications to become part of a cluster. The server could only be created from a template. V6.0.2 provides the capability to federate a stand-alone server that has been running in a production environment as the first member of a cell or a cluster. This allows the applications, buses, JDBC providers, databases and configuration from the stand-alone server to be preserved during the move into a cell or cluster. The federated server will continue to use the existing database that WebSphere Process Server used as a stand-alone server.

These changes now provide the capability to grow a production environment. The stand-alone server can be used in production until the applications grow to the point of requiring the network deployment environment for better control. Once the application usage grows beyond the servers capacity, the server and applications can be clustered to provide workload distribution and failover.

This support is provided only for the first member of a cell or cluster. All subsequent nodes to be federated must be empty servers. Promoting a server into a cluster is accomplished by creating a cluster and specifying the existing server as the first cluster member. The original server is used as the template for all subsequent cluster members.

There are also several configuration updates that allow guided configuration of the WebSphere Process Server for promotion of a server to a cluster, federating a stand-alone server into a node and configuration of both environments. The guides also allow direct access to the configuration panels, along with the instructions to change the configuration while using the guides.

In V6.0.1 the Business Process Choreographer Explorer used a single context root, which was the same across all servers. The Business Process Choreographer explorer now sets the URL to the server host name and allows the administrator to further modify the URL during configuration, which is used to uniquely identify each instance of the BPC explorer.

These enhancements have greatly improved the process of migrating a stand-alone production environment to a WebSphere Process Server distributed environment.

WebSphere Process Server enhancements

Existing 6.0.1	▪No copying of Business Ruleset or decision table ▪Multiple data sources used
New 6.0.2	▪Business Rules enhancements ▪Ease of Database Management Improvements ▪Product Installation
Benefits	▪Improve Business Rule support and configuration ▪Consolidate data base usage

6

The business rules support for WebSphere Process Server V6.0.2 has been enhanced with capabilities to make setup and configuration easier. Business rules that have been exported from one server can now be imported into another and support to use the java.util.List APIs on 'business object list attributes' has been added, allowing access to the rules using the list methods. A new *return* keyword forces execution of the rule list to end early and return the results, and an *otherwise* clause can be specified on decision table conditions from within WebSphere Integration Developer.

Audit logs are added to track additions or changes to the rule information and the audit log entries can be sent to either the SystemOut.log file, to a separate audit log file, or both. When auditing is enabled through the administrative console or WebSphere Studio Application Developer, it is enabled for all business rules. Filtering of the log entries is not supported. There is a separate audit log file for each server in a network deployment environment.

Improvements in database management now allow all WebSphere Process Server databases to be supported by business rules and a new database selection list is used across all panels for improved usability. The use of a common database for many of the base components improves the database usage and configuration.

The following information about product installation applies to distributed platforms. The z/OS installation has been improved to take advantage of installation wizards within the administrative console. The Network Deployment configuration has also been improved and Cloudscape databases are no longer created for the Service Component Architecture messaging engines. For more information on z/OS product installation, see the z/OS installation presentations.

The distributed product installation now includes the WebSphere Process Server, WebSphere Enterprise Service Bus and the application client installations, providing a consistent look across each of these product installations.

If a previous installation of WebSphere Process Server V6.0 is present, the 6.0.2 installation now allows the 6.0.2 version to be installed without removing the previous installation. An option of using the same port numbers is offered if the installations are intended to be run separately.

On systems that require root access to allow the product installation, a profile and user can be set up during root access, allowing the product installation to occur without root access. This is important to allow protection of the system and the flexibility to have a non-root user capability to install and manage WebSphere Process Server.

There are new profile creation panels for the service component architecture, common event infrastructure, business process choreographer and database configuration. Each of these allow individual security access configuration and database selection. The common database selected in the database configuration panel will be used for the recovery, relationship, mediation, AppScheduler, customization and lock manager components, which are installed as part of WebSphere Process Server.

