

# Service-Oriented Architecture

Slides are based on the book:

“Service-Oriented Architecture: Concepts, Technology, and Design”

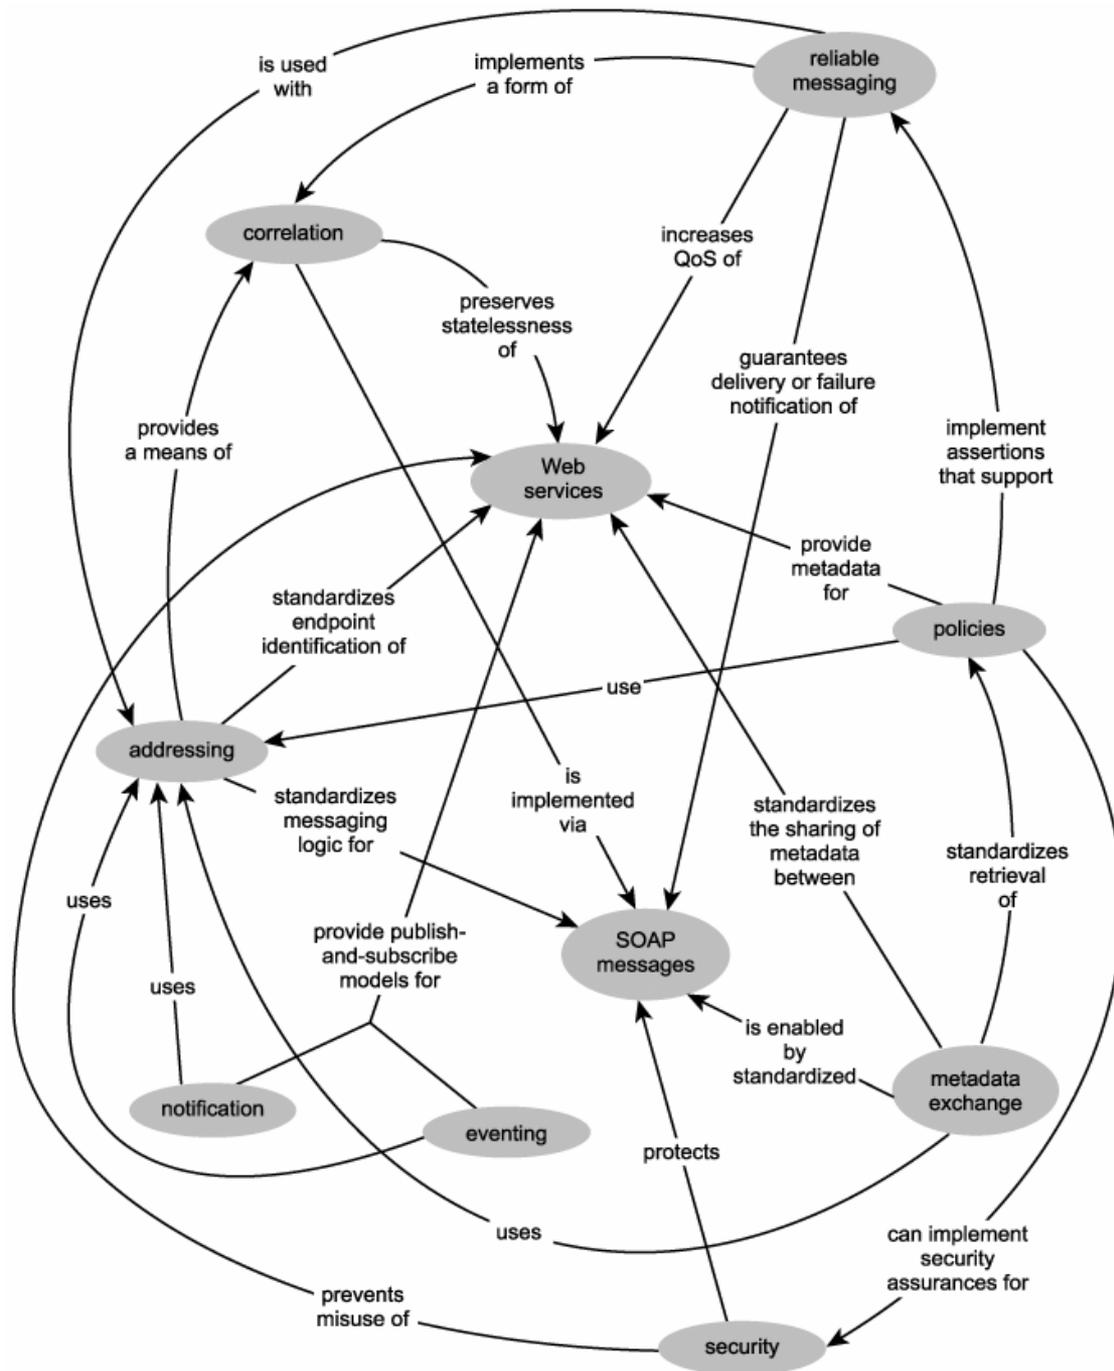
By Thomas Erl,

Publisher: Prentice Hall PTR

# **Web Services and Contemporary SOA**

(Part II: Advanced Messaging, Metadata, and Security)

- Additional WS-\* extensions that govern specific areas of the
  - SOAP messaging framework,
  - the creation and exchange of metadata, and
  - the introduction of message-level security.



# Current Web services specifications

- WS-Addressing
- WS-ReliableMessaging
- WS-Policy Framework (including WS-PolicyAttachments and WS-PolicyAssertions)
- WS-MetadataExchange
- WS-Security (including XML-Encryption, XML-Signature, and SAML)
- WS-Notification Framework (including WS-BaseNotification, WS-Topics, and WS-BrokeredNotification)
- WS-Eventing

# Addressing

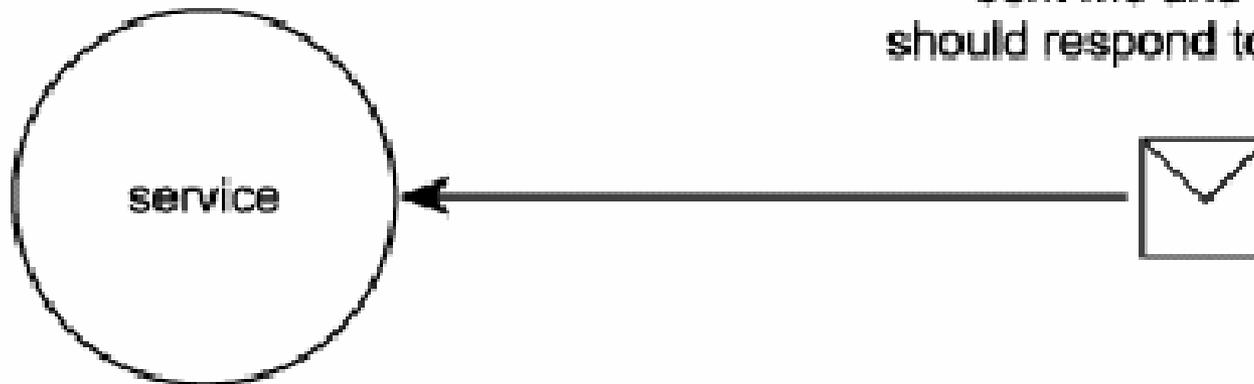
# Addressing

- What addressing brings to SOAP messaging is much like what a waybill brings to the shipping process.
- Regardless of which ports, warehouses, or delivery stations a package passes through en route to its ultimate destination, with a waybill attached to it, everyone it comes into contact with knows:
  - where it's coming from
  - the address of where it's supposed to go
  - the specific person at the address who is supposed to receive it
  - where it should go if it can't be delivered as planned

# Addressing

- The WS-Addressing specification implements these addressing features by providing two types of SOAP headers (explained shortly).
- Though relatively simple in nature, these addressing extensions are integral to SOA's underlying messaging mechanics.
- Many other WS-\* specifications implicitly rely on the use of WS-Addressing.

