



IBM Software Group

WebSphere Process Server V6.2
WebSphere Enterprise Service Bus V6.2
WebSphere Integration Developer V6.2

Service Component Architecture overview



@business on demand.

© 2009 IBM Corporation
Updated June 25, 2009

This presentation will provide an overview of Service Component Architecture (SCA).

Goals

- Introduce Service Component Architecture (SCA) as it applies to:
 - ▶ WebSphere® Process Server
 - ▶ WebSphere Enterprise Service Bus
 - ▶ WebSphere Integration Developer



The goal of this presentation is to introduce Service Component Architecture, or SCA. The presentation specifically addresses the implementation of SCA that is provided by WebSphere Process Server and WebSphere Enterprise Service Bus and is enabled through the WebSphere Integration Developer tool. There are places in the presentation that reference WebSphere Process Server, but although not specifically stated, the material is equally applicable to WebSphere Enterprise Service Bus.

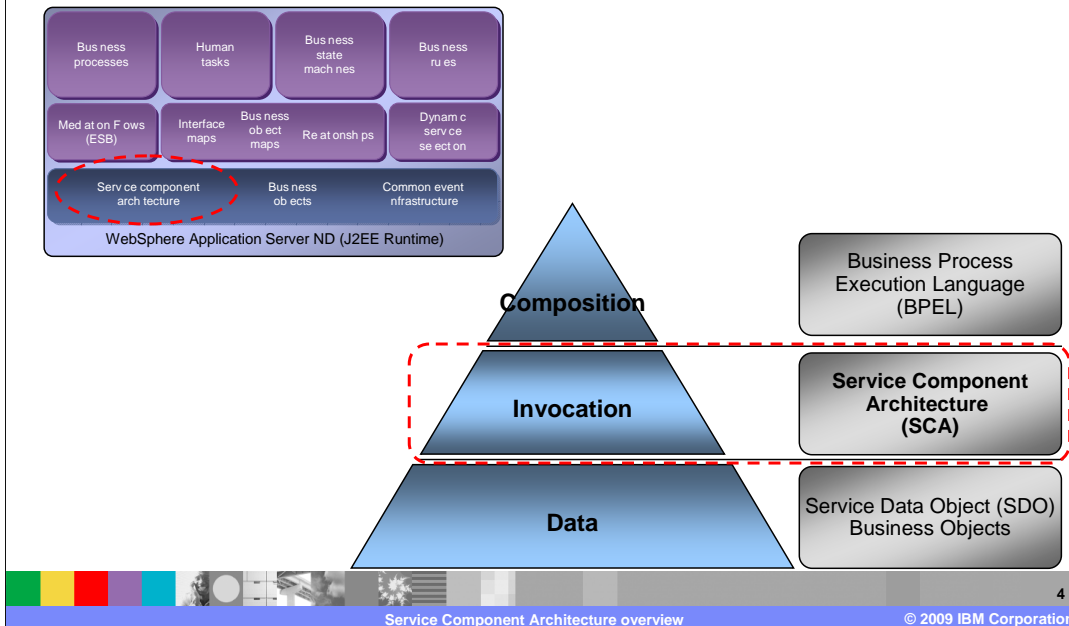
Agenda

- Overview
- Architecture
- Summary and references



This section will provide an overview of Service Component Architecture (SCA).

Service Oriented Architecture overview



This is a simple way to look at the important architectural constructs that make up a service oriented architecture. Specifically, there must be a way to represent the data that is exchanged between services, a mechanism for invoking services, and a way to compose services into larger integrated business applications.

Today there are many different programming models for supporting each of these. This situation presents developers with the challenge of not only solving a particular business problem, but also choosing and understanding the appropriate implementation technology. One of the important goals of the WebSphere Process Server SOA solution is to mitigate these complexities. This is done by converging the various programming models used for implementing service oriented business applications into a simplified programming model.

This presentation focuses specifically on the Service Component Architecture (SCA) in WebSphere Process Server as the service oriented component model for defining and invoking business services. SCA plays an important role in providing an invocation model for the SOA solution in WebSphere Process Server. You will also learn in this presentation that it plays a role in composing business services into composite business applications.

Service Component Architecture description

- SCA is a service oriented component model for defining business services that publish or operate on business data
- SCA provides a single abstraction for service types that might already be expressed as other types
 - ▶ Session bean, Web service, Java™ class, BPEL
- Separates “business logic” from “infrastructure logic”



SCA is a service oriented component model for defining and invoking business services that publish or operate on business data. SCA is aimed at providing a simplified programming model for writing applications that run in a J2EE runtime environment, and is based upon concepts and techniques that are refinements of existing J2EE technology. One of the important aspects of SCA is to enable the separation between application business logic and the implementation details. In order to accomplish this, SCA provides a single abstraction for service types that might already be expressed as session beans, Web services, Java classes, or BPEL. The ability to separate business logic from infrastructure logic is important to help reduce the IT resources needed to build an enterprise application, and give developers more time to work on solving a particular business problem rather than focusing on the details of which implementation technology to use.

Section

Architecture



This section will provide the architectural details of SCA.