System administration improvements

Existing 6.0.1	▪Integrated applications were not connected
	▪No cursor for large query sets
New 6.0.2	▪Task filtering, guided activities

	▪Administrative Console
	  ▸ Failed Event Manager
	  ▸ Common Base Event Browser
	  ▸ Relationship Manager

	▪Relationship Manager Improvements
Benefits	▪Ease of use
	▪References between events
	▪Reduction of memory use

System administration usability improvements have been added for this release. The task filtering selector has been enhanced to allow restriction of menu items based on the selection of either application integration, server and bus or all. The guided activities has added the process to configure network deployment and clustered environments. This process will guide you through the configuration with instructions and provide each configuration panel as needed.

In the V6.0.1 release, there was no linkage between the Failed Event Manager, Common Base Event Browser, and BPC explorer. In the V6.0.2 release, the viewing can go directly from one tool to the other using a common session ID.

The relationship manager has added the 'cache size hint' to limit the amount of results in memory per page. By setting the 'cache size hint', the relationship service will only retrieve the specified amount of relationship instances. The relationship manager now shows the next and back buttons which only get additional rows that can be shown at once in the table. The block cursor support addresses a major potential problem where a possible large return set could either leave the relationship manager in an unresponsive state or potentially the server could run out of memory.

There can be one data source specified for all the relationships and this data source is configured as part of the Relationship Service.

Relationship roles have been added to the relationship instance panel to group the participating instances together, allowing you to see all relevant data at once. In the new section for role information, each role type associated with this relationship instance has a table that lists all instances created for that role type and you can delete or create new instances from here. Clicking on the instance id link shows the role instance detail, where a you can change the 'key attribute values' and other properties and add and delete role instances. A create button has been added to the relationships panel, and a 'go to page' entry field allows navigating directly to a specific instance in the relationship viewer.

These are just a few of the system administration enhancements designed to improve usability and performance for the V6.0.2 release.

System integration

Existing 6.0.1	▪Limited integration with Tivoli® ▪Artifact access limited to local only
New 6.0.2	▪PMI and ARM statistics for SCA and Component implementations ▪ITCAM support to track SCA operations ▪Remote Artifact Loader ▪Data Transformation Framework
Benefits	▪Better integration with Tivoli tools ▪Access to remote artifacts using other locators ▪Allow data transformation to a XML or BO format

Better integration with Tivoli tools are provided by enabling Performance Monitoring Infrastructure (PMI) and Application Response Measurement (ARM) statistics for SCA and component operations. This will enable Request Metrics on a subset of servers using an ARM application, IBM Tivoli Composite Application Manager for SOA or IBM Tivoli Composite Application Manager for Response Time Tracking to track SCA operations. In previous version of WebSphere Process Server, durations were supported only if the event pair entry/exit were subsequently reached on the same Java thread. In V6.0.2, ARM and PMI durations can be generated when the entry/exit or failure occur on the same logical thread. The good request counter is incremented per element when the exit event point is reached. The bad request counter is incremented per element when the FAILURE event point is reached. The duration is computed between entry and exit, but the duration from entry to failure is not computed.

In V6.0.1, application artifacts had to be installed on each client or server for that application. The Remote Artifact Loader provides a mechanism to access application artifacts in a remote system. After an application is deployed, if another application must use any artifacts of the deployed application, the second application does not need to copy the first application's artifacts to its class path. Instead, it could use Remote Artifact Loader to load those artifacts remotely. For the V6.0.2 release, the remote artifact loader is provided for use only by the Business Process Choreographer.

The BPC tools will use the context setting APIs of the remote artifact loader to provide the application name to the client. Without the application name, the client does not use the remote artifact loader. This name cannot be set through the configuration for this release. The remote artifact loader server application will be installed on all WebSphere Process Servers that have the SOA CORE installed. It will rely on a normal WebSphere Process Server cell/cluster scenario to achieve high availability.