"I can tell you exactly who sent me and how you should respond to my arrival."



"And I know exactly where to send that information."

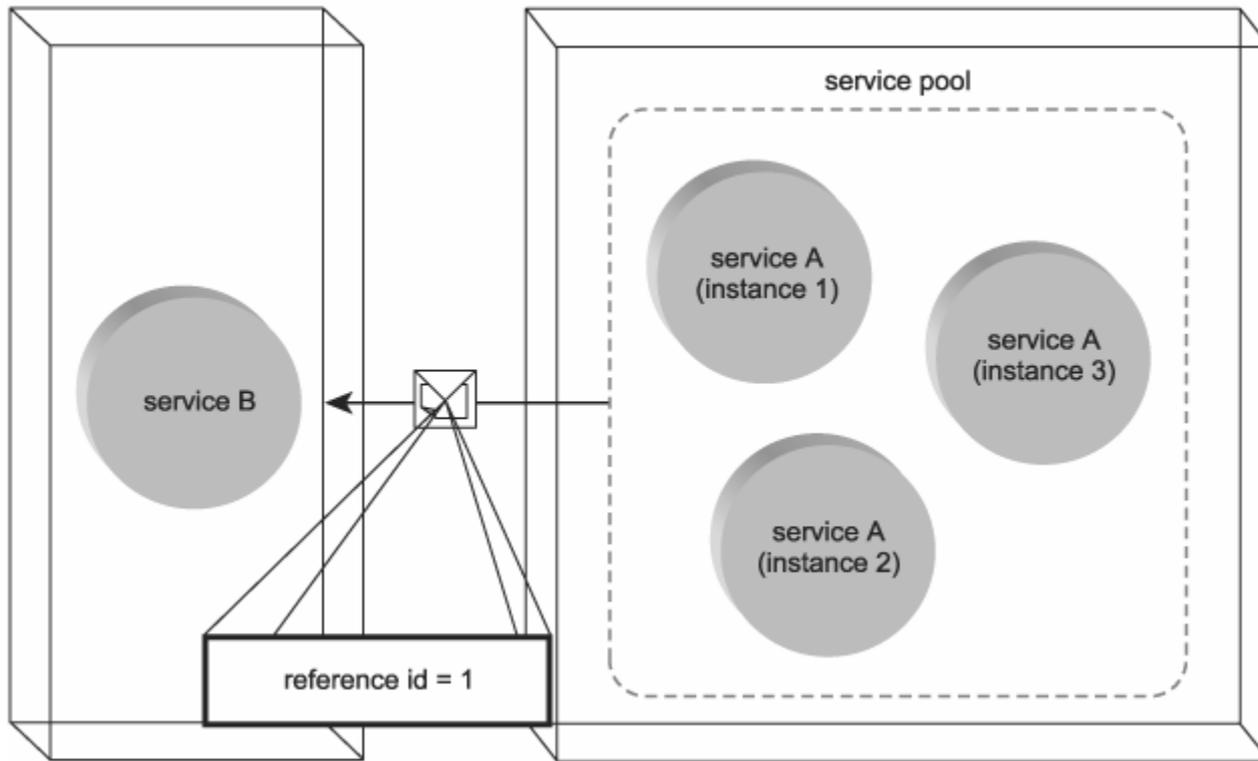
# Endpoint references

- The loosely coupled nature of SOA was implemented through the use of service descriptions. In other words, all that is required for a service requestor to contact a service provider is the provider's WSDL definition.
- This document, among other things, supplies the requestor with an address at which the provider can be contacted.
- What if, though, the service requestor needs to send a message to a specific instance of a service provider? In this case, the address provided by the WSDL is not sufficient.
- Traditional Web applications had different ways of managing and communicating session identifiers. The most common approach was to append the identifier as a query string parameter to the end of a URL.
- While easy to develop, this technique resulted in application designs that lacked security and were non-standardized.

# Endpoint references

- The concept of addressing introduces the endpoint reference, an extension used primarily to provide identifiers that pinpoint a particular instance of a service (as well as supplementary service metadata).
- The endpoint reference is expected to be almost always dynamically generated and can contain a set of supplementary properties.

A SOAP message containing a reference to the instance of the service that sent it



# Endpoint references

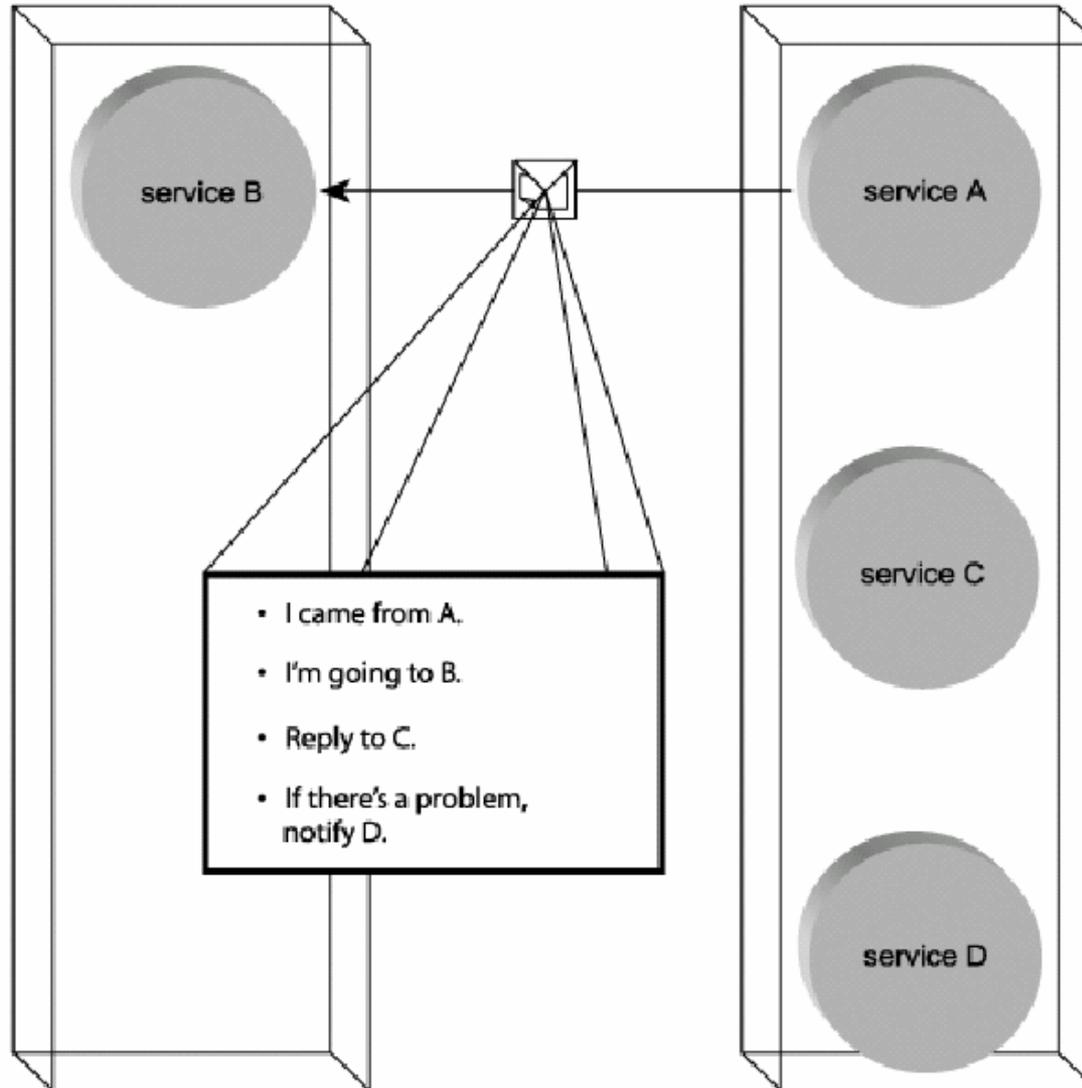
An endpoint reference consists of the following parts:

- **Address:** The URL of the Web service.
- **reference properties:** A set of property values associated with the Web service instance.
- **reference parameters:** A set of parameter values that can be used to further interact with a specific service instance.
- **service port type and port type:** Specific service interface information giving the recipient of the message the exact location of service description details required for a reply.
- **Policy:** A WS-Policy compliant policy that provides rules and behavior information relevant to the current service interaction

# Message information headers

- Various primitive message exchange patterns are used in complex activities.
- These MEPs have predictable characteristics that can ease the manner in which Web services are designed but also can limit the service interaction scenarios within which they participate.
- In sophisticated service-oriented solutions, services often require the flexibility to break a fixed pattern. For example, they may want to dynamically determine the nature of a message exchange.
- The extensions provided by WS-Addressing were broadened to include new SOAP headers that establish message exchange-related characteristics within the messages themselves.
- This collection of standardized headers is known as the message information (or MI) headers

A SOAP message with message information headers specifying exactly how the recipient service should respond to its arrival



# Message information headers

- The MI headers provided by WS-Addressing include:
  - **destination** The address to which the message is being sent.
  - **source endpoint** An endpoint reference to the Web service that generated the message.
  - **reply endpoint** This important header allows a message to dictate to which address its reply should be sent.
  - **fault endpoint** Further extending the messaging flexibility is this header, which gives a message the ability to set the address to which a fault notification should be sent.
  - **message id** A value that uniquely identifies the message or the retransmission of the message (this header is required when using the reply endpoint header).
  - **relationship** Most commonly used in request-response scenarios, this header contains the message id of the related message to which a message is replying (this header also is required within the reply message).
  - **action** A URI value that indicates the message's overall purpose (the equivalent of the standard SOAP HTTP action value).

# Message information headers

- Outfitting a SOAP message with these headers further increases its position as an independent unit of communication.
- Using MI headers, SOAP messages now can contain detailed information that defines the messaging interaction behavior of the service in receipt of the message.
- The net result is standardized support for the use of unpredictable and highly flexible message exchanges, dynamically creatable and therefore adaptive and responsive to runtime conditions.

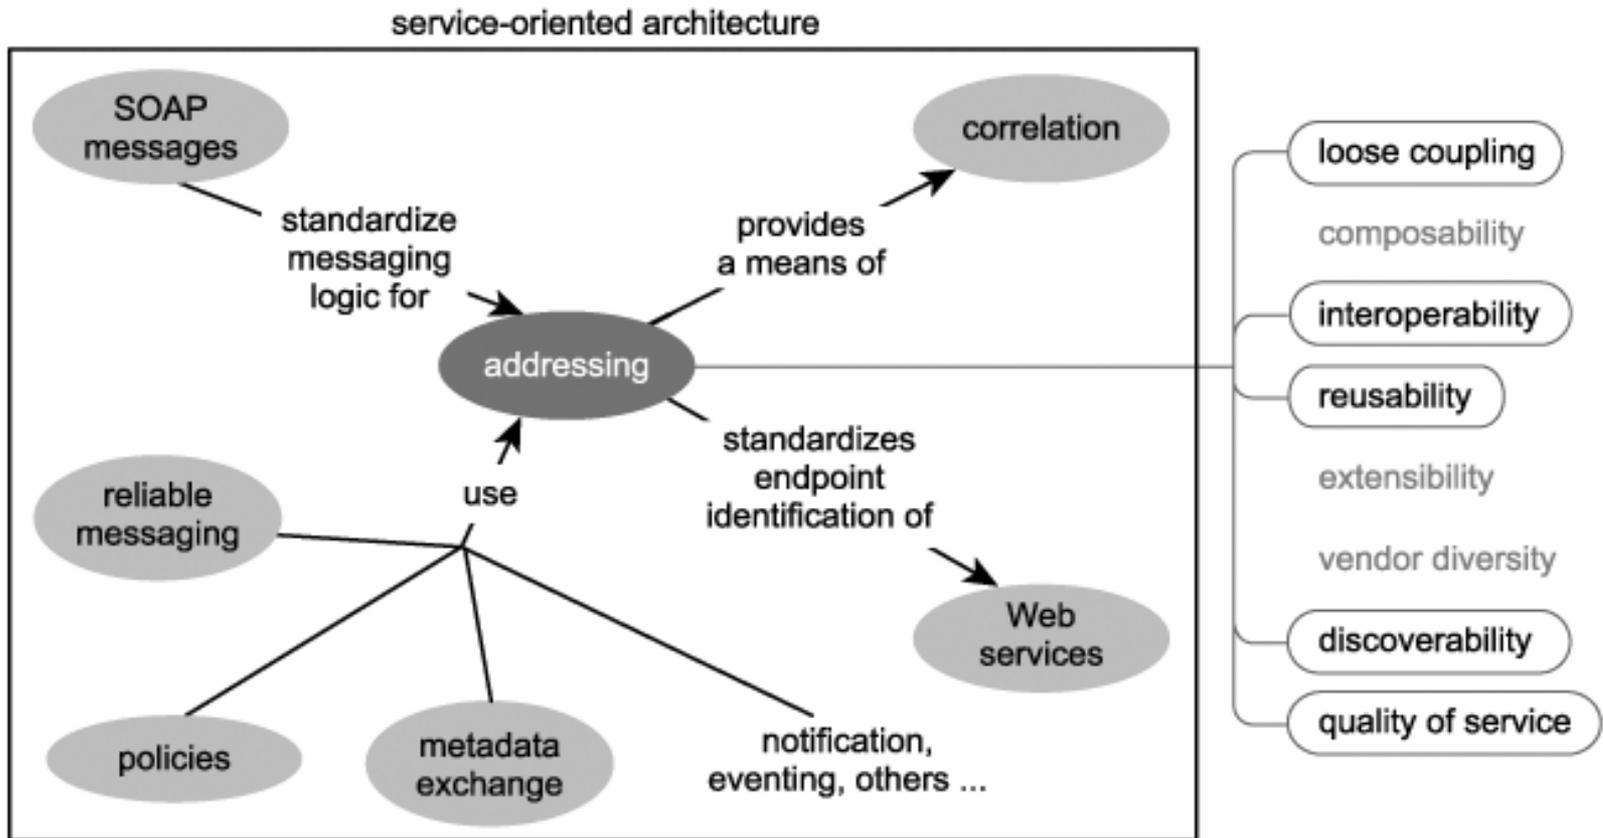
# Addressing and transport protocol independence

- Historically, many of the details pertaining to how a unit of communication arrives at point B after it is transmitted from point A was left up to the individual protocols that controlled the transportation layer.
- While this level of technology-based abstraction is convenient for developers, it also leads to restrictions as to how communication between two units of processing logic can be achieved.
- The standardized SOAP headers introduced by WS-Addressing remove much of this protocol-level dependence.
- These headers put the SOAP message itself in charge of its own destiny by further increasing its ability to act as a standalone unit of communication.