Service Component Architecture features

- Provides the Service Component Definition Language (SCDL) for defining service components
- Provides the ability to:
 - ▶ Define service components
 - ▶ Make services available to clients outside current module
 - ▶ Import and reference external services in current module
 - ▶ Compose services into larger application components
- Provides a client programming model allowing client access to service components



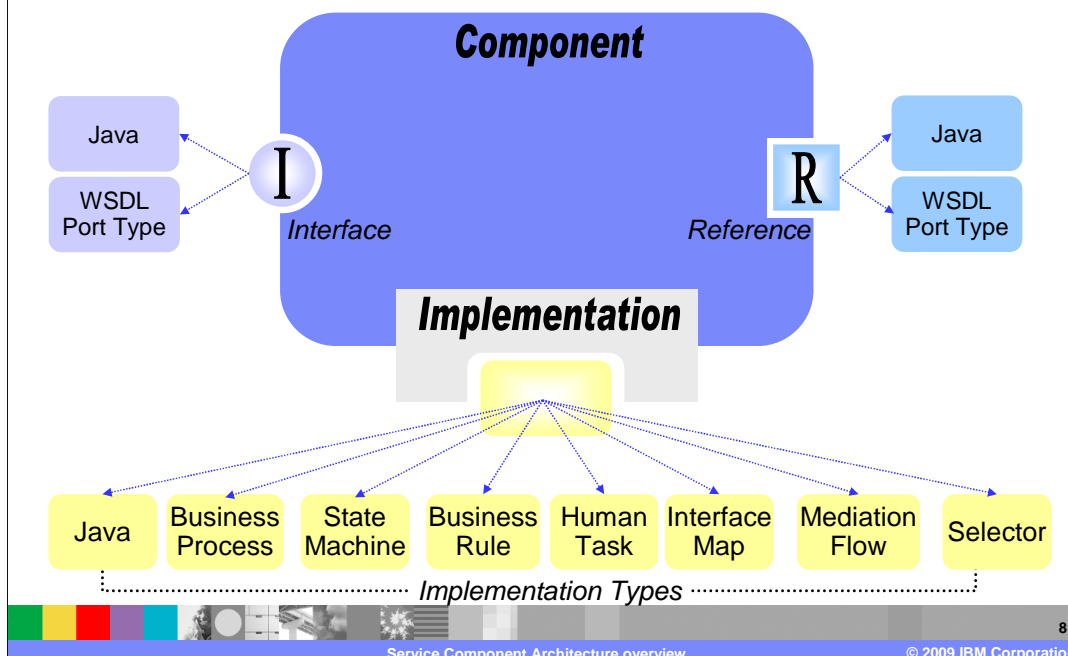
When learning a new technology or programming model, it is often useful to look at the pieces that compose the overall architecture of that technology. This slide lists some of the important features of SCA that you should be aware of as you begin learning about SCA.

First, the Service Component Definition Language (SCDL) provides the basis of SCA. SCDL is an XML based definition language used to define all SCA artifacts in a project. The WebSphere Integration Developer support of SCA takes care of generating the appropriate SCDL definitions when building an SCA-based application. However, a basic familiarity with SCDL can certainly help you to understand the overall architecture and help when debugging applications.

The next important part of SCA is the different types of artifacts that can be defined using SCDL. The various artifact types that exist in SCA were designed to support some of the basic requirements of this service oriented architecture. To start with, SCA needs a mechanism for defining a basic service component. Once there is a mechanism for defining service components, it is important to have the ability to make these services available to clients both inside or outside of the current SCA module. In addition to this, a construct designed to import and reference services external to the current SCA module must exist. Finally, SCA provides constructs for composing services and modules into larger applications. In the remaining slides of this section, you will learn about each of these SCA artifacts and how they can be composed into larger applications.

One final feature of SCA is the client programming model that allows clients to access and invoke service components in an SCA module. Later in this presentation, you will learn about the Java interfaces that are available for invoking services from a client.

Service component overview



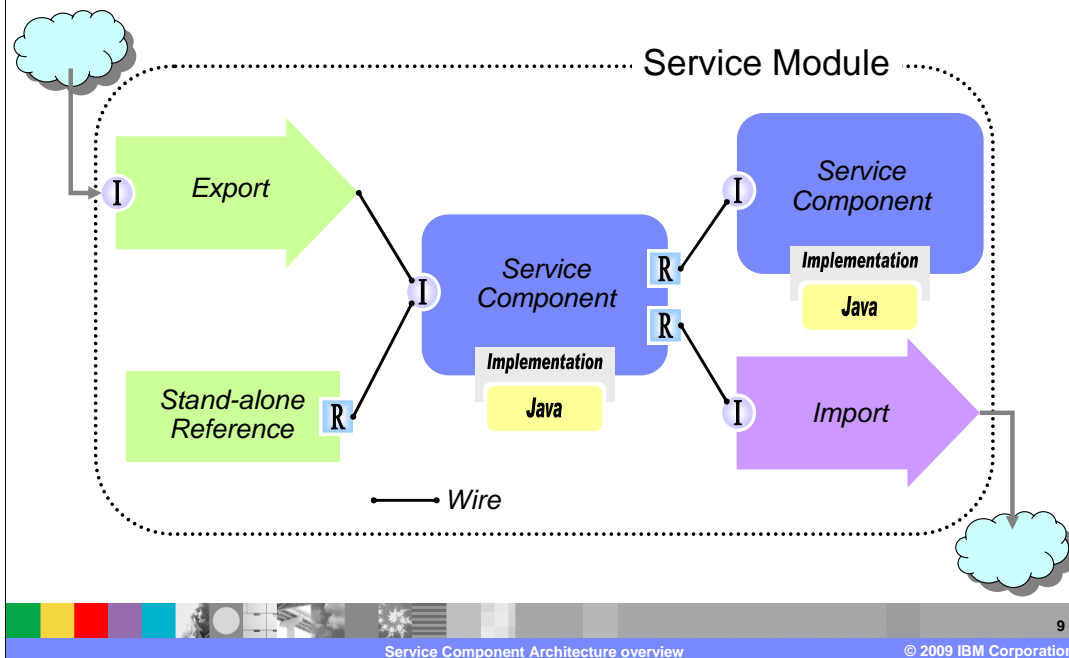
The basic building block in SCA is the service component. The service component represents a business service that publishes or operates on business data. The diagram on this slide introduces the essential pieces of a service component definition.

