WebSphere® Process Server for z/OS

Version 6.0.1

Product Overview

Note
Before using this information, be sure to read the general information in "Notices" on page 35.

23 June 2006

This edition applies to version 6, release 0, modification 1 of WebSphere Process Server for z/OS (product number 5655-N53) and to all subsequent releases and modifications until otherwise indicated in new editions.

To send us your comments about this document, email doc-comments@us.ibm.com. We look forward to hearing from you.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Product overview

You can use links to overview and introduction materials to gain a high-level understanding of IBM WebSphere Process Server.

WebSphere Process Server documentation (in PDF format)

Hardware and software requirements for WebSphere® Process Server are available on the WebSphere Process Server system requirements Web site.

Information roadmaps for WebSphere Process Server are available on the IBM® developerWorks® Web site under WebSphere Business Integration information roadmaps.

Overview materials include an introduction to the product, new features, information about other IBM products that work with WebSphere Process Server, a technical introduction to the architecture and components, and information about samples, standards compliance, and globalization.

Introduction to WebSphere Process Server

IBM WebSphere Process Server is the next generation business process integration server that has evolved from proven business integration concepts, application server technologies, and the latest open standards.

IBM WebSphere Process Server, which supports a service-oriented architecture (SOA), is the ideal platform for business applications that require business integration using different technologies. Using the WebSphere Integration Developer tool set, business integration solutions can be created using simplified integration mechanisms, such as the Service Component Architecture (SCA) programming model and the Service Data Objects (SDO) data model. SDO business objects can be defined, transformed, routed, and mapped using SCA components. WebSphere Adapters supply connectivity to back-end Enterprise Information Systems (EIS). With WebSphere Process Server, business integration applications may define business logic and processes based on Web Services - Business Process Execution Language (BPEL), human tasks, and business rules. For the runtime monitoring of the business integration solutions, WebSphere Process Server provides Common Event Infrastructure (CEI), which centralizes the monitoring of the various events that can occur in these applications.

WebSphere Process Server enables deployment of standards-based integration applications in an SOA structure. SOA is an application framework that takes everyday business applications and breaks them down into individual business functions and processes, called services. SOA is a conceptual description of the structure of a software system in terms of its components and the services they provide, without regard for the underlying implementation of these components, services and connections between components. Loosely coupled integration applications that are based on SOA provide flexibility and agility. You can implement integration solutions independent of platform, protocols and products. For more information about SOA, refer to the Service-Oriented Architecture (SOA) from IBM Web site.

Hardware and software requirements

To view the official statement of supported hardware and software for WebSphere Process Server, go to the WebSphere Process Server system requirements Web site.

Information Roadmaps

WebSphere Process Server can be used by a variety of users including *solution deployers*, who deploy solutions developed in WebSphere Integration Developer to WebSphere Process Server, *solution administrators*, who use the administrative console to administer solutions, and *operators*, who use the administrative console to operate a solution.

To help you to navigate through the available information sources, both within and beyond the product information centers, WebSphere Process Server roadmaps are available online from IBM developerWorks. These roadmaps list high level goals based on your user role and point to documentation resources that are useful to accomplishing your goals in WebSphere Process Server.

To access information roadmaps for WebSphere Process Server, go to the WebSphere Business Integration information roadmaps page on IBM developerWorks.

What is new in this release

IBM WebSphere Process Server for z/OS, version 6.0.1, includes several new features.

Welcome to IBM WebSphere Process Server for z/OS, version 6.0.1, which includes the following new features:

- Service Component Architecture - one simplified integration framework that leverages existing IT.
- A deployment environment for applications developed in IBM WebSphere Integration Developer, version 6.0, a new simplified development tool with visual editors for component development, assembly, integrated testing, and deployment.
- Support for all styles of integration - including human tasks, role-based task assignments, and multilevel escalation.
- Ability to change business processes with minimal programming skills, without redeploying the application.
- Business rules, business state machines, and selectors to dynamically choose interfaces based on business scenarios.
- Broad reach in integration and support for IBM WebSphere Adapters.
- Support for business-to-business (B2B) applications through a limited license of IBM WebSphere Partner Gateway included with WebSphere Process Server.

6.0.1+ Enhancements include:

- Clustered servers for workload management.
- Enterprise service bus integration capabilities: service-oriented, message-oriented and event-driven technologies that provide standards-based, messaging infrastructure.

Product family overview

IBM WebSphere Process Server works with several other IBM products including WebSphere Integration Developer, WebSphere Application Server, WebSphere Adapters, WebSphere Business Integration Adapters, WebSphere Application Server Toolkit, WebSphere Business Modeler, WebSphere Business Monitor, Rational Application Developer, Rational Software Architect, WebSphere Partner Gateway, and WebSphere Portal.

IBM WebSphere Process Server is based on the robust J2EE 1.4 infrastructure and associated platform services provided by WebSphere Application Server, version 6.0. WebSphere Process Server is built on WebSphere Application Server, Network Deployment, version 6.0. WebSphere Process Server also works with infrastructure and platform services from WebSphere Application Server, version 6.0. For more information about WebSphere Application Server, see the WebSphere Application Server for z/OS information center.

IBM WebSphere Integration Developer is the development environment for WebSphere Process Server. For more information about WebSphere Integration Developer, see the WebSphere Integration Developer information center.

WebSphere Process Server and WebSphere Integration Developer include additional capabilities that make it possible to model, build, deploy, install, configure, run, monitor, and manage integration applications. WebSphere Integration Developer complements IBM WebSphere Business Modeler, version 6.0, and IBM WebSphere Business Monitor, version 6.0, and can be used in conjunction with IBM Rational® Application Developer, version 6.0, or IBM Rational Software Architect, version 6.0, to create a unique, integrated and powerful integration development platform. For more information about these products, see the WebSphere Business Modeler information center, the WebSphere Business Monitor information center, the Rational Application Developer information center, and the Rational Software Architect information center.

IBM WebSphere Adapters, version 6.0., and IBM WebSphere Business Integration Adapters (based on IBM WebSphere Business Integration Framework, version 2.6) allow for integration of existing Enterprise Information System infrastructure and applications that are deployed on WebSphere Process Server. For more information about these products, see the WebSphere Adapters information center and the WebSphere Business Integration Adapters information center.

In addition, you can extend existing applications for reuse in enterprise processes with an IBM enterprise modernization portfolio that includes IBM CICS® Transaction Gateway and IBM WebSphere Host Access Transformation Services. For more information about these products, see the CICS Transaction Gateway information center and the WebSphere Host Access Transformation Services information center.

The IBM WebSphere Application Server Toolkit is a set of basic tools that help you to assemble, test, and deploy Web services in WebSphere Process Server. For more information, see the WebSphere Application Server Toolkit documentation on the WebSphere Application Server for z/OS information center.

IBM WebSphere Partner Gateway used with WebSphere Process Server supports business-to-business (B2B) applications. A limited license of WebSphere Partner

Gateway is included with WebSphere Process Server. For more information about WebSphere Partner Gateway, see the WebSphere Partner Gateway information center.

IBM WebSphere Portal provides access to various administrative functions and allows portlets to have access to business processes and other Service Component Architecture services in WebSphere Process Server. For more information about WebSphere Portal, see the WebSphere Portal information center.

Architectural overview of WebSphere Process Server

IBM WebSphere Process Server combines integration capabilities with a composite application platform to deliver an integration platform with a fully converged, standards-based business process engine, using the full power of WebSphere Application Server.

IBM WebSphere Process Server is a service-oriented architecture (SOA) integration platform built on a uniform invocation programming model and a uniform data representation model.

The base runtime infrastructure for WebSphere Process Server is WebSphere Application Server. The Service Component Architecture and business objects that are part of the SOA core provide uniform invocation and data-representation programming models. The SOA core includes the Common Event Infrastructure for generating events for the monitoring and management of applications running on WebSphere Process Server. Supporting services provide the foundational business object and transformation framework for WebSphere Process Server. Service components represent the functional components required for composite applications.

The combination of a powerful foundation (WebSphere Application Server and the SOA Core) and service components in WebSphere Process Server allows quick development and deployment of sophisticated composite applications that run on WebSphere Process Server.

One component-based framework addresses all styles of integration.

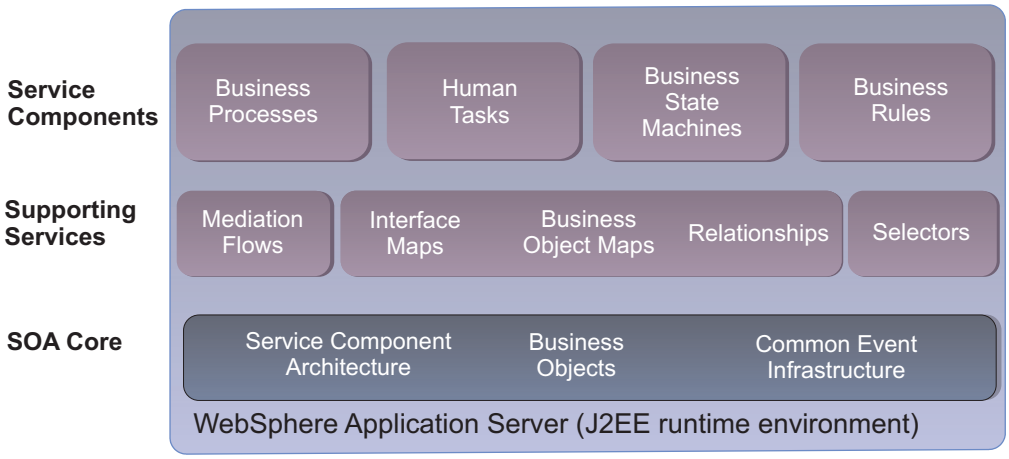

Service-oriented architecture core

The service-oriented architecture (SOA) core of WebSphere Process Server provides both uniform invocation and data-representation programming models and monitoring and management capabilities for applications running on WebSphere Process Server.