# Addressing and SOA

- Addressing achieves an important low-level, transport standardization within SOA, further promoting open standards that establish a level of transport technology independence .
- The use of endpoint references and MI headers deepens the intelligence embedded into SOAP messages, increasing message-level autonomy.

# Addressing relating to other parts of SOA



# Addressing and SOA

- Empowering a message with the ability to self-direct its payload, as well as the ability to dictate how services receiving the message should behave, significantly increases the potential for Web services to be intrinsically interoperable.
- It places the task-specific logic into the message and promotes a highly reusable and generic service design standard that also facilitates the discovery of additional service metadata.
- Further, the use of MI headers increases the range of interaction logic within complex activities and even encourages this logic to be dynamically determined. This, however, can be a double-edged sword.
- Even though MI headers can further increase the sophistication of service-oriented applications, their misuse (or overuse) can lead to some wildly creative and complex service activities.

# Addressing and SOA

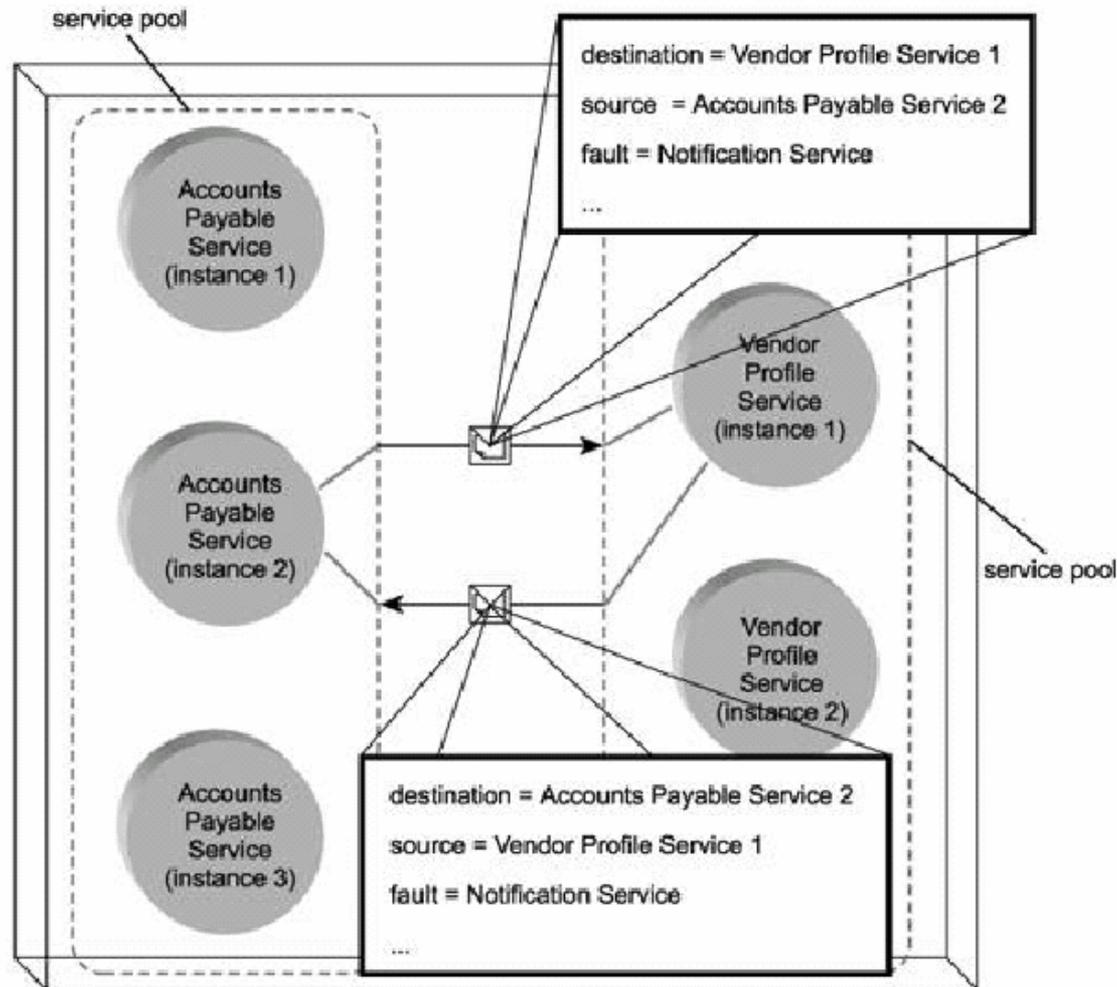
- Finally, by supporting the referencing of service instances, SOAs can be scaled in a standardized manner, without the need to resort to custom or proprietary application designs (scalability is a key QoS contribution).
- Having stated that, it should be pointed out that by providing functionality that enables communication with service instances, WS-Addressing indirectly supports the creation of stateful services.
- This runs contrary to the common service-orientation principle of statelessness and emphasizes the need for this feature to be applied in moderation.

# Case Study

The vendor invoice submission process

- One of these steps required that, upon receiving an invoice from a vendor, the TLS Accounts Payable Service interacts with the TLS Vendor Profile Service to have the received invoice validated against vendor account information already on file.
- Due to the volume of invoice submissions received by TLS, there can be, at any given time, multiple active instances of the Accounts Payable Service. Therefore, as part of the message issued by the Accounts Payable Service to the Vendor Profile Service, a SOAP header providing a reference id is included.
- This identifier represents the current instance of the Accounts Payable Service and is used by the Vendor Profile Service to locate this instance when it is ready to respond with the validation information.

# Separate service instances communicating using endpoint references and MI headers across two pools of Web services within TLS



# Summary of Key Points

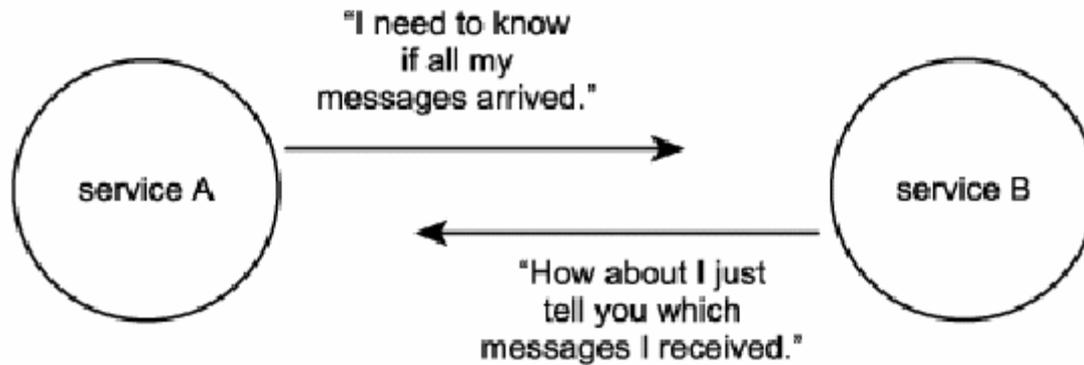
- Addressing extensions, as implemented by the WS-Addressing specification, introduce two important concepts:
  - endpoint references and
  - message information headers.
- Endpoint references provide a standardized means of identifying a specific instance of a Web service.
- Message information headers add message exchange properties to a specific message, conveying interaction semantics to recipient services.
- Though simple in comparison to other WS-\* specifications, WS-Addressing inserts a powerful layer of messaging autonomy within a service-oriented architecture.