A service component has one or more interfaces with which it is associated. The interfaces associated with a service component advertise the business operations associated with this service. These interfaces can be specified as either Java interfaces or WSDL port type interfaces. However, you can not mix Java and WSDL port type interfaces on the same service component definition. The arguments and return types for these interfaces are specified as simple Java types, Java classes, Service Data Objects, or XML Schema (for WSDL port type interfaces).

Also associated with a service component definition is an implementation. As the diagram indicates, there are multiple language types available for implementing a service component. This presentation will primarily focus on the Java implementation type. However other presentations are available to discuss the details of the other implementation types that are available.

Each service component can access other services in their implementation. For this, a service component definition can include zero or more references to other service components or imports included in the current module.

Service module overview



The previous slide introduced the service component as the basic building block in SCA. This slide provides a broader look at SCA and the other pieces that make up the architecture.

This discussion begins with the service module, which provides the basic unit of deployment and administration in an SCA-enabled runtime. A service module encapsulates the various artifacts available with SCA and is illustrated in the diagram on this slide. The following is a summary of the elements that make up a service module.

A service module may have zero or more service components included with it. In order to access these services by a client (SCA or non-SCA) there must exist at least one reference to the service or the service must be exposed with an export.

A service module can have zero or more imports included with it. An import is used to access services that are outside the current SCA module. Once an import has been defined, other services from within the module can reference the imported service as if it was a regular service component defined in the module.

A service module can have zero or more exports included with it. An export is used to expose a particular service to clients outside the current SCA module.

A service may include a stand-alone references file that includes references to services in the module that can be used by SCA and non-SCA services.

Service module artifacts

Artifact	Comments
Module Definition	<ul style="list-style-type: none"> ▪ Contained in the sca.module file at the root SCA project JAR
Service Components	<ul style="list-style-type: none"> ▪ A module can contain 0..n service definitions ▪ Each component definition is contained in a <SERVICE_NAME>.component file
Imports	<ul style="list-style-type: none"> ▪ A module can contain 0..n import definitions ▪ Each import definition is contained in a <IMPORT_NAME>.import file
Exports	<ul style="list-style-type: none"> ▪ A module can contain 0..n export definitions ▪ Each export definition is contained in a <EXPORT_NAME>.export file
References	<ul style="list-style-type: none"> ▪ Two types of References <ul style="list-style-type: none"> ▶ In-line (contained within a service component definition) ▶ Stand-alone ▪ Stand-alone references are defined in the sca.references file
Other Artifacts	<ul style="list-style-type: none"> ▪ Other artifacts include: Java Classes, WSDL files, XSD files, BPEL.

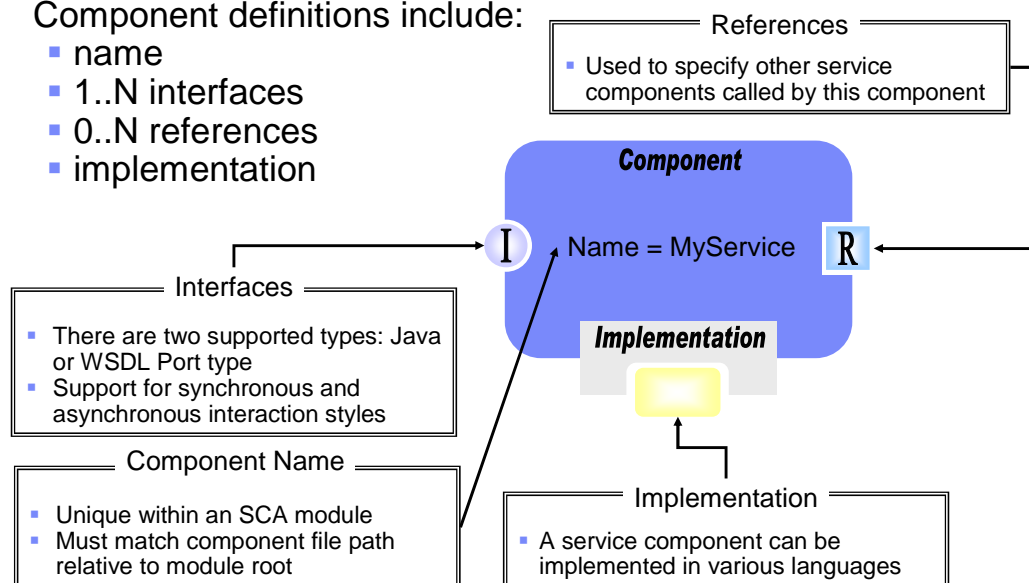


The table on this slide includes the primary artifacts that make up an SCA service module. Listed in the comments column is the name of the file for the artifact that includes the SCDL definition for that particular artifact type.