SOA is a conceptual description of the structure of a software system in terms of its components and the services they provide, without regard for the underlying implementation of these components, services and connections between components. The Service Component Architecture and business objects that are part of the SOA core provide uniform invocation and data-representation programming models for applications deployed on IBM WebSphere Process Server. The SOA core also includes the Common Event Infrastructure for generating events for the monitoring and management of applications on WebSphere Process Server.

One component-based framework addresses all styles of integration.

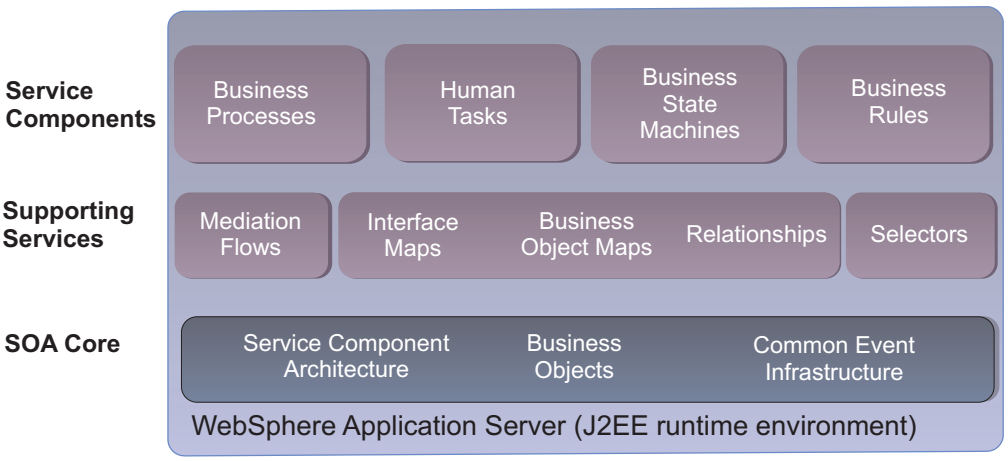

Service Component Architecture

Service Component Architecture presents all elements of business transactions – access to web services, Enterprise Information System (EIS) service assets, business rules, workflows, databases and so on – in a service-oriented way.

Service Component Architecture separates business logic from implementation, so that you can focus on assembling an integrated application without knowing implementation details. The implementation of business processes is contained in service components.

Service components can be assembled graphically in the IBM WebSphere Integration Developer tools, and the implementation can be added later. The Service Component Architecture programming model narrows what developers must know about Java™ and J2EE or other implementation in particular scenarios to a core set of language concepts that are familiar to all who develop business applications in other programming languages today. This allows developers to quickly and easily integrate technologies.

Developers switching from classical application development environments face a much smaller learning curve; they can quickly become productive with this programming model. The Service Component Architecture programming model also helps experienced J2EE developers be more productive.

Service Component Architecture supports several standard service implementation types:

- Java objects, which implement a Java class. As in the Java programming language, instances of Java components at run time are referred to as Java objects.
- Business process components, which implement a business process. The implementation language is the Business Process Execution Language (BPEL) and its IBM extensions.
- Human task components, which represent and implement a task typically performed by a person in a business process or an integration application.
- Business state machine components, which are used when applications work with artifacts that have a set of states. A state machine defines what the artifacts can do at a point in time.
- Business rule components, which determine the outcome of a business process based on a context and can be designed as if-then rules, decision tables, or decision trees. Business rules within a business process allow applications to respond quickly to changing business conditions. The rules are independent of the business process itself, and you can change them at any time without having to redo your process.

Service qualifiers govern the interaction between a service client and a service on the WebSphere Process Server runtime environment. Service qualifiers are quality of service specifications that define a set of communication characteristics required by an application for transmission priority, level of route reliability, transaction management, and security level. An application communicates its quality of service needs to a runtime environment by specifying service qualifiers. Quality of service qualifiers can be specified when wiring components in the assembly editor in WebSphere Integration Developer. These specifications, when running on WebSphere Process Server, determine how the clients interact with the target components. Depending on the qualifiers specified, the run time can supply additional required processing.

WebSphere Process Server solutions rely upon the underlying WebSphere Application Server capabilities for transaction, security, and workload management to provide a scalable integration environment.

For business processes, WebSphere Process Server offers support for transactions involving multiple resource managers using the two-phase commit process to ensure atomic, consistent, isolated, and durable (ACID) properties. This capability is available for both short-running flows (single transaction) and long-running flows (multiple transactions). You can group multiple steps in a business process into one transaction by modifying transaction boundaries in WebSphere Integration Developer.

Because not all service invocations support two-phase-commit transactions, WebSphere Process Server also includes recovery capabilities. If a failure occurs in the middle of running an integration application, the server detects it and allows an administrator to manage the failed event from the failed event manager.

Service Data Objects and business objects

Service Data Objects and business objects define the data flowing between components that are defined in Service Component Architecture.

Service Data Objects (SDOs), part of WebSphere Application Server capabilities that are built into WebSphere Process Server, provide a framework for data application development that simplifies the J2EE data programming model.

WebSphere Process Server includes business objects, which are enhanced SDOs. Business objects are based on a data-access technology called Service Data Objects. SDOs provide a universal means of describing disparate data (like JDBC ResultSet, XML Schema described data, for example). Business objects include some extensions that are important for integration solutions and are used to further describe the data that is being exchanged between Service Component Architecture services. Business objects are part of the Service-oriented architecture (SOA) core of WebSphere Process Server.

A business object is a set of attributes that represent a business entity (such as Employee), an action on the data (such as a create or update operation), and instructions for processing the data. Components of the integration application use business objects to exchange information and trigger actions. Business objects are flexible because they can represent many kinds of data. For example, in addition to supporting the data canonicalization model of traditional integration servers, they also can represent data returned from a synchronous EJB Session Bean facade or a synchronous business process, and then they can be bound to IBM WebSphere Portal portlets and JSF components.

Business objects are the primary mechanism for representing business entities, or documenting literal message definitions, enabling everything from a simple basic object with scalar properties to a large, complex hierarchy or graph of objects.

In WebSphere Process Server, business object framework is made up of these elements:

- Business object definition
- Business graph definition
- Business object metadata definition
- Business object services (service APIs)

A business object definition is the name, set of ordered attributes, properties, version number, and application-specific text that specify a type of business object. A business graph definition is the wrapper added around a simple business object or a hierarchy of business objects to provide additional capabilities, such as carrying change summary and event summary information related to the business objects in the business graph. A business object metadata definition is the metadata that can be added to business object definitions to enhance their value when running on WebSphere Process Server. This metadata is added to the business object's XML schema definition as well known xs:annotation and xs:appinfo elements. Business object services are a set of capabilities provided on top of the basic capabilities provided by WebSphere Application Server Service Data Objects. Examples are services such as create, copy, equality, and serialization.

For more information about WebSphere Application Server Service Data Objects, see the WebSphere Application Server for z/OS information center.

Common Event Infrastructure

Support for Common Event Infrastructure and Common Base Events enhances tracking, auditing and monitoring of business processes for applications running on WebSphere Process Server.

Common Event Infrastructure provides facilities for the IBM WebSphere Process Server runtime environment to persistently store and retrieve events from many different programming environments. Events are represented using the Common Base Event model a standard, XML-based format that defines the structure of an event.

For more information about Common Event Infrastructure, refer to the *Administering WebSphere Process Server* PDF file.

Supporting services

Supporting services in IBM WebSphere Process Server address a number of transformation challenges for connecting components and external artifacts.

You can use mediation flows, interface maps, business object maps, relationships, and selectors to integrate applications running on IBM WebSphere Process Server.

One component-based framework addresses all styles of integration.

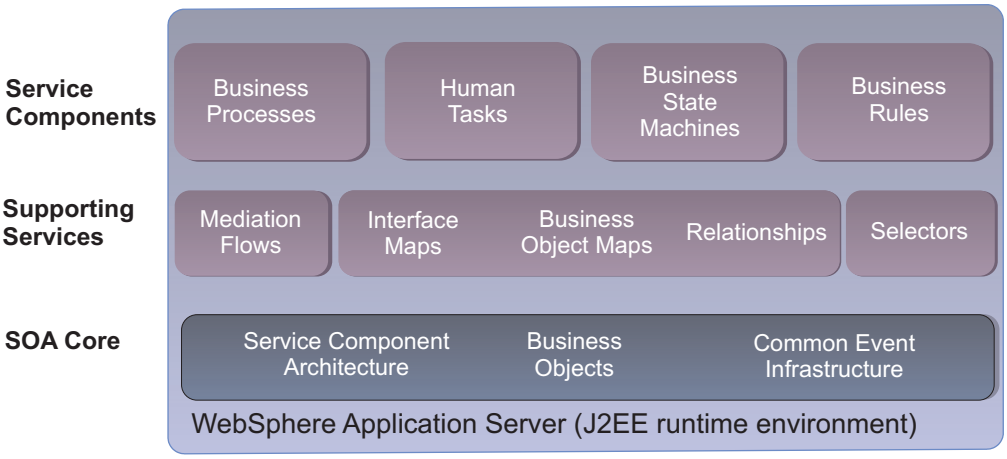

Mediation flows

Mediation flows intercept and modify messages that are passed between existing services (providers) and clients (requesters) that want to use those services.

A mediation flow mediates or intervenes between an export and import to provide functions such as message logging, data transformation and routing. Mediation flows are created in the IBM WebSphere Integration Developer and deployed as part of a mediation module in WebSphere Process Server.

Related concepts

Overview of enterprise service bus
Enterprise service bus capabilities in WebSphere Process Server support the integration of service-oriented, message-oriented and event-driven technologies to provide standards-based, messaging infrastructure to companies wanting a fast start to an enterprise service bus.