# Reliable messaging

# Reliable messaging

- The benefits of a loosely coupled messaging framework come at the cost of a loss of control over the actual communications process.
- After a Web service transmits a message, it has no immediate way of knowing:
  - whether the message successfully arrived at its intended destination
  - whether the message failed to arrive and therefore requires a retransmission
  - whether a series of messages arrived in the sequence they were intended to
- Reliable messaging addresses these concerns by establishing a measure of quality assurance that can be applied to other activity management frameworks

Reliable messaging provides a guaranteed notification of delivery success or failure



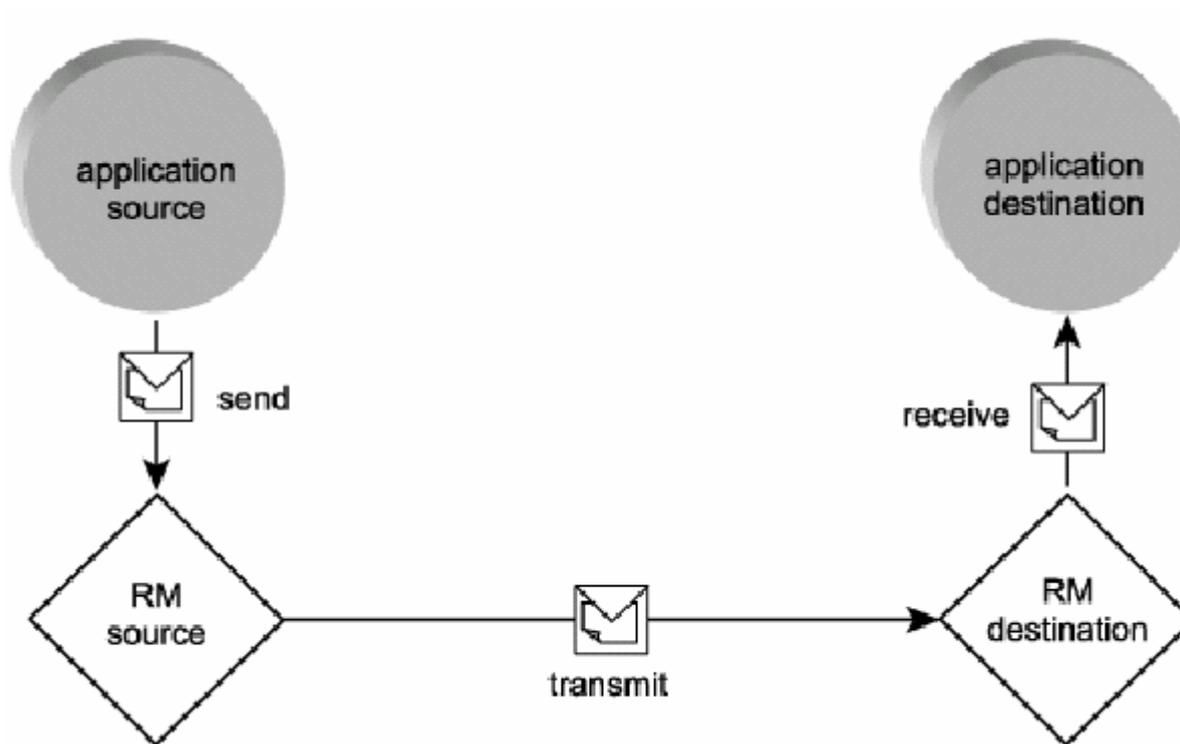
# Reliable messaging

- WS-ReliableMessaging provides a framework capable of guaranteeing:
  - that service providers will be notified of the success or failure of message transmissions
  - that messages sent with specific sequence-related rules will arrive as intended (or generate a failure condition)
- Although the extensions introduced by reliable messaging govern aspects of service activities, the WS-ReliableMessaging specification is different from the activity management specifications.
- Reliable messaging does not employ a coordinator service to keep track of the state of an activity; instead, all reliability rules are implemented as SOAP headers within the messages themselves.

# RM Source, RM Destination, Application Source, and Application Destination

- WS-ReliableMessaging makes a distinction between the parts of a solution that are responsible for initiating a message transmission and those that actually perform the transmission.
- It further assigns specific descriptions to the terms "send," "transmit," "receive," and "deliver," as they relate differently to these solution parts. These differentiations are necessary to abstract the reliable messaging framework from the overall SOA.
- An application source is the service or application logic that sends the message to the RM source, the physical processor or node that performs the actual wire transmission.
- Similarly, the RM destination represents the target processor or node that receives the message and subsequently delivers it to the application destination.

An application source, RM source, RM destination, and application destination



# RM Source, RM Destination, Application Source, and Application Destination

## Sequences

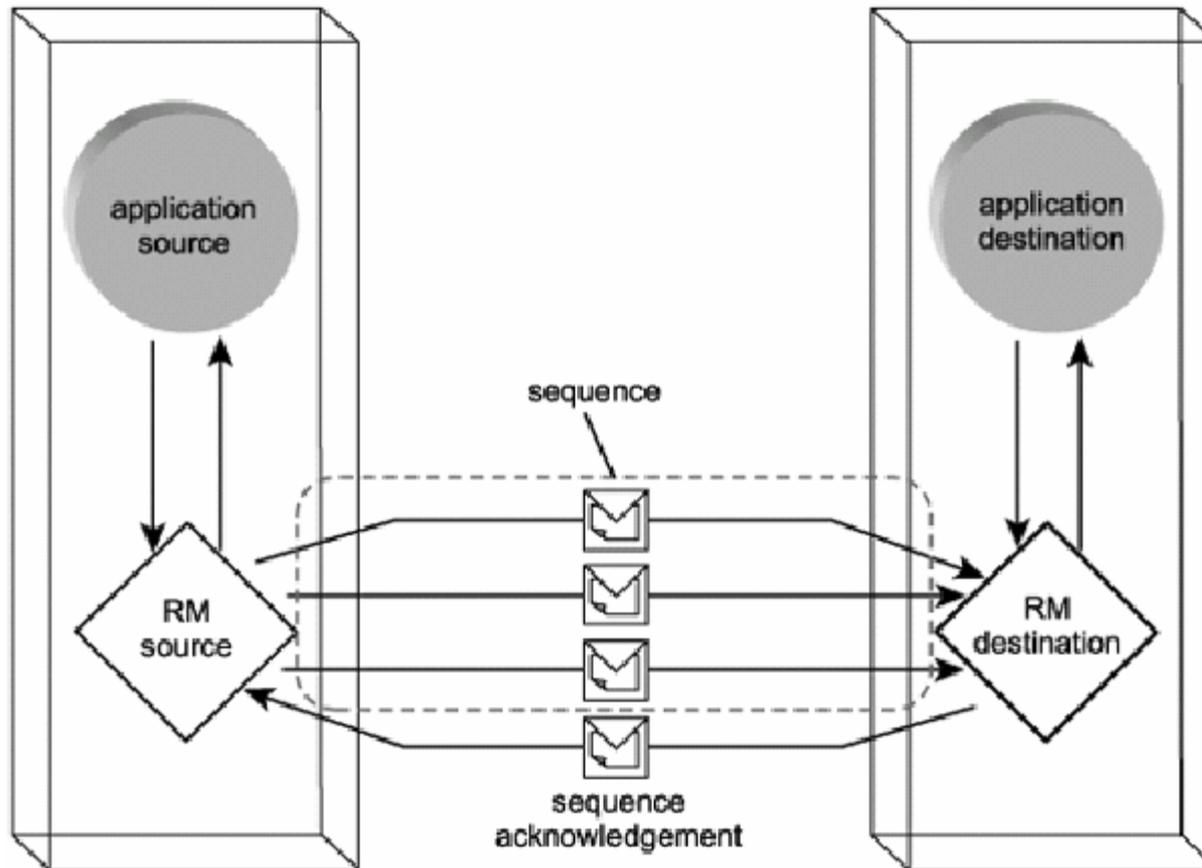
- A sequence establishes the order in which messages should be delivered.
- Each message that is part of a sequence is labeled with a message number that identifies the position of the message within the sequence.
- The final message in a sequence is further tagged with a last message identifier