Service component definition

Component definitions include:

- name
- 1..N interfaces
- 0..N references
- implementation

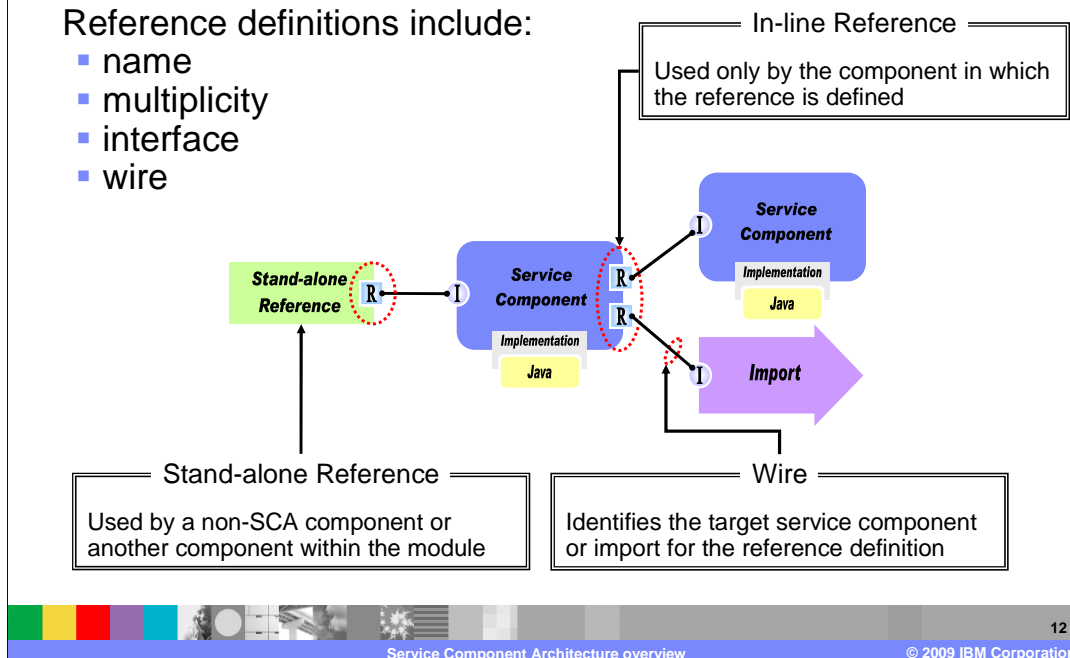


This slide provides a more detailed look at the service component definition introduced earlier in this presentation. Each service component must have a unique name within the SCA module and it must match the file path relative to the module root. As noted on the previous slide, the service component definition is included in a file called {SERVICE NAME}.component. Next, each service component can have one or more interfaces associated with it, which can be either Java or WSDL port type interface definitions. The interfaces associated with a service component can support either a synchronous or asynchronous interaction style with clients calling the service. This feature is discussed in more detail in upcoming slides in this presentation. As noted earlier, each service component can be implemented in various ways, specified by the implementation definition. Finally, service components can invoke other service components or imports defined in the current service module. In this case, the appropriate reference must be defined to indicate which service is used. Often this type of reference is in-lined in the service component definition, although it may alternatively be placed in the stand-alone references file. Each service component definition can have zero or more references to other services called by the service component being defined.

References and wires

Reference definitions include:

- name
- multiplicity
- interface
- wire



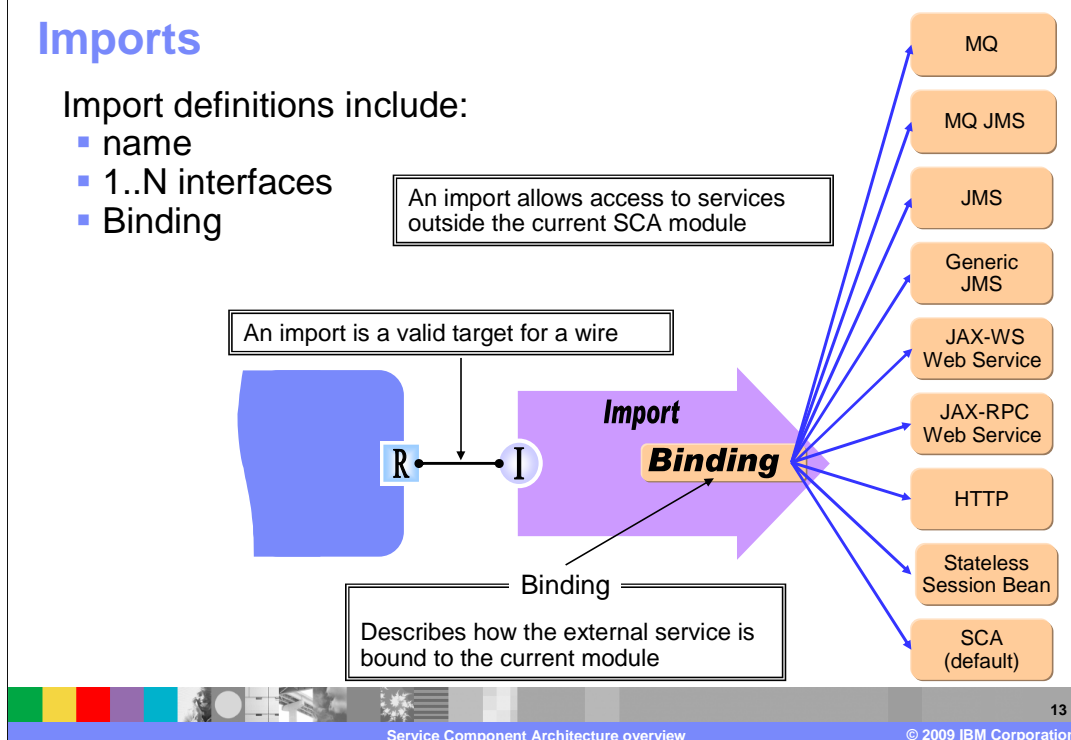
SCA and non-SCA clients calling a service component need a reference to that service in order to invoke it. This slide reveals some of the details of defining references. Each reference has a name, used to look up the appropriate service by a client using the client programming model. Details of this look up code will be covered later in this presentation. In addition to the name, a reference also includes an interface element. The multiplicity for a reference indicates how many wire definitions can name this reference as the source. Finally, the wire definition specifies the name of the target service component or Import that will resolve the reference.