Interface maps

Interface maps reconcile the differences between components that have different interfaces.

Interface maps are supporting service components in IBM WebSphere Process Server that resolve and reconcile differences between interfaces in other Service Component Architecture (SCA) components to enable them to communicate. The

interface map captures a first-class pattern that allows module designers in IBM WebSphere Integration Developer to reconcile differences across multiple interfaces using transforms and other rudimentary operations. Interface maps are deployed on WebSphere Process Server as part of modules, also called SCA modules.

Business object maps

Business object maps support mappings between the source and target business objects.

Business object maps are supporting service components in IBM WebSphere Process Server that assign values to the target business objects service components based on the values in the source business objects service components.

Developers create the business object maps in IBM WebSphere Integration Developer.

Relationships

Relationships are supporting services in WebSphere Process Server applications that establish an association between data from two or more data types.

In IBM WebSphere Process Server, relationship manager is a tool for manually manipulating relationship data to correct errors found in automated relationship management or provide more complete relationship information. In particular, it provides a facility for retrieving as well as modifying relationship instance data. Relationship manager allows you to configure, query, view, and perform operations on relationship runtime data, including participants and their data. You create relationship definitions with relationship designer. At run time, instances of the relationships are populated with the data that associates information from different applications.

For more information about relationship manager, refer to the *Administering WebSphere Process Server* PDF file.

Selectors

Selectors provide flexibility at points in the processing of components in an application running on IBM WebSphere Process Server.

Selectors are supporting services that provide flexibility at points in the processing of service components during run time. A selector takes one invocation and allows different targets to be called based on the selection criteria.

Selectors add additional flexibility beyond business rules. Business rules are a fundamental part of businesses. Business rules drive the general processing of an application, invoking certain services to get the data through the application. For example, a rule may be: Two weeks before school starts, offer a back-to-school special price on our school-related merchandise. A selector takes one invocation and allows different targets to be called based on the selection criteria. For example, if the time is just before school starts, then the previous back-to-school offer would be called. However, if the season is the just as school ends, then a get-your-kids-ready-for-summer offer would be called.

The application is portable because it calls the same thing all the time. The business rule never changes. The actual processing differs (and calls different service components) because of the selector.

For more information about selectors, also called selector components, refer to the *Administering WebSphere Process Server* PDF file.

Service components

All integration artifacts running on IBM WebSphere Process Server (for example, business processes, business rules, and human tasks) are represented as components with well defined interfaces. Within the Service Component Architecture (SCA), a service component defines a service implementation.

Because all integration artifacts are represented as service components, also called SCA components, IBM WebSphere Process Server creates an environment with unparalleled flexibility. SCA components each have an interface and can be wired together to form a module deployed to WebSphere Process Server. This enables changing any part of an application without affecting the other parts. It is possible, for example, to replace a human task for an approval with a business rule for automatic approval simply by replacing the components in the assembly diagram without changing either a business process or the caller of the business process.

Components can interact with existing applications, using the following programming constructs:

- Java Beans
- Enterprise Java Beans
- Web services
- JMS messages

In addition, components can interact with other applications on enterprise information systems (EIS) with IBM WebSphere Adapters, version 6.0 and WebSphere Business Integration Adapters, based on WebSphere Business Integration Framework, version 2.6.

On top of the runtime infrastructure supporting services and the service-oriented architecture core, WebSphere Process Server offers a variety of ready-to-use SCA service components that can be used in integration applications.

One component-based framework addresses all styles of integration.

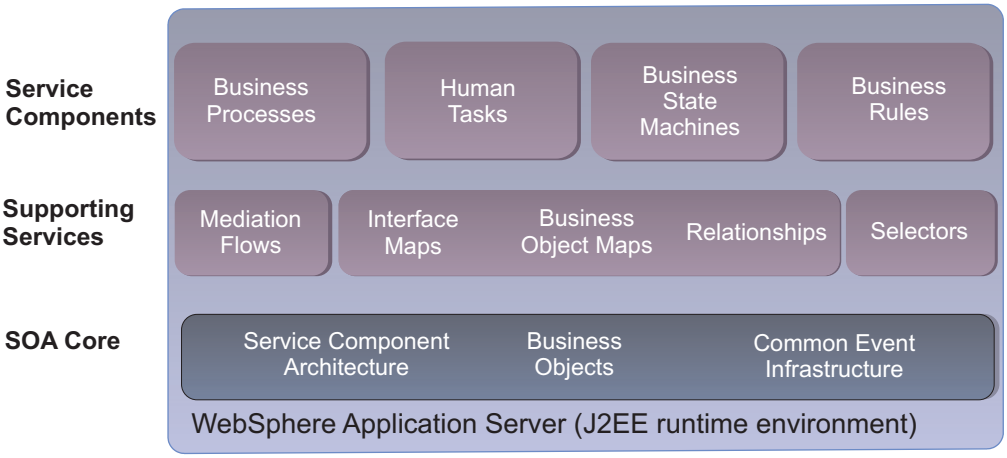

Business processes

Business processes are service components that provide the primary means through which enterprise services are integrated.

A business process is any system or procedure that an organization uses to achieve a larger business goal. When you break it down, you see that a business process is actually a series of individual tasks, and each task is executed in a specific order. As an integral part of applications running on IBM WebSphere Process Server, business processes provide the primary means through which enterprise services are integrated.

Business process components implement a fully supported Web Services Business Process Execution Language (BPEL) engine. WebSphere Process Server includes a business process choreography engine on top of the WebSphere Application Server. You can develop and deploy complex business processes in a simple development model with sophisticated support for long and short running business processes in a highly scalable infrastructure. You can either create BPEL models in WebSphere Integration Developer, version 6.0, or import them from a business model you created in WebSphere Business Modeler, version 6.0.

Web Services Business Process Execution Language (BPEL) is used to choreograph the flow of business processes. Business process integration services build on BPEL4WS version 1.1 and add major capabilities of the upcoming WS-BPEL version 2.0 specification.

For more information about business processes, refer to the WebSphere Process Server *Installing* PDF file.

Human tasks

Human tasks are stand-alone service components that can be used to either assign work to employees or to invoke other services.

The Human Task Manager, available in IBM WebSphere Process Server, supports ad-hoc creation and tracking of tasks. Existing LDAP directories (as well as operating system repositories and the WebSphere user registry) can be used to access user and group information. WebSphere Process Server supports multi-level escalation for human tasks including e-mail notification. It also includes a Web client to manage human tasks, and a set of Java Server Faces (JSF) components that can be used to create custom clients or to embed human task functionality into other Web applications.

Human task services allow role-based task assignment, invocation and escalation.

For more information about about human tasks, refer to the WebSphere Process Server *Installing* PDF file.

Business state machines

Business state machines are service components that specify the sequences of states, responses, and actions that an object or an interaction goes through during its life in response to events.

A business state machine provides another way of modeling a business process. This gives you the choice of representing business processes based on states and events rather then a sequential business process model.

For more information about monitoring business state machines in WebSphere Process Server, refer to the WebSphere Process Server *Monitoring* PDF file.

Business rules

Business rules are service components that declare policy or conditions that must be satisfied within your business.

Business rules make business processes more flexible. Because business rules determine the outcome of a process based on a context, using business rules within a business process allows applications to respond quickly to changing business conditions.

Business rule authoring is supported with IBM WebSphere Integration Developer. IBM WebSphere Process Server includes the business rules manager, a Web-based runtime tool for business analysts to update business rules as business needs dictate, without affecting other components or Service Component Architecture (SCA) services.

For more information about administering business rules in WebSphere Process Server, refer to the *Administering WebSphere Process Server* PDF file.

Imports, exports and adapters

Importing and exporting capabilities within the Service Component Architecture define a service module's external interfaces or access points for IBM WebSphere Process Server. Imports and exports can be either to other modules within a same application, or to other applications on enterprise information systems (EIS).

Imports identify services outside of a module, making them callable from within the module. *Exports* allow components in a module to provide their services to external clients.

A *module level* import or export lets modules access other modules. A *system level* import or export lets your application access applications on EIS systems as if they were local components, which allows working with IBM WebSphere Adapters and IBM WebSphere Business Integration Adapters.

WebSphere Adapters, version 6.0, and WebSphere Business Integration Adapters (based on WebSphere Business Integration Framework, version 2.6) provide a service-oriented approach to EIS integration.

WebSphere Adapters are compliant with J2EE Connector Architecture (JCA 1.5). JCA is the J2EE standard for EIS connectivity. EIS Import and EIS Export provide SCA components with the uniform view of the services external to the module. This allows components to communicate with the variety of external EIS systems using the consistent SCA programming model. WebSphere Adapters are assembled in WebSphere Integration Developer from imported RAR files and then exported as an enterprise archive (EAR) file and deployed on WebSphere Process Server.

WebSphere Adapters include the following:

- IBM WebSphere Adapter For Flat Files, version 6.0
- IBM WebSphere Adapter for JDBC, version 6.0
- IBM WebSphere Adapter for PeopleSoft Enterprise, version 6.0
- IBM WebSphere Adapter for Siebel Business Applications, version 6.0
- IBM WebSphere Adapter for SAP Applications, version 6.0

For more information about these products, see the WebSphere Adapters information center.

WebSphere Business Integration Adapters consist of a collection of software, Application Programming Interfaces, and tools to enable applications to exchange business data through an integration broker. Each business application requires its own application-specific adapter to participate in the business integration process. You can install, configure, and test the adapter using current WebSphere Business Integration Adapter Framework and Development Kit System Manager tools. You can use WebSphere Integration Developer to import existing business objects and connector configuration files, to generate artifacts, and to assemble the solution for WebSphere Process Server. Operational commands for the WebSphere Business Integration Adapters are part of the WebSphere Process Server administrative console. For more information about working with these adapters and WebSphere Process Server, see the WebSphere Business Integration Adapters information center.