There is currently no support for data transformation with the version 6.0 of the adapters or the WebSphere Process Server runtime. The data transformation framework provides very similar functionality as the WebSphere Business Integration Adapter data handlers, allowing the format of the data to be transformed to and from what the business object expects. The framework also allows pre-parsing of the data stream to identify the business object type. The framework can also be used with JMS and other SCA components by creating custom bindings. The data transformation framework is based on the open service component architecture and the common model for transformation standards.

Developer support ease of use

Existing 6.0.1	▪No easy way integrate with existing Java/J2EE components. ▪Event sequencing is not supported.
New 6.0.2	▪SCA Java mapper ▪Event sequencing for asynchronous MDBs
Benefits	▪Easy integration and reuse of existing Java™ and J2EE components ▪Methods are invoked in the order they arrive

9

Several new features and enhancements to the V6.0.2 release allow you to build applications easier. The SCA Java mapper feature enables process server components to easily integrate with existing Java POJO and J2EE components without writing code. This requires that the imported component provide a WSDL specification and a component description file of the component and interfaces to be used to access it. The SCA Java mapper converts this information into Java types that can be used to communicate with the component.

Event sequencing has been added to control application invocations and provide migration support. Since Message Driven Beans (MDB) are invoked on separate threads in the server, it is possible to receive several requests from a client that must be ordered and that may not be invoked in the specified order. Event sequencing insures that the invocations will be invoked on the server in the order in which they are invoked on the client.

Event sequencing can be accomplished by specifying a qualifier that uses the operation and the attributes to form a key. The interface for the business object must be W-typed to allow this key to be included in the WSDL specification.

Event sequencing solves the sequencing problem by requiring events to acquire a lock before they are dispatched to the target component for invocation of the business logic. This solves the problem without losing scalability. When an event acquires a lock, it is dispatched to the target component for execution of business logic. When execution of business logic completes, the event releases the lock.

If an event cannot acquire a lock, the execution of the invocation is suspended. At a later point in time, the event may acquire a lock and then it will be dispatched to the target component.

WebSphere Integration Developer is enabled to add the 'event sequencing qualifier' to an interface method of an SCA component. When this qualifier is specified, the component is enabled for event sequencing.

Event sequencing maintains the correct sequencing order in case of system failures. The order is kept in the system database and will continue where it left off when the system comes back up. All component kinds and module types are supported. Event sequencing also works in a clustered environment to provide failover capability as long as the underlying messaging destination is not partitioned. For the WebSphere Process Server V6.0.2 release, event sequencing is supported only for components that will be invoked using the SCA asynchronous invocation style.

The esAdmin command line tool allows viewing and controlling locks within a server and can be used to resolve deadlocks, which may occur if a circular dependency is created.

Migration enhancements

Existing 6.0.1	▪Limited WICS to WebSphere Process Server migration ▪Limited BPEL support
New 6.0.2	▪Add migration capabilities for: 　▸ Isolation, Security, Custom Assembly Editor Template 　▸ Event Sequencing support 　▸ WICS APIs supported ▪Enhanced the BPEL processing
Benefits	▪More complete migration

The migration from WebSphere InterChange Server to WebSphere Process Server for V6.0.1 had some limitations that required manual steps. For V6.0.2, additional capabilities for Event Sequencing, Isolation, Migration of WICS Security Support and for Custom Assembly Editor Template are included, but there are still a few manual steps required.

The migration capabilities for WebSphere Process Server are provided as follows:

•Adding event sequencing to WebSphere Process Server allows the migration of WebSphere Interchange Server (WICS), WebSphere Business Integration and MQ Work Flow artifacts to map their use of event sequencing to the WebSphere Process Server runtime. Event sequencing in WebSphere Process Server has the functional equivalence to these products.

•Documentation is provided for Failed Events outlining the differences between WICS failed events and WebSphere Process Server Server failed events. This document will provide recommendations for how you can implement and manage failed as you did with WICS.

•Documentation is provided for outlining the security options that WebSphere Business Integration Server makes available. It describes how you can setup WebSphere Process Server Server security to provide a level of control similar to that used in WICS.