# RM Source, RM Destination, Application Source, and Application Destination

## **Acknowledgements**

- A core part of the reliable messaging framework is a notification system used to communicate conditions from the RM destination to the RM source.
- Upon receipt of the message containing the last message identifier, the RM destination issues a sequence acknowledgement.
- The acknowledgement message indicates to the RM source which messages were received.
- It is up to the RM source to determine if the messages received are equal to the original messages transmitted.
- The RM source may retransmit any of the missing messages, depending on the delivery assurance used

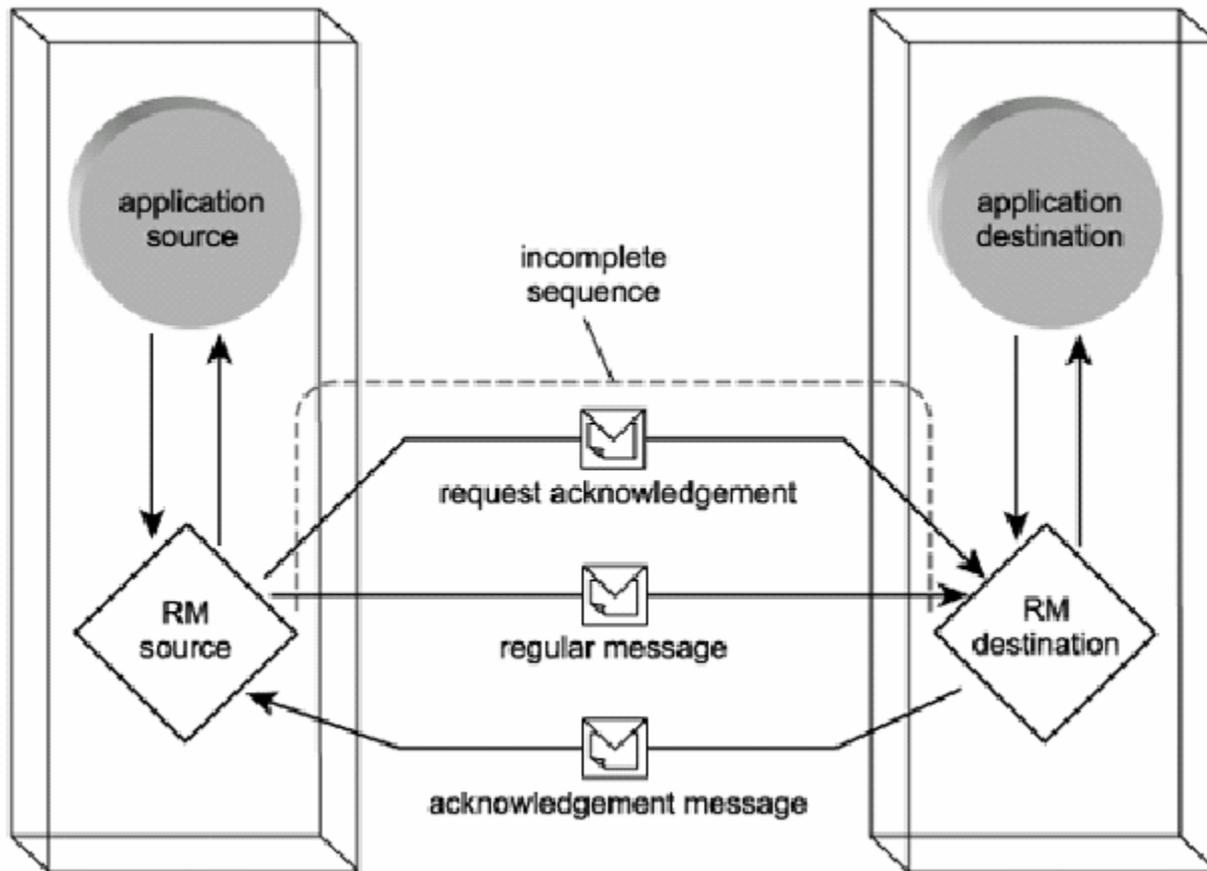
A sequence acknowledgement sent by the RM destination after the successful delivery of a sequence of messages



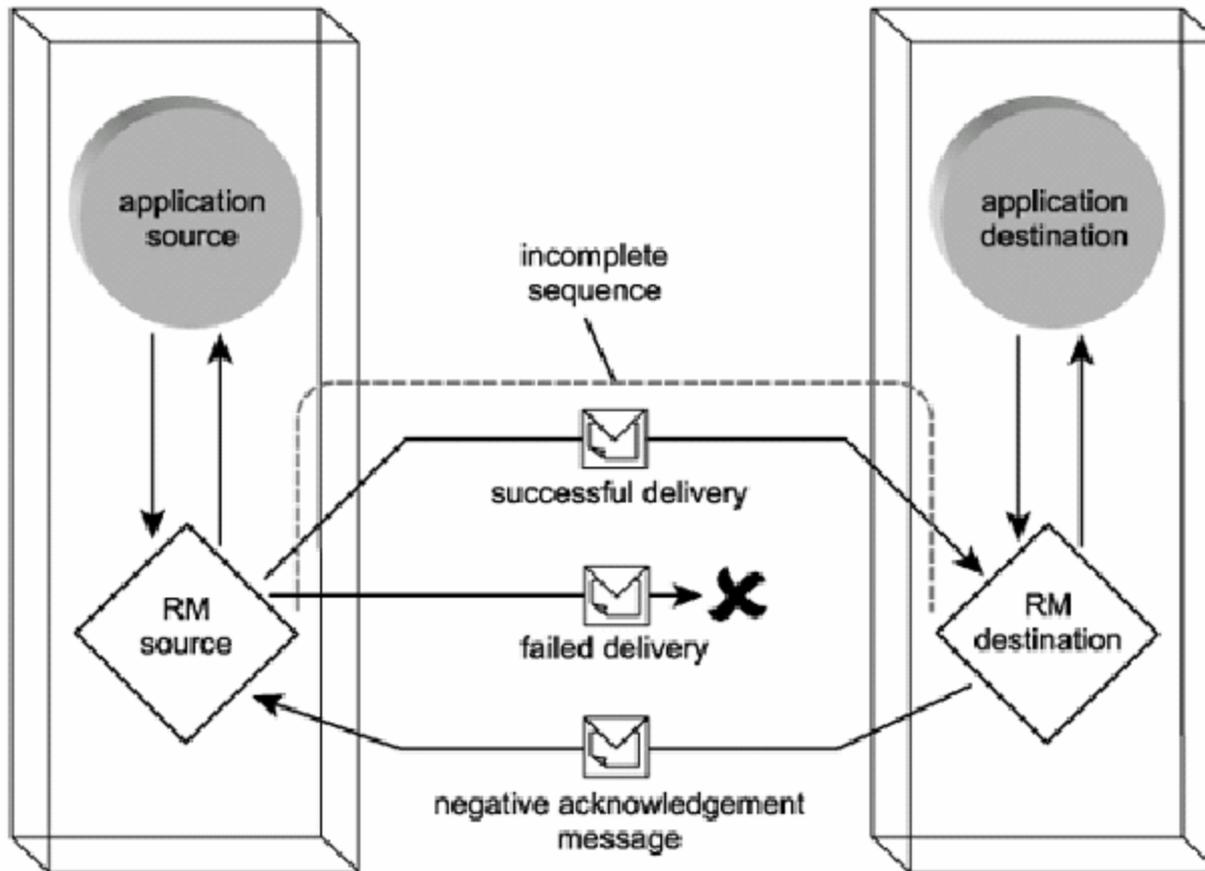
# RM Source, RM Destination, Application Source, and Application Destination