There are two ways to define references. The first way is to in-line the reference in the service component definition. Using this approach, the references are only available to the service component in which the references are included. Another approach is to include reference definitions within the stand-alone references file. For this approach, the references can be used by a non-SCA client or by another component within the module. An example of a non-SCA component that may use a reference in the stand-alone references file is a user interface component such as a JSP that needs the ability to invoke a particular service. In order to invoke a service component, the client needs a reference so that it can use the SCA runtime to lookup the appropriate service to invoke.

Imports

Import definitions include:

- name
- 1..N interfaces
- Binding



SCA imports allow clients in an SCA module to access services that are outside the current SCA module. Like service components, imports have a name and a set of from one to N interfaces with which they are associated. Imports also have a binding attribute, which is used to describe how the external service is bound to the current module. The common binding types are indicated on this slide.

There is a group of binding types that are referred to as the messaging bindings. One of these is the MQ binding, which enables clients to access a service using WebSphere MQ native protocols, including access to MQ headers and message payloads. The remaining messaging bindings are all based on JMS. The MQ JMS binding provides access to services using the WebSphere MQ JMS provider. The JMS binding makes use of the JMS default messaging provider which uses the systems integration bus that is built into WebSphere Application Server. The last of the messaging bindings is the generic JMS binding which can be used with any JMS provider that is compliant with the JMS 1.1 specification, such as SonicMQ.

Web service bindings allow clients to access external Web services using the SCA programming model. Support is included for JAX-WS using SOAP 1.1 or 1.2 over HTTP. Support is also included for JAX-RPC using SOAP 1.1 over either HTTP or JMS.

The HTTP bindings provide access to HTTP based applications. These are different from the Web services bindings in that the payload does not have to be SOAP, but can be any format. Also, unlike the Web services bindings, the HTTP bindings provide access to the HTTP headers.

Stateless session bean bindings enable clients, using the SCA programming model, to access services that have been exposed using a stateless session bean.

Finally, there is the SCA binding. Imports with SCA bindings allow clients to access SCA services that reside in another SCA module and have an export which also has an SCA binding type. Because all types of bindings are a part of SCA, this binding is often referred to as the SCA default binding.

There is also an EIS binding type that is used in association with adapter components. However, a discussion of this binding type is outside the scope of this presentation.

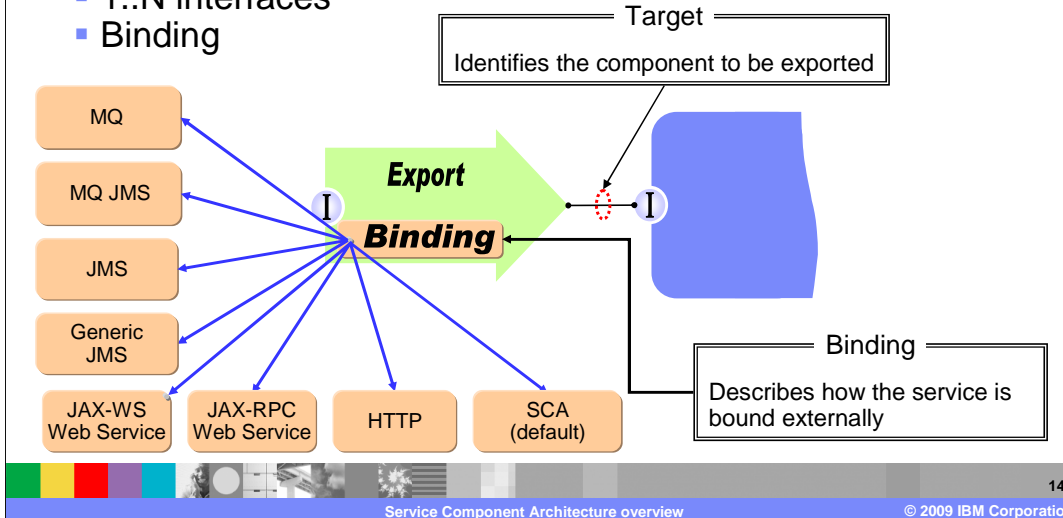
Imports can be thought of as a special type of service component in an SCA module. Imports are valid targets in a wire definition for a service reference. This means that to a client invoking a target service the client programming model is the same whether the reference points to an import or a component.

Exports

Export definitions include:

- name
- target
- 1..N interfaces
- Binding

An export allows access to services for use outside the current SCA module



SCA exports provide access to service components defined in an SCA module for use by clients outside of the current SCA module. Exports include a name and a target attribute, which names the service component that is to be exported. Like imports, exports have a binding attribute that indicates how the service is bound externally. The common binding types are indicated on this slide.

The binding types for exports are the same as the binding types for imports, with the exception of the stateless session bean binding, which is only available on imports.

You can see the messaging binding types on the left side of the slide. The MQ binding enables clients using WebSphere MQ native protocols to invoke SCA services, with the binding providing access to the MQ headers and message payloads. The MQ JMS binding enables clients using the WebSphere MQ JMS provider to invoke SCA services and likewise the JMS binding enables clients using the JMS default messaging provider to invoke SCA services. Clients using other JMS 1.1 compliant JMS providers make use of the generic JMS binding in order to access SCA services.

The Web service bindings type allows SCA services to be exported and made available to external clients as a Web service. The external clients can use JAX-WS with SOAP 1.1 or 1.2 over HTTP to invoke the SCA service. They can also use JAX-RPC with SOAP 1.1 over HTTP or JMS.

The HTTP bindings enable HTTP based applications to access SCA services. Similar to how they are used with imports, the payload does not have to be SOAP but can be any format and access to the HTTP headers is provided.

The SCA default binding type allows SCA services to be exported to other SCA clients in modules external to the current SCA module. This binding type is used in conjunction with a corresponding import with an SCA binding type in another SCA module.