•WICS Collaboration Templates and Maps can contain your custom Java code referencing a set of supported APIs that were delivered with WICS. Many of these supported APIs are implemented in a deprecated fashion for the WebSphere Business Integration Server runtime. These APIs will take advantage of the Adapter usage pattern in WebSphere Process Server.

•The WebSphere Process Server Application Install will read the administrative artifacts data created by the migration and create WebSphere Business Integration Server Administrative Objects matching those found in WICS.

The artifact-level migration support from WICS to WebSphere Business Integration Server is provided by a WebSphere Business Integration Server command line interface and a WebSphere Integration Developer plug-in. These two interfaces convert the WICS artifacts into their corresponding target artifact form for consumption by WebSphere Process Server Server runtime or WebSphere Integration Developer tool.

After running the V6.0 migration utilities, the BPEL service is not fully re-wired and the assembly diagram has an unwired BPEL component for each migrated BPEL process. After running the enhanced migration utilities in WebSphere Integration Developer V6.0.2, the BPEL service is fully re-wired. The inbound SCA protocols and bindings and the SCA component for BPEL partner services for outbound are wired correctly.

The WebSphere Studio Application Developer non-BPEL services that can be invoked externally still require manual migration. These services cannot be automatically migrated or re-wired because WebSphere Studio Application Developer does not store .component or .wiring files for them.

All V6.0.2 enhancements are made to the migration utilities only. There are no user interface changes needed to the existing wizard.

Platform support

Existing 6.0.1	▪No access to databases on z/OS® for distributed platforms
New 6.0.2	▪New Platforms supported ▪z/OS support for Remote DB2®
Benefits	▪New platform support ▪DB2 on z/OS support for distributed platforms

11

© 2007 IBM Corporation

The platform updates for WebSphere Process Server include:

- •iSeriesi5/OS
- •Solaris 10 (Sparc, Opteron)
- •HP/UX 11i2 (PA-RISC)
- •SUSE Linux Enterprise Server (SLES) 10
- •Red Hat Enterprise Linux (RHEL) 4

The distributed platforms now support the use of DB2 on z/OS as a remote Database Management System for WebSphere Process Server and ESB production data only. This functionality will be specific to the creation or augmentation of stand-alone and deployment manager profiles. There is no database configuration done for the creation or augmentation of a custom profile. There are two places that the current Profile wizard will be impacted by adding remote DB2 database support on a z/OS. The first is during the CEI configuration process and the second is during the Common Database configuration process, both of which are minimized by the use of the current function provided in the common database code that provides support of remote DB2 database on z/OS.

During profile creation, the panels for either the Common Event Infrastructure or the Database Configuration for common components, the DB2 UDB for z/OS database can be selected. When the z/OS database is selected, there are two additional database configuration panels shown to provide information needed to access the z/OS database.

Section

Business Process Choreographer enhancements

12

This section covers the latest V6.0.2 enhancements for Business Process Choreographer, including the Business Flow Manager and Human Task Manager. There are also some enhancements to the Business Process Choreographer Clients that come with the product.

IBM

Business Flow Manager enhancements

Existing 6.0.1	• Limited functions and performance
New 6.0.2	• Information service activity • Cleanup for completed processes • Uninstall applications with running instances • Parameter passing by reference • Generic Web Service Client Interfaces • Remote Client Support
Benefits	• Functionality expanded and performance improved