- An RM source does not need to wait until the RM destination receives the last message before receiving an acknowledgement.
- RM sources can request that additional acknowledgements be transmitted at any time by issuing request acknowledgements to RM destinations.
- Additionally, RM destinations have the option of transmitting negative acknowledgements that immediately indicate to the RM source that a failure condition has occurred

A request acknowledgement sent by the RM source to the RM destination, indicating that the RM source would like to receive an acknowledgement message before the sequence completes



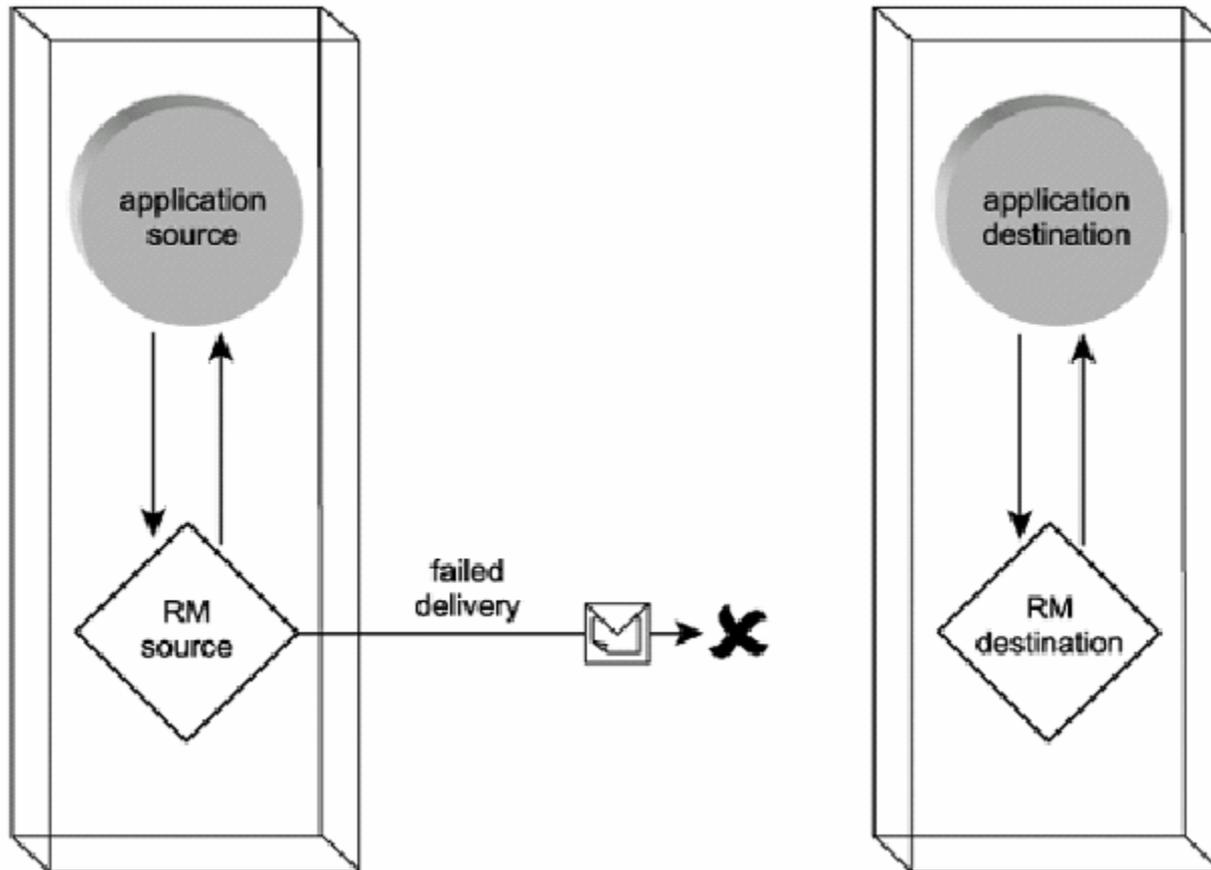
A negative acknowledgement sent by the RM destination to the RM source, indicating a failed delivery prior to the completion of the sequence



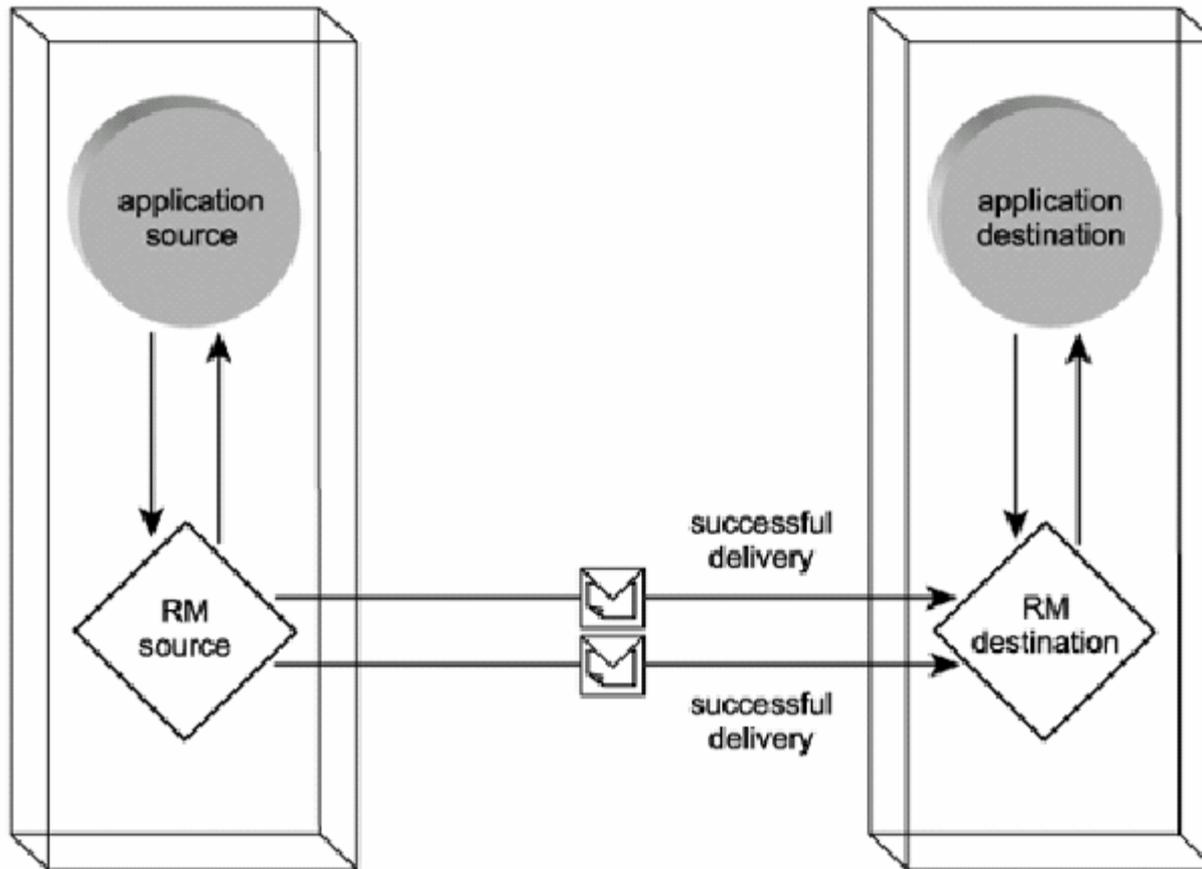
# Delivery assurances

- The nature of a sequence is determined by a set of reliability rules known as delivery assurances.
- Delivery assurances are predefined message delivery patterns that establish a set of reliability policies.
- The following delivery assurances are supported:
- **The AtMostOnce delivery assurance** promises the delivery of one or zero messages. If more than one of the same message is delivered, an error condition occurs.
- **The AtLeastOnce delivery assurance** allows a message to be delivered once or several times. The delivery of zero messages creates an error condition
- **The ExactlyOnce delivery assurance** guarantees that a message only will be delivered once. An error is raised if zero or duplicate messages are delivered
- **The InOrder delivery assurance** is used to ensure that messages are delivered in a specific sequence. The delivery of messages out of sequence triggers an error. Note that this delivery assurance can be combined with any of the previously described assurances.