There is also an EIS binding type for exports that is used in association with adapter components. Information on this binding type is addressed in the presentation on the various WebSphere adapters.

Client programming model

- Client programming model allows clients to
 - ▶ Locate services
 - ▶ Invoke methods on services
- Clients locate services with the ServiceManager
 - ▶ Key class is
 - `com.ibm.websphere.sca.ServiceManager`
 - ▶ Two ways to instantiate a ServiceManager depending on required lookup scope for service
 - ▶ Method to locate a service
 - `com.ibm.websphere.sca.Service locateService(String);`

The SCA client programming model provides two primary functions for clients. The programming model exposes an interface that allows clients to locate services within the current module, and once a service is located the client programming model provides a way for the client to invoke operations on that service.

The key interface that clients should be aware of for locating services is `com.ibm.websphere.sca.ServiceManager`. This interface includes a `locateService` method that returns a reference to the service implementation for the service requested. The string parameter that is passed into the `locateService` method represents the reference name for the service that the client wants to locate. The Java documentation for the SCA programming model is included in the WebSphere Process Server information center, and is also included if you choose to install the Java documentation as part of the WebSphere Process Server installation.

Client programming model (continued)

- Two service invocation models

Dynamic Invocation

```
Service myService = (Service) serviceManager.locateService("myService");
DataObject input = ...
DataObject result = (DataObject) myService.invoke("someMethod", input);
```

Type Safe Invocation

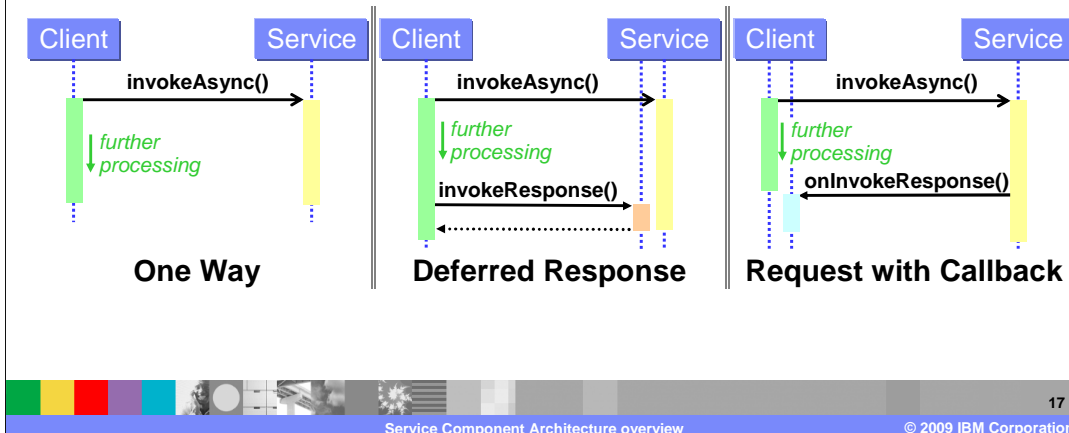
```
MyService myService = (MyService) serviceManager.locateService("myService");
String input = ...
String result = myService.someMethod(input);
```

```
public Interface MyService {
    public String someMethod(String input);
}
```

Once a client has located the appropriate service, there are two types of invocation models that can be used to make a call to an operation or method offered by the service. First, there is a dynamic invocation style of interaction. The key interface for this style of interaction is `com.ibm.websphere.sca.Service`. The `invoke()` method on this interface takes the name of the operation that you are going to invoke, along with the parameters needed to call that operation. Clients can also use a type safe invocation method to call a particular operation associated with a service. This type of invocation only works for interface definitions that are specified as Java. In this situation, the client casts the return from the `locateService()` call to the appropriate interface and can proceed calling the appropriate type safe method calls on that interface.

Asynchronous model

- SCA provides the ability for services to be called synchronously or asynchronously
- There are three types of asynchronous invocation models



So far, this presentation has focused on the synchronous invocation model. However, SCA provides the ability for services to be called either synchronously or asynchronously. The next several slides will present more about the asynchronous programming model.

With asynchronous invocation in SCA, there are three types of asynchronous interaction styles available. With all three types of asynchronous invocation, the client receives control back immediately from the SCA runtime upon an `invokeAsync()` call. However, there are three different ways that the client can capture the response at a later time. First, the client can choose to discard the response entirely or if it is a call to a void method. In this case, the asynchronous invocation is said to be “one way”. Another option is for the client to call `invokeAsync()` and then continue processing until some later time when the client makes a request to capture the response. This scenario is termed “deferred response”. Finally, the client also has the option of doing an asynchronous “request with callback”. To do this, the client must first implement the `ServiceCallback` interface. Then, after making a call to `invokeAsync()`, the SCA runtime provides a callback to the `ServiceCallback` handler to provide the response to the client.

The next slide provides a summary of the interfaces needed to support these three asynchronous invocation models for the dynamic interaction style.

Dynamic client invocation

Interface	Methods	Description
Service	Object invoke(String, Object)	Used to invoke synchronous service requests
	Ticket invokeAsync(String, Object)	Used to invoke one-way or deferred response asynchronous service requests
	Ticket invokeAsyncWithCallback(String, Object)	Used to invoke request with callback asynchronous service requests. The client must implement the ServiceCallback interface
	Object invokeResponse(Ticket, long)	Used to get response in the case of deferred response invocation
ServiceCallback	void onInvokeResponse(Ticket, Object, Exception)	Callback interface must be implemented by the client using a request with callback asynchronous service invocation

This slide provides a summary of some of the key methods and interfaces needed to support both synchronous and asynchronous interaction when using dynamic client invocation. The Java documentation for these interfaces is included with the WebSphere Process Server installation and in the product information center.