Imports and exports require binding information, which specifies the means of transporting the data from the modules. The assembly editor in WebSphere Integration Developer sets up imports and exports, lists the bindings supported and simplifies the creation of them. A properties view displays the binding information.

Clusters and workload management

Clusters are sets of servers that are managed together and participate in workload management. The servers that are members of a cluster can be on different host machines, as opposed to the servers that are part of the same node and must be located on the same host machine.

Creating clusters for servers, also known as clustering servers, makes additional processing power available to your applications by distributing requests between the servers that comprise the cluster.

Servers that belong to a cluster are members of that cluster set and must all have identical application components deployed on them. Other than the applications configured to run on them, cluster members do not have to share any other configuration data. One cluster member might be running on a huge multi-processor enterprise server system, while another member of that same cluster might be running on a smaller system. The server configuration settings for each of these two cluster members are very different, except in the area of application components assigned to them. In that area of configuration, they are identical. This allows client work to be distributed across all the members of a cluster instead of all workload being handled by a single application server.

When you create a cluster, you make copies of an existing application server template. The template is most likely an application server that you have previously configured. You are offered the option of making that server a member of the cluster. However, it is recommended that you keep the server available only as a template, because the only way to remove a cluster member is to delete the server. When you delete a cluster, you also delete any servers that were members of that cluster. There is no way to preserve any member of a cluster. Keeping the original template intact allows you to reuse the template if you need to rebuild the configuration.

For more information about clusters, refer to Introduction: Clusters in the WebSphere Application Server for z/OS information center.

Overview of enterprise service bus

Enterprise service bus capabilities in WebSphere Process Server support the integration of service-oriented, message-oriented and event-driven technologies to provide standards-based, messaging infrastructure to companies wanting a fast start to an enterprise service bus.

This capability is based on the robust J2EE 1.4 infrastructure and associated platform services provided by WebSphere Application Server Network Deployment, version 6.0.2.

Related concepts

"Mediation flows" on page 8
Mediation flows intercept and modify messages that are passed between existing services (providers) and clients (requesters) that want to use those services.

Overview of the bus environment

The bus environment comprises an one or more service integration buses, ESB servers, and their resources, organized into logical administrative domains of cells and nodes.

If you create a complete (default) installation for WebSphere ESB, you get a stand-alone server on which you can deploy Service Component Architecture (SCA) modules without having to do any configuration of the server.

However, administrators may still want to act on the bus environment, so would benefit from some detail about the environment.

- The SCA runtime (exploited by mediation modules) uses queues on an SCA.SYSTEM service integration bus as a robust infrastructure to support asynchronous interactions between components and modules. The queues are hosted by the server as a member of the SCA.SYSTEM bus.
- The ESB server provides the integration technologies, infrastructure services, configuration, and runtime administration needed to run mediation modules and service applications in WebSphere ESB. As a bus member, the server has a messaging engine that provides the core messaging functionality of the SCA.SYSTEM bus.

Both the server and SCA.SYSTEM are configured with default properties that may be suitable for you to deploy and run your SCA modules.

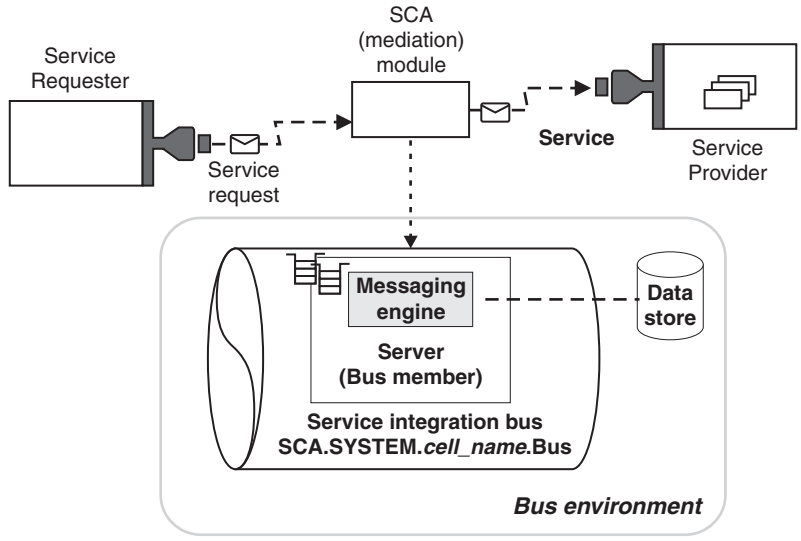

Figure 1. A bus environment with one server assigned to the SCA.SYSTEM service integration bus. As a bus member, the server is assigned one messaging engine, which has a data store for storing state data and messages. This figure also shows a mediation module deployed into the bus environment and assigned to the bus member.

For more advanced usage, you may want to change the configuration of the bus environment for a stand-alone server, or create a bus environment for a deployment manager cell. For example:

- You can configure a variety of quality of service from secure, assured delivery (where messages are guaranteed not to get lost and are transported securely) to best-effort (where messages may get lost in case of a system failure).
- You may want to set up a deployment manager cell to provide several servers to host mediation modules. This provides advantages of scalability, the ability to handle more client connections, and greater message throughput. You can also create server clusters, which enables you to manage a group of servers together and enables those servers to participate in workload management.
- Your complete bus environment may be made up of several stand-alone and deployment manager profiles, to provide separate administrative domains for different departments or to separate test and production facilities. Each profile has its own SCA.SYSTEMservice integration bus.

Besides the SCA.SYSTEM bus used for SCA modules, you can also create other service integration buses that you can use to support the service integration logic provided by the modules. For example, the SCA.APPLICATION.*cell_name*.Bus is provided and used to define JMS queue destinations and other JMS resources for modules deployed with JMS bindings.

You can create other buses for use as in WebSphere Application Server; for example, for applications acting as service requesters and providers within WebSphere ESB, or to link to WebSphere MQ. You can also use a WebSphere ESB deployment manager to manage separate application servers for use with applications and modules deployed onto WebSphere Application Server.

Related concepts

"Management of the server and bus environment" on page 22
With WebSphere ESB, administrators can create an environment of ESB servers and service integration buses that support the deployment of service applications.

Clients

WebSphere ESB provides Message Service clients that extend the connectivity of the enterprise service bus.

- Message Service Clients for C/C++ and .NET enable non-Java applications to connect to WebSphere ESB.

  Message Service Clients for C/C++ and .NET provide an API called XMS that has the same set of interfaces as the Java Message Service (JMS) API. Message Service Client for C/C++ contains two implementations of XMS, one for use by C applications and another for use by C++ applications. Message Service Client for .NET contains a fully managed implementation of XMS, which can be used by any .NET compliant language.
- Web Services Client for C++ provides a set of libraries and Java tools that enable you to build ANSI C++ web service client applications from existing Web Service Description Language (WSDL) files.

  The ANSI C++ web service client applications that you build from existing WSDL files, using the Web Services Client for C++ libraries and Java tools, are able to communicate with other similarly configured applications over HTTP using TCP/IP with SOAP protocols.

You can also install and use the J2EE client support from WebSphere Application Server Network Deployment, including Web services Client, EJB Client, and JMS Client. For information about installing J2EE client support, see **Installing Application Client for WebSphere Application Server** in the WebSphere Application Server for z/OS information center.

Mediation modules

Mediation modules are Service Component Architecture (SCA) modules that can change the format, content or target of service requests.

Mediation modules operate on messages that are in flight between service requesters and service providers. They allow you to route messages to different service providers. They also let you transform messages: you can amend message content or form. In addition, mediation modules can provide functions such as message logging, and error processing that is tailored to your requirements.

Components of mediation modules

Among the items that mediation modules contain are the following:

- Imports.
  - Imports define interactions between Service Component Architecture (SCA) modules and service providers.
  - Imports allow SCA modules to call external services as if they were local.
  - Mediation module imports can be viewed from WebSphere ESB and if the import binding is an SCA binding, then it can be modified to point to another SCA module.
- Exports.
  - Exports define interactions between Service Component Architecture (SCA) modules and service requesters.
  - Exports allow an SCA module to offer a service. Exports define the external interfaces (access points) of an SCA module.
  - Mediation module exports can be viewed from WebSphere ESB.
- Service Component Architecture (SCA) components.

- SCA components, or service components, are SCA building blocks. You build SCA modules such as mediation modules, using SCA components. You can create and customize SCA modules and components graphically, using WebSphere Integration Developer.
- Typically, mediation modules contain a specific type of SCA component called a mediation flow component. A mediation module can contain, at most, one mediation flow component.
- A mediation flow component can contain one mediation primitive, a number of mediation primitives or no mediation primitives. WebSphere ESB supports a supplied set of mediation primitives that provide functionality for message routing and transformation. One of the mediation primitives that WebSphere ESB supports allows you to invoke custom logic.
- A mediation module does not have to contain a mediation flow component. The purpose of a mediation module that does not contain a mediation flow component is to transform service requests from one protocol to another. For example, a service request might be made using SOAP/JMS but need transforming to SOAP/HTTP before sending on.

Note: You can view mediation modules from WebSphere ESB. You can also make limited changes to mediation modules from WebSphere ESB. However, you cannot view or change SCA components or mediation primitives from WebSphere ESB. Use WebSphere Integration Developer to customize SCA components and mediation primitives.