13

Overview of new functionality

© 2007 IBM Corporation

V6.0.1 included a few functional and performance enhancements that version V6.0.2 improves upon. The Information Service Activity, which was previously available from Alphaworks, now provides Integration Information for BPEL (II4BPEL) support with direct interaction with the Information Server product. This activity supports discovering and invoking operations for information services, and allows referencing relational tables, rather than copying the data to improve data access performance. Another performance improvement is improved cleanup of process information. Completed, top-level process instances and all their associated data, such as activity instances, child process instances, and inline task instances, can be selectively deleted from the Business Process Choreographer runtime database. Both an MBean function and a Jython script are provided for this administrative task. Other than deleting data, V6.0.2 adds a scriptable option that allows for last resort uninstall of applications with existing processes or tasks. Normal uninstallation of process applications requires that no corresponding process or task instances are running. A new force flag in the bpcTemplate.jacl script allows the server to run in development mode, where the server can uninstall applications with existing processes or tasks. This simplifies testing of running applications and allows last resort option for production systems. Another performance improvement is the by-reference semantics used for synchronous, module-local BPEL interactions. If an SCA service component is invoked by a client contained in the same SCA service module and it uses synchronous service invocations, data is exchanged by-reference. However, for all other asynchronous or out-of-module invocations, the data is still passed by-value. Probably one of the most important functional improvement is the addition of generic Web service client Interfaces. WebSphere Web services, .Net Web services, and custom Web service clients can now interact with Business Process Choreographer components using generic Web services interfaces. The Business Process Choreographer EJB API methods are exposed as generic Web service operations. Lastly, V6.0.2 introduces remote client support from the WebSphere Process server where remote BPC clients can exchange business objects and business object schemas from remote WebSphere Process Server installations. The Business Process Choreographer artifacts are loaded dynamically using the Remote Artifact Loader (sometimes referred to as RAL or AL).

Business Flow Manager enhancements

Existing 6.0.1	▪Limited usability and dynamic customization
New 6.0.2	▪Query properties for process-level variables ▪Enhanced support for custom properties ▪Dynamic invocation of sub-processes ▪Access to current fault in fault handler
Benefits	▪Improved usability and more dynamic function

14

In V6.0.1, Business Process Choreographer had some limitations on usability and dynamic customization that V6.0.2 improves upon. First, query properties for process-level variables allows developers to declare which parts of a process variable should be accessible through the query() function. In addition to enabling customized retrieval and filtering of business processes based on business data, it improves developer ease of use by eliminating the need to write custom code to copy process variables into custom properties. Inline with the query properties for process-level variables, V6.0.2 adds enhanced support for process-level custom properties and activity-level custom properties. For process-level custom properties, new APIs and database views can be used to retrieve the custom properties of a process template. For activity-level custom properties, the existing APIs can now be used to get and set custom properties of an activity instance even if the activity has not yet been reached by process navigation. Following the dynamic theme, version 6.0.2 also allows dynamic invocation of sub-processes. That is, child processes can be dynamically determined at runtime instead of hard coded at development time. The sub-process template to be invoked at runtime is identified dynamically using a unique template name, which is held in a BPEL variable. One major improvement for developers is the addition of access to the current fault in fault handler. This gives developers the flexibility to provide one common catch-all fault handler, rather than a specific individual catch for each fault type. A new Java snippet method called `getCurrentFaultAsException() provides access to the current fault being handled in the form of a Java exception.`

Human Task Manager enhancements

Existing 6.0.1	▪Limitations in a few key dynamic operations
New 6.0.2	▪Ad-hoc human tasks ▪Suspend human tasks with duration ▪Post-processing of staff query results
Benefits	▪Dynamic functionality improved

15

Version 6.0.2 introduces a few key dynamic operations that limited version V6.0.1. The ability to create and start subtasks and follow-on tasks in an ad-hoc way through existing task templates. APIs like createAndStartTaskAsSubTask() and completeWithNewFollowOnTask() are used to start a task in a more dynamic way. Suspending tasks is not a new feature, but scheduling suspended tasks to be hidden from your task list for a duration of time is. A follow-up or suspend task can be scheduled when you need to organize or delay tasks without giving up ownership, instead of rejecting work items with a cancel claim. The task is then temporarily hidden from a user's task list and can optionally specify a duration for hiding the task. The task reappears on task list after the suspended duration. Though the task is suspended, timers continue to run which means expirations and escalations are still valid during the suspended state. Another dynamic addition to the version 6.0.2 lineup is the post-processing of staff query results, which provides the flexibility to change staff query results. This allows you to add or remove users, change the result type, or plug in customer-specific workforce management policies like changing a result set to not assign work to users with an already high work queue list.