SCA interactions

Interface Type	Invocation Model				Invocation Methods	
	Synchronous	One Way	Deferred Response	Request with Callback	Dynamic	Type Safe
WSDL Port Type	●	■	■	■	YES	NO
Java Interface	●	■	■	■	YES	YES

● Data passed by reference in the same SCA Module

■ Data passed by value

The table shown on this slide lists the various invocation models and the method by which the data is passed, whether by reference or by value. For synchronous invocation, data is passed by reference within the same SCA module, while for asynchronous calls the data is passed by value. The table on this slide also summarizes when it is possible to use either type safe or dynamic invocation based upon the interface type. The dynamic invocation methods are always available for either WSDL port type or Java interfaces. However, in order to have type safe invocation methods available to the client a Java interface type must be used for the interface definition on the appropriate client reference.

SCA quality of service

- Qualifiers are used to specify quality of service requirements on the SCA runtime
- In WebSphere Integration Developer, the qualifiers are grouped into these categories
 - ▶ Reliability
 - ▶ Activity session
 - ▶ Security
 - ▶ Other asynchronous
 - ▶ Miscellaneous



Qualifiers are an important part of SCA because they allow you to place quality of service requirements on the SCA runtime. There are several different categories into which SCA qualifiers are grouped. The qualifier categories are reliability, activity session, security, other asynchronous and miscellaneous. In WebSphere Integration Developer, the qualifiers are presented to you in these groupings.

SCA quality of service

- Qualifiers are specified at various levels (scopes)
- Interface level qualifiers
 - ▶ Apply to components, often to imports and occasionally to exports
 - ▶ Are scoped to
 - All interfaces
 - A single interface
 - An individual method on an interface
- Reference level qualifiers
 - ▶ Are scoped to
 - All references
 - A single reference
- Implementation level qualifiers



Each SCA qualifier has a particular level, or scope, where the qualifier is specified. Some qualifiers are specified at the interface level. All SCA interface level qualifiers apply to SCA components, many apply to imports and one also applies to exports. These interface level qualifiers are further specifiable at different scopes within the supported interfaces for the SCA component or import. They can apply to all supported interfaces, and they can also be specified with a narrower scope applying only to a specific interface or to just a specific method.

Other qualifiers are specified at the reference level. Reference qualifiers apply only to SCA components and can be scoped to all references defined on a component, or only for a specific reference.

Finally, there are some qualifiers specified at the component implementation level.

The next slides provide tables showing the various qualifiers that are available and the level at which each is specified. The qualifiers are sorted by the type of quality of service they provide, such as reliability or security.

Qualifiers

Type	Qualifier	Scope	Description
Reliability	Transaction	Implementation	global – A global transaction must be present to run the component local – A global transaction must not exist to run the component any – Component is unaffected by transactional state
	Join transaction	Interface (component or import)	true – Hosting container joins client transaction false – Hosting container will not join client transaction
	Suspend transaction	Reference	true – Synchronous invocations of target component do not run within client global transaction. false – Synchronous invocations of target component run within client global transaction
	Asynchronous invocation	Reference	call – Asynchronous invocations of a target service occur immediately commit – Asynchronous invocations of a target service occur as part of a global transaction
	Reliability	Reference	Specifies the quality of service level for asynchronous message delivery. Reliability can be one of these values: best effort or assured

22

Reliability qualifiers allow you to request a particular transactional environment for the components and imports in an SCA module. They all relate to controlling transactions, asynchronous invocation and asynchronous reliability.

The **Transaction** qualifier is set at the implementation scope of a component. This qualifier can be set to either 'global', 'local', or 'any', with local being the default. When set to global, the component will run in the context of a global transaction. If a global transaction is present on invocation, the component is added to this global transaction scope. If set to local, the component will run in the context of a local transaction. Finally, if the value is set to any, when a global transaction is present the component will join the current global transaction. However, if a global transaction is not present, the component will run in the context of a local transaction.

The **Join transaction** qualifier is set at the interface scope of a component or import. This qualifier can be set to either true or false, false being the default. If set to true, it instructs the runtime not to suspend a global transaction, if present, at the interface boundary. If set to false, it instructs the runtime to suspend a global transaction, if present, at the interface boundary. Exposing the join transaction qualifier on an interface provides metadata that can be used by assemblers and deployers to ensure that the assembled application behaves as required. It is up to the assembler and deployer in addition to dynamic runtimes to reason about whether a target component will federate with a propagated transaction.

The **Suspend transaction** qualifier is set at the reference level of a component and identifies whether a global transaction should be suspended before invoking the target service associated with the reference. This qualifier can be set to either true or false, with the default being false.

The **Asynchronous invocation** qualifier is similar to the suspend transaction qualifier, except that it pertains to asynchronous interactions rather than synchronous types, as is the case with suspend transaction. The asynchronous invocation qualifier can have the value of call or commit, with call being the default. If set to call, it indicates to the runtime that the message for the asynchronous interaction should be committed to the queue immediately when the call has been made. Alternatively, the value of commit indicates that the message should be committed to the queue as part of a transaction associated with the current unit of work.

The **reliability** qualifier is used to specify the quality of service level for asynchronous message delivery. The reliability can be set to either best effort or assured, which is the default.