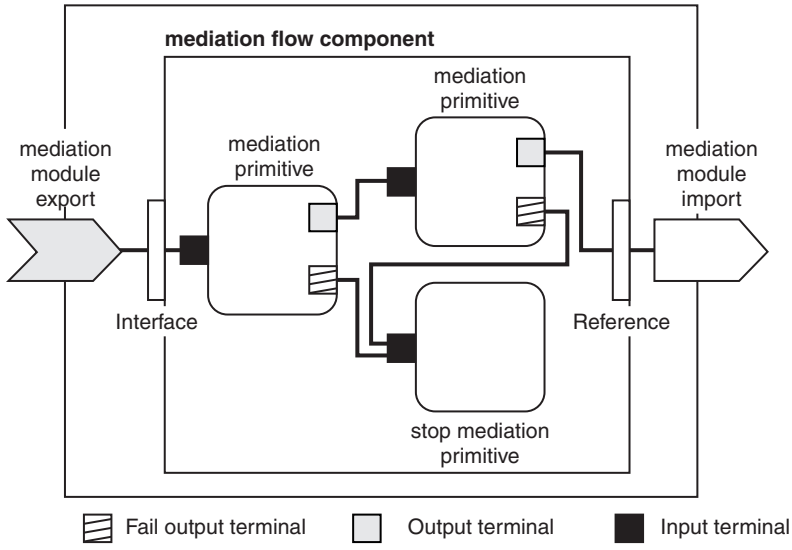

Figure 2. Simplified example of a mediation module. The mediation module contains one mediation flow component. The mediation flow component contains mediation primitives.

Deploying mediation modules

Mediation modules are created using WebSphere Integration Developer, and deployed to WebSphere ESB inside an enterprise archive (EAR) file. Therefore, a mediation module is deployed to WebSphere ESB in the same way you deploy enterprise applications.

WebSphere Integration Developer packages mediation modules inside Java archive (JAR) files, and the JAR files are then stored inside EAR files.

Logically, mediation modules can be thought of as one entity. In reality, SCA modules are defined by a number of XML files stored in one JAR file.

- EAR file.
  - Contains JAR file.
    - Contains Mediation module.

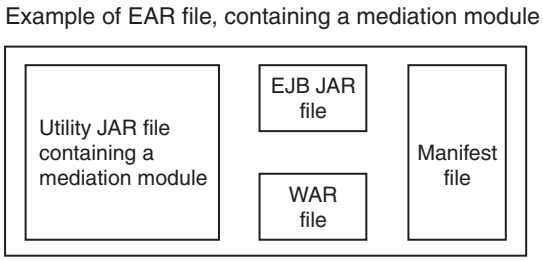

Figure 3. Simplified example of an EAR file containing a mediation module. The EAR file contains JAR files. The utility JAR file contains a mediation module.

Mediation primitives

Mediation components operate on message flows between service components. The capabilities of a mediation component are implemented by *mediation primitives*, which implement standard service implementation types.

A mediation component has one or more flows; for example, one for request and one for reply.

WebSphere ESB supports a supplied set of mediation primitives, which implement standard mediation capabilities for mediation modules deployed into WebSphere ESB. If you need special mediation capabilities, you can develop your own custom mediation primitives.

A mediation primitive defines one “in” operation that processes or handles messages that are represented by service message objects. A mediation primitive can also define “out” operations that send messages to another component or module.

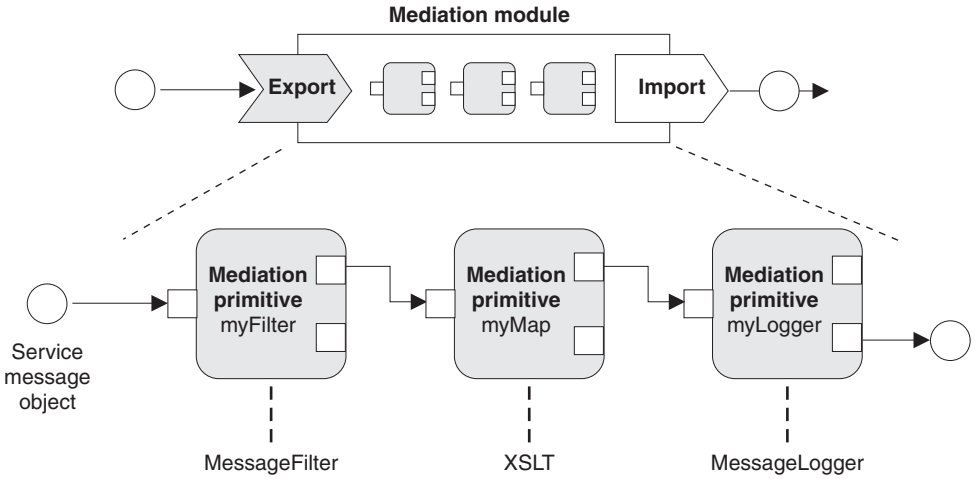

Mediation primitives typically function at the level of a single operation, with possible mediation of the request (input) and response (output). In some cases, you

can specify mediation primitives down to the level of a single parameter on an operation; for example, selectors can operate at the operation level or parameter level.

You can use WebSphere Integration Developer to graphically model and assemble mediation components from mediation primitives, and assemble mediation modules from mediation components.

The following set of mediation primitives are supported by WebSphere ESB:

Custom Mediation
Runs custom logic. The Custom Mediation primitive can call an external Service Component Architecture (SCA) component, that you provide.

- The operation that is called must be a two-way operation.
- The target SCA component must exist in the same mediation module as Custom Mediation primitive.

Database Lookup
Modifies messages, using information from a user-supplied database.

- You must set up a database, datasource and any server authentication settings for the Database Lookup mediation primitive to use.
- The Database Lookup mediation primitive can read from only one table.
- The specified key column must contain a unique value.
- The data in the value columns must be either a Java primitive or a Java String (or be able to be cast to a Java primitive or a Java String).

Fail Generates a failure in the flow.

Message Filter
Routes messages down different paths, based on the message content.

Message Logger
Logs messages in a database. The messages are stored as XML, therefore data can be post-processed by XML-aware applications.

- The database schema is defined by IBM.
- The default installation for WebSphere ESB creates a Cloudscape database and datasource. By default, WebSphere Integration Developer configures Message Logger mediation primitives to use this Cloudscape database.
- If you want to create your own database and datasource, using the administrative console, then WebSphere ESB provides data definition language (ddl) files that describe the table schema. The Table.ddl files are stored in: *install_root*/util/EsbLoggerMediation/*database_type*/Table.ddl. Where *database_ type* refers to the type of database, for example, CLOUDSCAPE_V50. If you create your own database and want to use the default JNDI name for your datasource, then you must remove the default datasource.

Stop Stops a particular path in the flow, without generating an exception.

XSLT Transforms messages.

- The XSLT mediation primitive can change the headers or the body of your messages.
- You transform messages using an XSLT (Extensible Stylesheet Transformations) 1.0 transformation. The transformation operates on an XML serialization of the message.

CEIEmitter
Emit a CEI event for messages processed ...

SplitPath
Select the target service (or another mediation), route to a specific target, modify the routing path ...

BOMapper
Select the target service (or another mediation), route to a specific target, modify the routing path ...

Service message objects

Service message objects (SMOs) are enhanced Service Data Objects (SDOs). SMO provides an abstraction layer for processing and manipulating messages exchanged between services.

SMO model

The SMO model is a pattern for using SDO DataObjects to represent messages. The SMO contains a representation of the following groups of data:

- The business payload of the message. The payload is the application data exchanged between service endpoints.
- Header information associated with the message. For example, Java Message Service (JMS) headers if a message has been conveyed using the JMS API.
- Context information (data other than the message payload).

All of this information is accessed as SDO DataObjects, and there is a schema declaration that specifies the overall structure of the SMO. The schema is generated by the WebSphere Integration Developer.

All SMOs have the same basic structure. The structure consists of a root data object called a ServiceMessageObject, which contains other data objects representing header, body and context data. The SMO body contains the message payload. The headers contain information that originates from a specific import or export binding. For example, a JMS Binding.

SMO provides an interface to access and modify message headers and message payloads. SMO can represent the logical content of many different types of message.

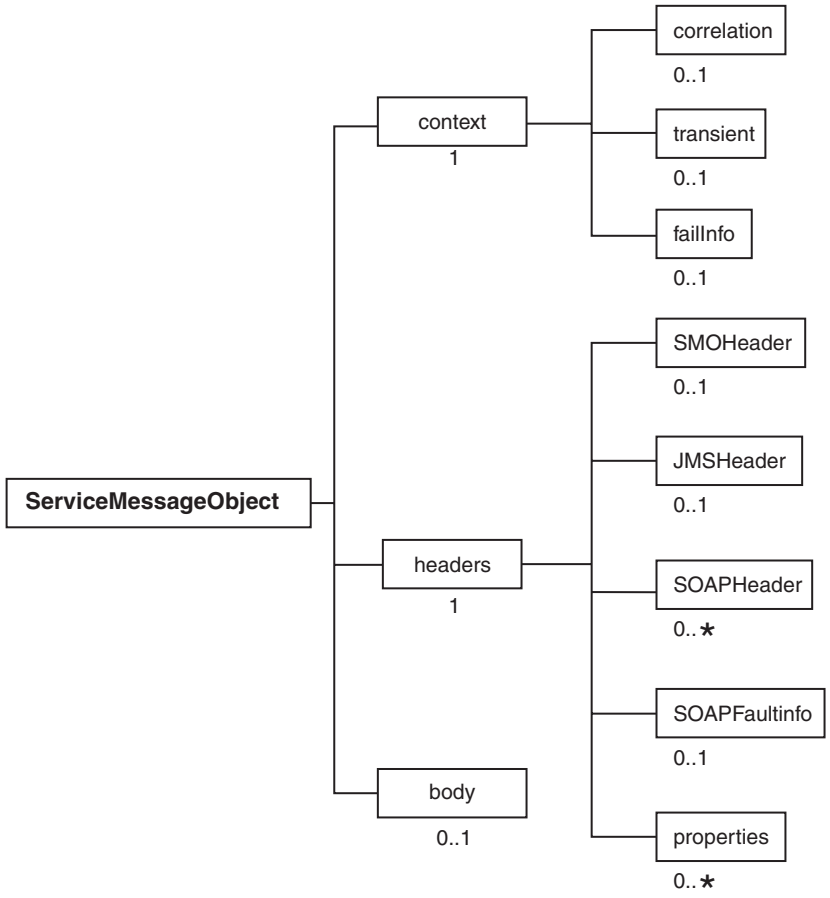

Figure 4. Overview of SMO structure. The headers, context and body of a ServiceMessageObject

WebSphere ESB and SMO

WebSphere ESB operates on messages that are in flight between interaction endpoints. Within WebSphere ESB, mediation flows process messages as SMOs.