Human Task Manager enhancements

Existing 6.0.1	▪Limitations for functional usability
New 6.0.2	▪Enhanced support for groups ▪Customizable e-mail notification ▪Server Controlled Page Flow ▪Binary custom properties ▪Enhanced runtime configuration ▪Human Workflow Web Client Generation
Benefits	▪Tasks easier to create, manage, and operate

In version 6.0.1 there were several functional and usability enhancements for human tasks. Enhanced support for groups significantly improves performance when dealing with large groups, as people assignments are managed on a per-group basis and not on an individual user basis. This allows administrators to easily and efficiently transfer work from one group to another group with a single operation instead of transferring work to every user of a group. Another group usability improvement is the work queue environments, which allow for custom views and queries that only show your tasks organized by the groups you are a part of. In addition to groups, the email notification feature is now more customizable. The E-Mail notification feature is not new, however the ability to craft E-Mails with an HTML Editor and add business process context variables for business relevant data to be displayed is. Furthermore, you can provide a link to the Business Process Choreographer Explorer from the E-Mail to directly solve the issue you were E-Mailed about. Another advancement is the server controlled page flow. The new completeAndClaimSuccessor() interface is added to the BusinessFlowManagerService interface API (this is not an HTM API). The Business Flow Manager API handles process work to claim next activity for the Human Task Manager, allowing one user working through inline human tasks to claim the next succeeding activity after completing an inline human task activity. The server handles the required logic and cuts complex logic on the client side, requiring multiple interactions with the server. Another useful feature of version 6.0.2 are binary custom properties. These custom properties move beyond string support to attach binary files to human tasks. For example, you can now add an image file of home to mortgage contract in a mortgage approval task. Enhanced runtime configuration for Business Process Choreographer, including additional items for E-Mail services, staff resolution, and group work items have been added to the administrative client for WebSphere Process Server. Human Task Manager E-Mail services now allows for setting sender e-mail address and URLs to be used for hot links for custom clients. The staff resolution has a staff query refresh schedule using CRONTAB syntax for post-processing and a timeout setting for the staff query. Also "group work items" has a selection added for turning on group work items from the server. Another new usability feature of version 6.0.2 is the human workflow Web client generation. A wizard has been added to the WebSphere Integration Developer context menu to quickly generate a JSF or JSP Web client that allows you to begin working with human tasks very quickly and easily. After generating the base client, you can add more custom features with JavaServer Faces and JavaServer Pages and edit Cascading Style Sheets (CSS) to change page themes, colors, and fonts on the Web client.

BPC Observer enhancements

Existing 6.0.1	▪Business Process Choreographer Observer ▪Shipped as sample only
New 6.0.2	▪BPC Observer becomes its own product ▪Fills gap between Tivoli ITCAM and WebSphere Business Integration Monitor ▪Contains both drill down and savable user defined reports ▪JSF based UI like BPC Explorer ▪V6.0.1 Observer sample not touched
Benefits	▪Ability to access key BPC statistics ▪Determine load and reliability of process infrastructure

The Business Process Choreographer Observer was a sample in version 6.0.1 and remains the same during the product update for version 6.0.2. It continues to work on a 6.0.2 runtime. However, the removal of the version V6.0.1 sample is highly recommended. Business Process Choreographer Observer is a JSF based Web client that provides customizable reports and graphical charts of historical and accumulated data of business processes and supports flexible drill-down capabilities to retrieve statistical data on processes and activities. With an easy setup, users of the Business Process Choreographer Observer can determine how many processes are running by the underlying BPEL infrastructure and what the actual load of process infrastructure is. Furthermore, it can determine how reliably processes run by keeping track of the percentage of failed instances and the duration of the execution of instances. If these statistics are collected regularly, the Business Process Choreographer Observer can detect changes in the usage pattern of the BPEL infrastructure. In this way, increasing load can be detected and resources can be added to the processing infrastructure. Business Process Choreographer Observer fills the gap between the IT monitoring of IBM Tivoli Composite Application Management (ITCAM) and the business level monitoring of WebSphere Business Integration Monitor.