Qualifiers (continued)

Type	Qualifier	Scope	Description
Activity session	Activity session	Implementation	<p>true – There must be an activity session established in order to run this component</p> <p>false – The component runs under no Activity Session</p> <p>any – The component is agnostic to the presence or absence of an activity session</p>
	Join activity session	Interface (component or import)	<p>true – Hosting container joins client activity session</p> <p>false – Hosting container will not join client activity session</p>
	Suspend activity session	Reference	<p>true – Methods on target component will NOT run as part of any client activity session</p> <p>false – Methods on target component will run as part of any client activity session</p>

The set of activity session qualifiers are similar to the reliability qualifiers for transactions introduced earlier. The activity session service is a Websphere Application Server programming model extension that provides an alternative unit of work when compared with global transactions. In fact, an activity session context can be longer lived than a global transaction and can even include global transactions. Here is a summary of the activity session qualifiers:

The **activity session** qualifier is specified at the implementation level and is used to indicate whether an activity session should or should not exist in order to run the service component with which it is associated. This qualifier can be set to either 'true', 'false', or 'any', with any being the default. If set to true, it indicates that the component will run as part of an activity session. If set to false, the component should not run as part of an activity session. Finally, if this qualifier is set to any, the component will run as part of an activity session if it is present, otherwise it will not.

The **join activity session** qualifier is set at the interface level, and indicates whether the component should join the activity session of the caller. There are two values for this qualifier, true and false, with false being the default. If set to true it indicates that the runtime should not suspend an activity session if present when the component is invoked. If set to false it indicates that an activity session should be suspended before invoking the component.

The **suspend activity session** qualifier is set at the reference level and is used to indicate whether a target service associated with a reference will get called as part of the calling activity session. If set to true, the activity session is suspended and the methods on the target component will not run as part of the client activity session. If set to false the activity session is not suspended and methods on the target component will run as part of the client activity session. False is the default.

Qualifiers (continued)

Type	Qualifier	Scope	Description
Security	Security identity	Implementation	Specifies a logical name for the identity under which the implementation executes at runtime.
	Security permission	Interface (component only)	The caller identity must have the role specified from this qualifier in order to have permission to run the interface or method
Other asynchronous	Request expiration	Reference	Specifies the length of time (milliseconds) after which an asynchronous request is to be discarded if not delivered
	Response expiration	Reference	Specifies the duration (milliseconds) between the time a request is sent and the time a response or callback is received

There are two qualifiers available for indicating quality of service related to security.

The **security identity** qualifier is used to specify the security identity under which the implementation for the service component should run at runtime. This qualifier must be placed at the implementation scope for the service component and the value given must match the logical name for the identity under which the component will run.

The **security permission** qualifier is specified at the interface level for components. The value for this qualifier indicates that a caller of this service must have the role that is specified in order to invoke the service.

For both the security permission and the security identity, the underlying implementation for these qualifiers is based on existing J2EE concepts.

In addition to asynchronous reliability, which was already covered, there are two more qualifiers available for controlling asynchronous request and response. Each of these asynchronous qualifiers are specified at the reference scope.

The **request expiration** qualifier is used to specify the length of time the runtime should hold onto an asynchronous request if it has not yet been delivered. After the time indicated for this qualifier, given in milliseconds, this request is discarded.

The **response expiration** qualifier is used to specify the length of time that the runtime must retain an asynchronous response or must provide a callback. The value for this qualifier is also given in milliseconds.

Qualifiers (continued)

Type	Qualifier	Scope	Description
Miscellaneous	Data validation	Interface (component, import and export)	Confirms that the data passed in to an operation matches the XSD types of the operation's inputs. Log error and continue – Errors logged and requested operation is performed. Throw exception – Errors result in exception and requested operation is not performed.
	Event sequencing	Interface (component only and at method scope only)	Controls the order in which the runtime environment processes events. You specify one or more operations and a key. Events for those operations with the same key are processed in the order they are received.

25

There are two additional qualifiers that don't fall into any of the preceding categories.

The **Data validation** qualifier is specified at the interface scope and can be used with components, imports and exports. This is the only qualifier that is applicable to exports. When data validation is specified, the input business object instances are checked to see if the data they contain conforms to the XSD type definition. For those cases where the input data does not conform, there are two options that can be used with this qualifier. The first is log error and continue which requests that a log be written but that no exception be thrown and that processing continues. The other option is called throw exception, which does just that. When data which is not valid is discovered an exception is thrown and the operation is not performed.

The **Event sequencing** qualifier is specified at the interface scope, but in a very limited way. It is only applicable to components and can only be specified at the individual method level. This qualifier controls the order in which events are processed by the component. This qualifier is specified on one or more operations along with a key. Any events which contain the same key arriving for the specified operations are processed in order. This qualifier is only valid when used with a WebSphere Process Server runtime, and has no affect in an WebSphere Enterprise Service Bus runtime.

Section

Summary

This section will provide a summary of service component architecture.

Summary

- SCA is the fundamental component model
 - ▶ Programming model for a Service Oriented Architecture
 - ▶ Used by WebSphere Process Server and WebSphere Enterprise Service Bus
- SCA separates business and implementation logic
 - ▶ Focus is on assembling solutions rather than implementation details
 - ▶ Mitigates need for integration developers to have deep knowledge of Java or J2EE
 - ▶ Aimed at helping J2EE developers become more productive



SCA is the fundamental component model for WebSphere Process Server and WebSphere Enterprise Service Bus and provides the basis of the service oriented architecture solution. SCA helps separate business logic from implementation logic and allows you to focus on assembling solutions rather than focusing on implementation details. This enables integration developers to develop applications without having a deep knowledge of J2EE and at the same time enables J2EE developers to be more productive.

Feedback

Your feedback is valuable

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

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

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_WBPMv62_SCA_Overview.ppt

This module is also available in PDF format at: [../WBPMv62_SCA_Overview.pdf](..WBPMv62_SCA_Overview.pdf)



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

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both: WebSphere

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of other IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>

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

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

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

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

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

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

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

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2009. All rights reserved.

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