Messages can come from a number of sources, so SMO has to be able carry different kinds of message header. The kinds of message headers handled by WebSphere ESB are:

- Web Services message headers.
- Service Component Architecture (SCA) message headers.
- Java Message Service (JMS) message headers.
- WebSphere Adapter message headers.

WebSphere ESB SMO runtime

WebSphere ESB creates SMO objects, which are then available to mediation flows.

Some of the SMO objects created by the runtime are implemented by classes supplied by the runtime. For example, the ServiceMessageObject class is supplied by WebSphere ESB. Some of the SMO header classes are also supplied by the runtime. Other SMO objects created by the runtime are implemented by classes created by a developer.

When creating mediation flows, WebSphere Integration Developer specifies the type of message body for each terminal (input, output or fail) and, optionally, the type of context information. WebSphere ESB uses this information to convert messages into SMO objects of the specified type.

Management of the server and bus environment

With WebSphere ESB, administrators can create an environment of ESB servers and service integration buses that support the deployment of service applications.

Managing the bus environment includes setting up WebSphere ESB, sometimes as part of a larger system, typically as a production environment or realistic test environment. This includes some installation and customization activities, bus topology planning, and creating product configurations. The focus is on the administration of service integration buses, servers, and how to balance workload through clustering and high availability configurations.

If you create a Complete (default) installation for WebSphere ESB, you get a stand-alone server on which you can deploy Service Component Architecture (SCA) modules without having to do any configuration of the server.

You can choose to create a WebSphere ESB deployment manager profile to enable the use of multiple servers and server clusters. You can start with one server in the deployment manager cell and optionally add capacity and enhanced availability by build up to multiple servers or server clusters.

You can add other buses if you need, to deploy applications that connect to an ESB, or to enable integration with WebSphere MQ.

The main descriptions of administrative tasks are based on use of the administrative console. Each task in the administrative console is supported by one or more panels. You can choose to apply the Server and Bus task selection filter to restrict the range of tasks to those associated with management of the bus environment. This displays panels suitable for the following task areas:

- Listing the service integration buses, servers, server clusters, messaging engines, and local topologies needed to support the deployment of mediation modules and service applications
- Enabling and disabling infrastructure services
- Installing applications and mediation modules
- Creating resources (for example, JMS connection factories and Common Event Infrastructure profiles) needed by deployed service applications and mediation modules
- Operational control of the ESB runtime

This is most appropriate for an administrator interested in managing the server and bus environment needed to support the deployment of service applications and mediation modules. This includes defining the network and bus topology, defining appropriate resources and monitoring the runtime system and troubleshooting any runtime errors.

For more comprehensive administrative tasks, you can also choose to apply the All task selection filter. This displays all administrative console panels. This is most appropriate for an administrator interested in managing all parts of WebSphere ESB and the underlying WebSphere Application Server.

Besides the administrative console, you can use commands and scripts and administrative programs to manage the server and bus environment.

Related concepts

"Overview of the bus environment" on page 14
The bus environment comprises an one or more service integration buses, ESB servers, and their resources, organized into logical administrative domains of cells and nodes.

Developing and deploying service applications

You can develop service applications that, when deployed onto WebSphere ESB, enable message-based communication between services and can operate on and manipulate messages in flight between interaction endpoints.

Developing service applications

WebSphere Integration Developer is the separate development environment for WebSphere ESB. You can use WebSphere Integration Developer to graphically model and assemble mediation components from mediation primitives, and assemble mediation modules from mediation components.

- If the interface for SCA mediation components are not imported, you can use the Simplified Interface Editor to create the interface. You can use this editor to specify and edit interfaces (operations and parameters) of mediation modules.
- You can use the Mediation Flow Editor to map between operations on the end points of a mediation, to define the set of mediation flows needed for this application. You can use a set of predefined mediation primitives to visually compose a mediation flow.
- You can use the Business Object Editor to construct the messages that are used in mediations.
- You can use other editors to extend the development environment to meet your business needs; for example:
  - Create and edit custom mediation primitives, and add them to the Mediation Flow Editor.
  - Create and edit message descriptors.

You can develop some service components using other application development tools then import them into WebSphere Integration Developer for modelling, editing, testing, and packaging for deployment into WebSphere ESB.

Deploying service applications

Deploying is the act of enabling your applications in either a test or a production environment. While the concept of deploying is the same for both environments, there are a few differences between the deployment task in each environment.

After developing a service application, best practices state that you should deploy the application onto a test server for testing before committing it to the production requirement. Use WebSphere Integration Developer to deploy the applications into a test environment, and to package a service application as a standard enterprise application package, for deployment into WebSphere ESB.

Use WebSphere ESB to install and deploy the applications into a production environment. In WebSphere ESB, you can use the standard WebSphere administrative console, with role-based administration views that simplify the

experience for solution administrators to deploy and manage the components of service integration packages.

Related concepts

"Development and deployment of integrated applications"
Options for development and deployment of integrated applications on WebSphere Process Server include working in the WebSphere Integration Developer development environment, working with Service Component Architecture APIs, and enabling the applications in a test or production server environment using WebSphere Process Server.

Development and deployment of integrated applications

Options for development and deployment of integrated applications on WebSphere Process Server include working in the WebSphere Integration Developer development environment, working with Service Component Architecture APIs, and enabling the applications in a test or production server environment using WebSphere Process Server.

IBM WebSphere Integration Developer is the development environment for WebSphere Process server. For more information about developing integrated applications in WebSphere Integration Developer, refer to the WebSphere Integration Developer information center.

In addition to the WebSphere Integration Developer development environment, Service Component Architecture APIs are published for developers. For more information about Service Component Architecture APIs, refer to the WebSphere Process Server *Developing and Deploying Modules* PDF file.

Modules, also called Service Component Architecture (SCA) modules when deployed to WebSphere Process Server, determine what artifacts are packaged in enterprise archive (EAR) files that are deployed to the runtime environment.

Within WebSphere Integration Developer, you can use an assembly editor to group services into modules and specify which services are exposed by the module to outside consumers. The modules are then connected to form complete integration solutions.

Service Component Architecture enables you to encapsulate integration logic within modules so that a change to services within a module will not affect any of the other modules in the solution as long as the interface of the changed module stays the same.

Deploying is the act of enabling your applications in either a test or a production environment. While deploying is the same for both environments, there are a few differences between the deployment task in each environment. Because it is best to test any changes to your applications on a test server before committing them to the production environment, use WebSphere Integration Developer to deploy the applications into a test environment before using WebSphere Process Server to deploy the applications into a production environment.

If you need to deploy many application files, you may want to use a batch file. For more information about batch files, refer to the WebSphere Process Server *Developing and Deploying Modules* PDF file.

Related concepts

"Developing and deploying service applications" on page 23
You can develop service applications that, when deployed onto WebSphere ESB, enable message-based communication between services and can operate on and manipulate messages in flight between interaction endpoints.

Administration of applications on WebSphere Process Server

Administering IBM WebSphere Process Server involves preparing, monitoring, and modifying the environment into which applications and resources are deployed, as well as working with the applications and resources themselves.

For more information about administering applications, refer to the *Administering* PDF file.

WebSphere Process Server offers several interfaces for administering the runtime environment:

- Administrative console

  The administrative console is a browser-based interface where you can monitor, update, stop, and start a wide variety of applications, services, and resources for the applications running on WebSphere Process Server. The administrative console can also be used to work with relationships and to locate and resolve failed WebSphere Process Server events.

  The administrative console also provide administration capabilities for WebSphere Application Server and other customer-defined products.

  For more information about the administrative console, refer to the *Administering* PDF file.

- Business Process Choreographer Explorer

  Business Process Choreographer Explorer is a stand-alone Web application that provides a basic set of administration functions for managing business process and human tasks. You can view information about process templates, process instances, task instances, and their associated objects. You can also act on these objects; for example, you can start new process instances, repair and restart failed activities, manage work items, and delete completed process instances and task instances.

- Scripting program

  The WebSphere administrative (wsadmin) scripting program is a non-graphical command interpreter environment that enables you to run administrative options in a scripting language and to submit scripting language programs for execution. It supports the same tasks as the administrative console. The wsadmin tool is intended for production environments and unattended operations.

- Command-line tools

  Command-line tools are simple programs that you run from an operating system command-line prompt to perform specific tasks. Using these tools, you can start and stop application servers, check server status, add or remove nodes, and other tasks. The WebSphere Process Server command-line tools include the serviceDeploy command, which processes .jar, .ear, .war and .rar files exported from a WebSphere Integration Developer environment and prepares them for installation to the production server.

- Administrative programs

  A set of Java classes and methods under the Java Management Extensions (JMX) specification provide support for administering Service Component Architecture (SCA) and business objects. Each programming interface includes a description

of its purpose, an example that demonstrates how to use the interface or class, and references to the individual method descriptions.

Security on WebSphere Process Server

IBM WebSphere Process Server provides a security infrastructure and mechanisms based on IBM WebSphere Application Server security.

For more information about security, refer to the WebSphere Process Server *Securing Applications and their Environment* PDF file.

Performance monitoring on WebSphere Process Server

You can monitor applications to understand response times, basic health of the systems running on the server and resource usage of applications.