IBM

BPC Explorer enhancements

Existing 6.0.1	▪ Usability and customization limitations
New 6.0.2	▪Graphical Structure View ▪Expanded Query ▪Custom Views ▪Business State Machines ▪Linkage from WebSphere Business Integration Tools
Benefits	▪ Easier to view, operate, and customize

18

Several usability and customization limitations were addressed in version 6.0.2. The graphical view for process structure and instances closely resembles Business Process Editor inside WebSphere Integration Developer with a similar look and actions. The benefit is a comfortable GUI for seeing the big picture of running processes and the ability to navigate process details quickly from the runtime, and not in the development environment.

An expanded query also allows for better usability and customization. The Business Flow Manager and Human Task Manager support custom properties on human tasks and processes. You can extend the schemas to filter for or display custom information to act as filters in searches as columns in lists where custom properties are displayed as details of an entity. The Business Flow Manager introduces the concept of query properties for variables that allow queries on selected content of BPEL variables. Another addition to usability are custom views. First, a BPC Administrator customizes BPC Explorer for all users, meaning they set default welcome view, customize navigation area and views. They can also Delete, Edit, Copy, and Move placement of links up or down. Users have the ability to Delete, Edit, and Copy their custom created views. Therefore, each user can personalize the Business Process Choreographer Explorer within their scope. Once a user or administrator saves a view, they will be shown in the navigation area and persisted as stored queries. Public stored queries by Administrator can be seen and used by all users. Private stored queries are just for each User.

As an added bonus for those who use business state machines (BSM), the BPC Explorer supports 2 new attributes to display state and correlation information. This leverages the custom and query properties for variables in BPEL processes generated from BSMs in V6.0.2. BSMs become BPEL processes at runtime and carry BSM specific custom properties and query properties for variables. You can use these custom and query properties as filter criteria or list columns.

Lastly, linkage to and from WebSphere Business Integration Tools has been added to version 6.0.2. From the Failed Event Manager, you can see a BPC Explorer tab that allows you to work on the activity with the failed event. For the Common Base Event Browser, you can see detailed event information for Business Process Choreographer activities. With the Human Task Manager notification e-mails, you can configure E-Mail notifications with direct links to BPC Explorer so that you can immediately work on the human tasks for which you receive email notification.

Summary

- WebSphere Process Server
  - System administration improvements
  - System integration and migration
  - Usability improvements

- Business Process Choreographer
  - Business flow manager
  - Human task manager
  - BPC clients

19

In summary, this presentation reviewed improvements in server federation that allow growing business applications to move from a stand-alone environment to network deployment and then into a clustered environment. This represents a great improvement over redeploying and losing valuable application data and configuration.

V6.0.2 includes several improvements in business rules, product installation and database management. Many usability improvements were shown for event viewing between each viewer and reduction in the number of relationships viewed per screen.

Event sequencing was added to support MDB execution and migration from WICS, WebSphere Business Integration and MQWF.

Functionality, dynamicity, usability, customization, and performance are major version 6.0.2 themes for business process choreographer. The business flow manager provides enhanced functionality and performance with features like generic Web service interface exposure and the ability to run remotely off a WebSphere Process Server. The human task manager provides enhanced dynamicity and usability with features like ad-hoc subtasks and follow-on tasks. The business process choreographer clients also offer improved functionality and customization.

IBM Software Group

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

DB2   Tivoli   WebSphere   z/OS

J2EE, Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance 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 or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2007. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.

20

Overview of new functionality   © 2007 IBM Corporation

WPIv602_WPSWhatsNew602.ppt   Page 20 of 20