IBM WebSphere Process Server monitoring capabilities include performance monitoring with the Tivoli® Performance Viewer, and event monitoring using the Common Base Events and Common Event Infrastructure.

For more information, refer to the WebSphere Process Server *Monitoring* PDF file.

Samples and tutorials

To help you learn how to accomplish your goals with IBM WebSphere Process Server, educational materials that include tutorials and samples are available.

Samples and tutorials are available from the IBM Education Assistant Web site, and tutorials for administrative tasks are available in the IBM WebSphere Process Server information center.

Tutorials

You can find tutorials using the IBM Education Assistant.

IBM Education Assistant tutorials

The IBM Education Assistant site provides tutorials that you can use at your convenience. For more information about the IBM Education Assistant, see the IBM Education Assistant Web page.

WebSphere Process Server tutorials

WebSphere Process Server documentation contains tutorial topics to assist you with some administrative, security setup, and monitoring tasks.

For tutorials on WebSphere Process Server administrative, security, and monitoring tasks, refer to the *Administering WebSphere Process Server* PDF file, the *Securing Applications and their Environment* PDF file, and the *Monitoring* PDF file.

Accessing the Samples (Samples Gallery)

Samples of integration application artifacts are available in the Samples Gallery.

The Samples Gallery contains samples of simple artifacts such as those that would be generated by IBM WebSphere Integration Developer and deployed on IBM WebSphere Process Server.

To install and view the WebSphere Process Server Samples Gallery, perform the following steps.

1. Unload the contents of the WebSphere Process Server for z/OS install media. The samples package is loaded onto the system when you unload the installation media. Run the install script to create the symlinks to the files in the samples directory

   The samples are installed in the *install_root*/samples directory.

   Applications that are run on WebSphere Process Server have XML artifacts, such as business objects, relationship definitions, and business rules, which must be deployed before installing the application. WebSphere Process Server provides a utility named serviceDeploy to build and deploy these artifacts. The enterprise archive (EAR) file in *install_root*/samples/lib for each sample application contains these artifacts. The sampleDeploy utility invokes serviceDeploy with specific parameters required for the samples. Running sampleDeploy creates a second EAR file named *sample_ name*Deployed.ear in the same directory as the original EAR file. This new EAR file contains the Web archive (WAR) files that were in the original EAR file plus the additional Java archive (JAR) and WAR files that contain the deployed artifacts. The deployed EAR file may be installed as an enterprise application in WebSphere Process Server.

   For more information, see Creating an installable EAR file using serviceDeploy in the WebSphere Process Server for z/OS *Developing and Deploying Modules* PDF file.
2. Start the server.
3. The Samples Gallery can be accessed directly from a web browser via the link http://xxxxxxxx:9080/WSsamples/en/index.html, where xxxxxxxx represents a host name or IP address of the target z/OS system. The actions of deploying and installing each sample into the server is part of the samples themselves. WebSphere Process Server samples are initially listed as installable samples in the Samples Gallery.
4. Click each sample that you want to install and follow the instructions that appear in the browser window for installation of each sample.
5. Click **Refresh** in the Samples Gallery and the samples are listed as installed samples. For each installed sample, you can select the sample name to open a browser window with additional information and an option to run the sample.
6. 6.0.1+ If you are running the samples in a distributed WebSphere Process Server environment without clustering, complete the following steps.
   a. On the machine with the deployment manager node, run the command *install_root*/samples/bin/installwbi -node *node_name* -server *server_name* -samples SamplesGallery WBISamplesGallery.
   b. In the administrative console, under **Enterprise Applications**, start the SamplesGallery and WBISamplesGallery.
   c. Open a browser to access the Samples Gallery at http://*host_name*:*host_port*/WSsamples/index.jsp.
   d. Follow the instructions in the Samples Gallery to install and run each sample, making sure to use the **-node** *node_name* **-server** *server_name* parameters with the installwbi command.
7. 6.0.1+ If you are running the samples in a distributed WebSphere Application Server environment with clustering, complete the following steps.
   a. In the administrative console, click **Applications > Install New Application**.

b. Click the browse button and locate the SamplesGallery.ear file in *install_root*/samples/lib/SamplesGallery directory.
c. Install the EAR file, accepting all defaults, except for the target mapping panel, where you can designate a server or cluster on which to install the Samples Gallery.
d. Repeat the previous steps for the WBISamplesGallery.ear file in the *install_root*/samples/lib/SamplesGallery.
e. Start the applications that you just installed.
f. Open a browser to access the Samples Gallery at http://*host_name*:*host_port*/WSsamples/index.jsp.
g. Follow the instructions in the Samples Gallery to install and run each sample, but use **Install New Application** on the administrative console instead of the installwbi command, which does not support clusters. Locate the deployed EAR files in the *install_root*/samples/lib/*sample_name* directory for each sample.

Standards compliance

This product is compliant with several government and industry standards, including accessibility standards, information processing standards, software download security standards, and internet protocol standards.

Accessibility features

Accessibility features enable users with physical disabilities, such as restricted mobility or limited vision, to operate software products successfully.

The following accessibility features are built into IBM WebSphere Process Server installation and administration features.

- Installation

  You can install WebSphere Process Server either in graphical or silent form. The silent installation program is recommended for users with accessibility needs.

  You install WebSphere Process Server for z/OS silently from a command line. This silent installation meets accessibility needs.

  For instructions, see Installing the product silently in the WebSphere Process Server *Installing* PDF file.

  For instructions, see Running the install script in the WebSphere Process Server for z/OS *Installing* PDF file.
- Administration

  The administrative console is the primary interface for interacting with the product. This console is displayed within a standard Web browser. By using an accessible Web browser, such as Microsoft® Internet Explorer or Netscape Navigator, administrators are able to:

  – Use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen
  – Use voice recognition software, such as IBM ViaVoice®, to enter data and to navigate the user interface
  – Operate features by using the keyboard instead of the mouse

You can configure and administer product features by using standard text editors and scripted or command-line interfaces instead of the graphical interfaces that are provided.

When appropriate, the documentation for specific product features contains additional information about the accessibility of the features.

Federal Information Processing Standard

Federal Information Processing Standards (FIPS) are standards and guidelines issued by the National Institute of Standards and Technology (NIST) for federal government computer systems.

IBM WebSphere Process Server relies on IBM WebSphere Application Server for all cryptographic functions, which are compliant with Federal Information Processing Standards.

FIPS are developed when there are compelling federal government requirements for standards, such as for security and interoperability, but acceptable industry standards or solutions do not exist. Government agencies and financial institutions use these standards to ensure that the products conform to specified security requirements. For more information on these standards, see the National Institute of Standards and Technology.

WebSphere Application Server integrates cryptographic modules including Java Secure Socket Extension (JSSE) and Java Cryptography Extension (JCE), which have undergone FIPS 140-2 certification. In the WebSphere Application Server documentation, the IBM JSSE and JCE modules that have undergone FIPS certification are referred to as IBMJSSEFIPS and IBMJCEFIPS.

To enable FIPS, see the **Configuring Federal Information Processing Standard Java Secure Socket Extension files** topic in the WebSphere Application Server for z/OS information center. When you enable FIPS, several components of the server are affected including the cipher suites, the cryptographic providers, the load balancer, the caching proxy, the high availability manager, and the data replication service.

Common Criteria

The National Institute of Standards and Technology (NIST) has developed Common Criteria to ensure that you have a safe option for downloading software to use on your systems.

IBM WebSphere Process Server is compliant with the Common Criteria standards developed by the National Institute of Standards and Technology.

Information held by IT products or systems is a critical resource that enables organizations to succeed in their mission. Additionally, individuals have a reasonable expectation that their personal information contained in IT products or systems remain private, be available to them as needed, and not be subject to unauthorized modification. IT products or systems should perform their functions while exercising proper control of the information to ensure it is protected against hazards such as unwanted or unwarranted dissemination, alteration, or loss. The term IT security is used to cover prevention and mitigation of these and similar hazards.

Many consumers of IT lack the knowledge, expertise or resources necessary to judge whether their confidence in the security of their IT products or systems is appropriate, and they may not wish to rely solely on the assertions of the developers. Consumers may therefore choose to increase their confidence in the

security measures of an IT product or system by ordering an analysis of its security (in other words, a security evaluation).

More information about Common Criteria can be found on the IBM Support site, including Recommended Updates for WebSphere Application Server and WebSphere Process Server.

Internet Protocol Version 6.0

IBM WebSphere Process Server relies on WebSphere Application Server for all Internet Protocol Version 6.0 compatibility.

IBM WebSphere Application Server Version 6.0 and its JavaMail component support Internet Protocol Version 6.0 (IPv6).

For more information about this compatibility in WebSphere Application Server, see IPv6 support in the WebSphere Application Server for z/OS information center.

For more information about IPv6, see www.ipv6.org.

Globalization

Globalized products can be used without language or culture barriers and can be enabled for a specific locale.

IBM WebSphere Process Server, version 6.0.0, and version 6.0.1, are available only in English but support all locales: data can be entered in other languages, but messages and interface elements are in English.

WebSphere Process Server, version 6.0.1.1, is available in other languages. For more information, refer to the WebSphere Process Server announcement letter at http://www.ibm.com/common/ssi/rep_ca/9/897/ENUS205-309/.

For information about the internationalization service available with WebSphere Application Server, refer to WebSphere extensions in the WebSphere Application Server for z/OS information center.

Bidirectional language support

The data that you process in WebSphere Process Server needs to be in the bidirectional language format of ILYNN (implicit, left-to-right, on, off, nominal), which is also the Windows® bidirectional language format. All other bidirectional language formats for applications you are running on WebSphere Process Server must be converted prior to being introduced to WebSphere Process Server.

Bidirectional attributes

There are five attributes that must be set for the proper bidirectional language format. The attributes and settings are listed in the table below.

Letter Position	Purpose	Values	Description	Default Setting
1	Order Schema	I or V	Implicit (Logical) or Visual	I
2	Orientation	L or R	Left-to-Right or Right-to-Left	L

Letter Position	Purpose	Values	Description	Default Setting
3	Symmetric Swapping	Y or N	Symmetric Swapping is on or off	Y
4	Shaping	Y or N	Text is shaped or not shaped	N
5	Numeric Shaping	H, C, N	Hindi, Contextual, Nominal	N

It is the responsibility of any client applications, external components (such as Web services, stateless session beans, and custom code), or anyone building solutions to run on WebSphere Process Server, to transform the data into the supported bidirectional language format.

For an example of bidirectional language transformation of a string, refer to Example: Using bidirectional transformation on string-type data.

For an example of bidirectional language transformation of a Service Data Object, refer to Example: Using the bidirectional transformation on DataObject-type data.

Note: The locale setting of the user interface (browser) defines the bidirectional language display and edit format.

For more information on bidirectional language, see the technical articles on IBM developerWorks, available at www.ibm.com/developerworks/websphere/library/techarticles/bidi/bidigen.html.

Example: Using the bidirectional (bi-di) transformation on DataObject-type data

When using information that is in a bidirectional language script, it might be necessary to transform the format of the data. This is a step-by-step example of the coding that performs a transformation on DataObject-type data.

6.0.0 ONLY For IBM WebSphere Process Server, version 6.0.0, if you are using JDK 1.4.1 for Windows , AIX® or Linux® operating systems, you do not need to install these classes.

The module requires that DataObject-type information is transformed from one bidirectional format to another.

Note: If you are unfamiliar with the formats, see "Values for the bidirectional (bi-di) format string" on page 34.

1. Include all bi-di classes that contain the bi-di engine implementation.
   ```
   import com.ibm.bidiTools.bdlayout.*;
   ```
2. Include all the classes you need to manipulate the DataObject-type object.
   ```
   import commonj.sdo.DataObject;
   import commonj.sdo.Type;
   import commonj.sdo.Property;
   ```
3. Define string variables to contain the different types of strings that a DataObject-type object contains. This step facilitates filtering the attributes of `type String` while transversing recursively the DataObject.

```
String STRING_STR_TYPE = "String";
String NORM_STRING_STR_TYPE = "normalizedString";
String TOKEN_STR_TYPE = "token";
String LANG_STR_TYPE = "language";
String NAME_STR_TYPE = "Name";
String NMTOKEN_STR_TYPE = "NMTOKEN";
String NCNANE_STR_TYPE = "NCName";
String ID_STR_TYPE = "ID";
String IDREF_STR_TYPE = "IDREF";
String IDREFS_STR_TYPE = "IDREFS";
String ENTITY_STR_TYPE = "ENTITY";
String ENTITIES_STR_TYPE = "ENTITIES";
```

4. Define the function that verifies if the type of a property is String.

```
private static boolean isStringFamilyType (Property property) {
    boolean rc = false;
    if ((property.getType().getName().equalsIgnoreCase(STRING_STR_TYPE)) ||
    (property.getType().getName().equalsIgnoreCase(NORM_STRING_STR_TYPE)) ||
    (property.getType().getName().equalsIgnoreCase(TOKEN_STR_TYPE)) ||
    (property.getType().getName().equalsIgnoreCase(LANG_STR_TYPE)) ||
    (property.getType().getName().equalsIgnoreCase(NAME_STR_TYPE)) ||
    (property.getType().getName().equalsIgnoreCase(NMTOKEN_STR_TYPE)) ||
    (property.getType().getName().equalsIgnoreCase(NCNANE_STR_TYPE)) ||
    (property.getType().getName().equalsIgnoreCase(ID_STR_TYPE)) ||
    (property.getType().getName().equalsIgnoreCase(IDREF_STR_TYPE)) ||
    (property.getType().getName().equalsIgnoreCase(IDREFS_STR_TYPE)) ||
    (property.getType().getName().equalsIgnoreCase(ENTITY_STR_TYPE)) ||
    (property.getType().getName().equalsIgnoreCase(ENTITIES_STR_TYPE)))
       rc = true;
     return rc;
}
```

5. Define the recursive function that applies the bi-di transformation on the entire DataObject.

   Note: The following explain basic assumptions that code logic follows.

   - bi-di transformation is applied on properties of string type only
   - the properties of type string in the DataObject are stored in one bi-di format

```
DataObject BiDiDataObjTransformationBO(DataObject boIn, String formatIn, String formatOut){

   Type type;
   Property property;

   if (boIn == null) return null;

   type = boIn.getType();
   List propertyList = type.getProperties();
   for (int propertyNumber = 0; propertyNumber < propertyList.size(); propertyNumber++){
        property = (Property) propertyList.get(propertyNumber);
        String propertyName = property.getName();
```

   a. Skip all non-string properties.

```
          if (!isStringFamilyType(property))
                continue;

          if (property.isContainment()) {
               if (property.isMany()) {
                     List childsList = boIn.getList(property);
```

   b. Recursively call the transformation to handle children objects.

```
                  for (int childNumber = 0; childNumber < childsList.size();
      childNumber++){                 BiDiDataObjTransformationBO(connectionContext,
    ((DataObject)childsList.get(childNumber)),formatIn, formatOut);
                }
           } else {
```

   c. Recursively call the transformation to handle children objects of any contained business objects.

```
                BiDiDataObjTransformationBO(connectionContext,
          ((DataObject)boIn.get(property)),formatIn, formatOut);
             }
          } else {
```

d. Transform the simple string attributes.

```
             String str = BiDiStringTransformation(
                 (boIn.getString(propertyName),formatIn, formatOut);
             boIn.setString(propertyName, str);
          }
       }
       return boIn;
    }
```

Example: Using bidirectional (bi-di) transformation on string-type data

When using information that is in a bidirectional language script, it might be necessary to transform the format of the data. This is a step-by-step example of the coding that transforms string-type data.

6.0.0 ONLY For IBM WebSphere Process Server, version 6.0.0, make sure that all the classes that contain the bi-di engine implementation are installed on the server on which you are developing your modules. If you are using JDK 1.4.1 for Windows , AIX or Linux operating systems, you will not have to install these classes.

The module requires that string information is transformed from one bidirectional format to another.

Note: If you are unfamiliar with the formats, see "Values for the bidirectional (bi-di) format string" on page 34.

1. Include all bi-di classes that contain the bi-di engine implementation.

   ```
   import com.ibm.bidiTools.bdlayout.*;
   ```

2. Define the strings to contain the data object to transform, and the input and output format values.

   The input format is the bi-di format in which the string object is currently stored. The output format is the bi-di format in which you want to store the string object.

   ```
   String strIn = new String("Hello world");
   String formatIn = "ILYNN";
   String formatOut = "VLYNN";
   ```

3. Call the BidiStringTransformation function.

   ```
   String strOut = BiDiStringTransformation(strIn, formatIn, formatOut);
   String BiDiStringTransformation(String strIn, String formatIn, String formatOut) {
   ```

   a. Test if input string is null.

      ```
      if (strIn == null) return null;
      ```

   b. Perform transformation

      ```
      BidiFlagSet flagsIn;
      BidiFlagSet flagsOut;
      formatIn = formatIn.toUpperCase();
      formatOut = formatOut.toUpperCase();

      if (formatIn != null)
          flagsIn = new BidiFlagSet(formatIn.toCharArray());
      else
          flagsIn = new BidiFlagSet();
      ```

```
    if (formatOut != null)
        flagsOut = new BidiFlagSet(formatOut.toCharArray());
    else
        flagsOut = new BidiFlagSet();

    if (flagsIn.equals(flagsOut)) return strIn;
    String strOut = BiDiStringTransformation(strIn, flagsIn, flagsOut);
    return strOut;
}
```

Values for the bidirectional (bi-di) format string

The values for the bidirectional (bi-di) language format string control the transformation of bidirectional scripts from one format to another.

Purpose

Use the table below to determine the correct value for either the input string or output string to use when transforming String-type or DataObject-type data from one format to another.

Table 1. Bidirectional language format string values

Letter position	Purpose	Allowable values	Default value	Meaning
1	Ordering schema	I	I	Implicit
		V		Visual
2	Orientation	L	L	Left to right
		R		Right to left
		C		Contextual left to right
		D		Contextual right to left
3	Symmetric swapping	Y	Y	Symmetrical swapping is on
		N		Symmetrical swapping is off
4	Shaping	S	N	Text is shaped
		N		Text is not shaped
		I		Initial shaping
		M		Middle shaping
		F		Final shaping
		B		Isolated shaping
5	Numeric	H	N	Hindi (National)
		C		Contextual
		N		Nominal

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

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

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
577 Airport Blvd., Suite 800
Burlingame, CA 94010
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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 and cannot confirm the accuracy of 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.

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

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming interface information

Programming interface information, if provided, is intended to help you create application software using this program.

General-use programming interfaces allow you to write application software that obtain the services of this program's tools.

However, this information may also contain diagnosis, modification, and tuning information. Diagnosis, modification and tuning information is provided to help you debug your application software.

Warning: Do not use this diagnosis, modification, and tuning information as a programming interface because it is subject to change.

Trademarks and service marks

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries, or both: IBM, IBM (logo), AIX, CICS, Cloudscape, DB2, DB2 Connect, DB2 Universal Database, developerWorks, Domino, IMS, Informix, iSeries, Lotus, MQSeries, MVS, OS/390, Passport Advantage, pSeries, Rational, Redbooks, Tivoli, WebSphere, z/OS, zSeries

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

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

This product includes software developed by the Eclipse Project (http://www.eclipse.org/).

IBM Websphere Process Server for z/OS version 6.0.1

Printed in USA