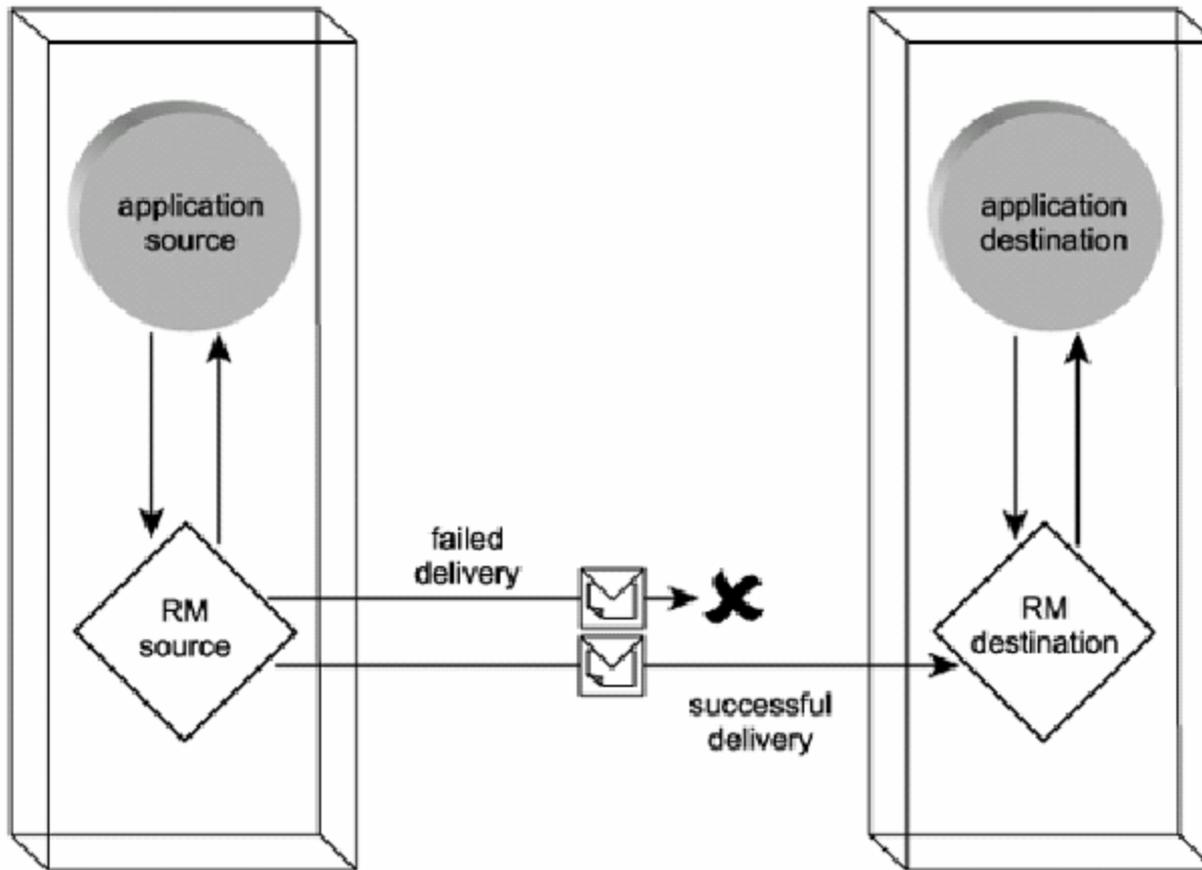
# The AtMostOnce delivery assurance



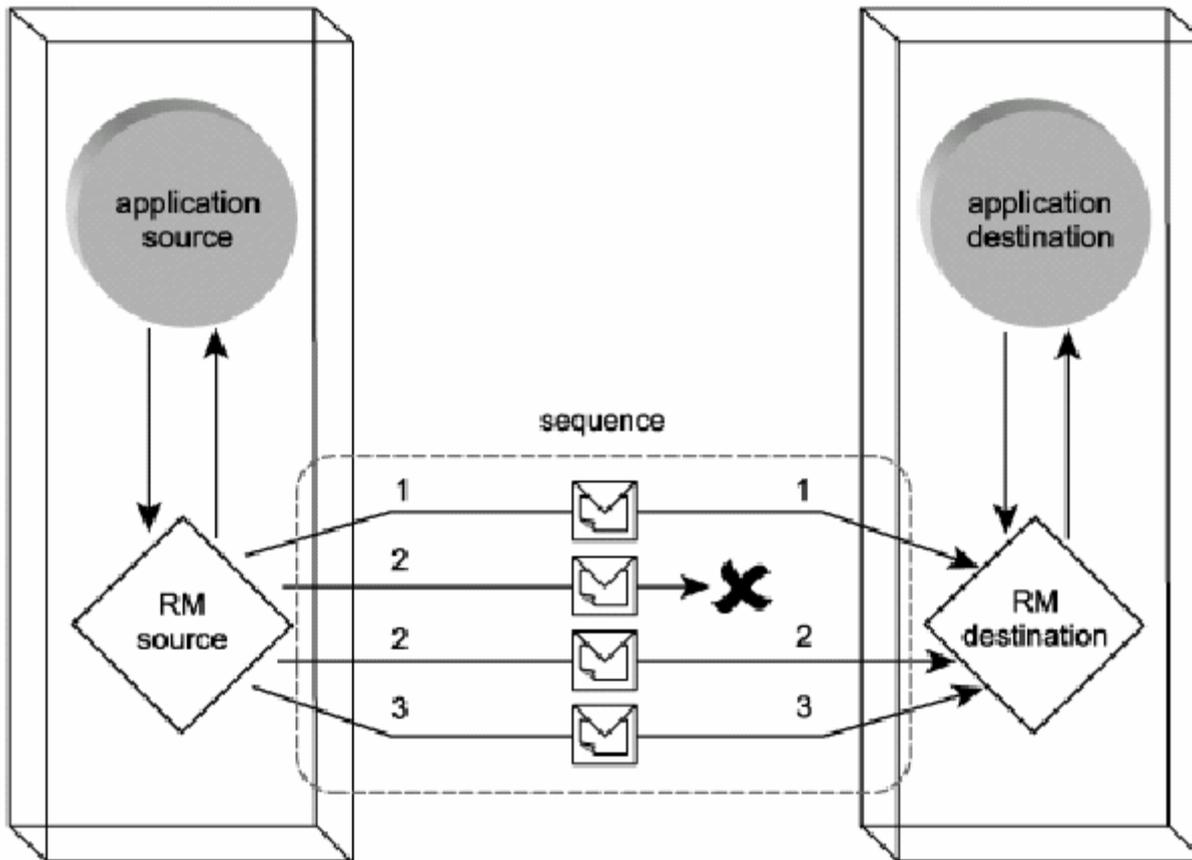
## The AtLeastOnce delivery assurance



# The ExactlyOnce delivery assurance



# The InOrder delivery assurance



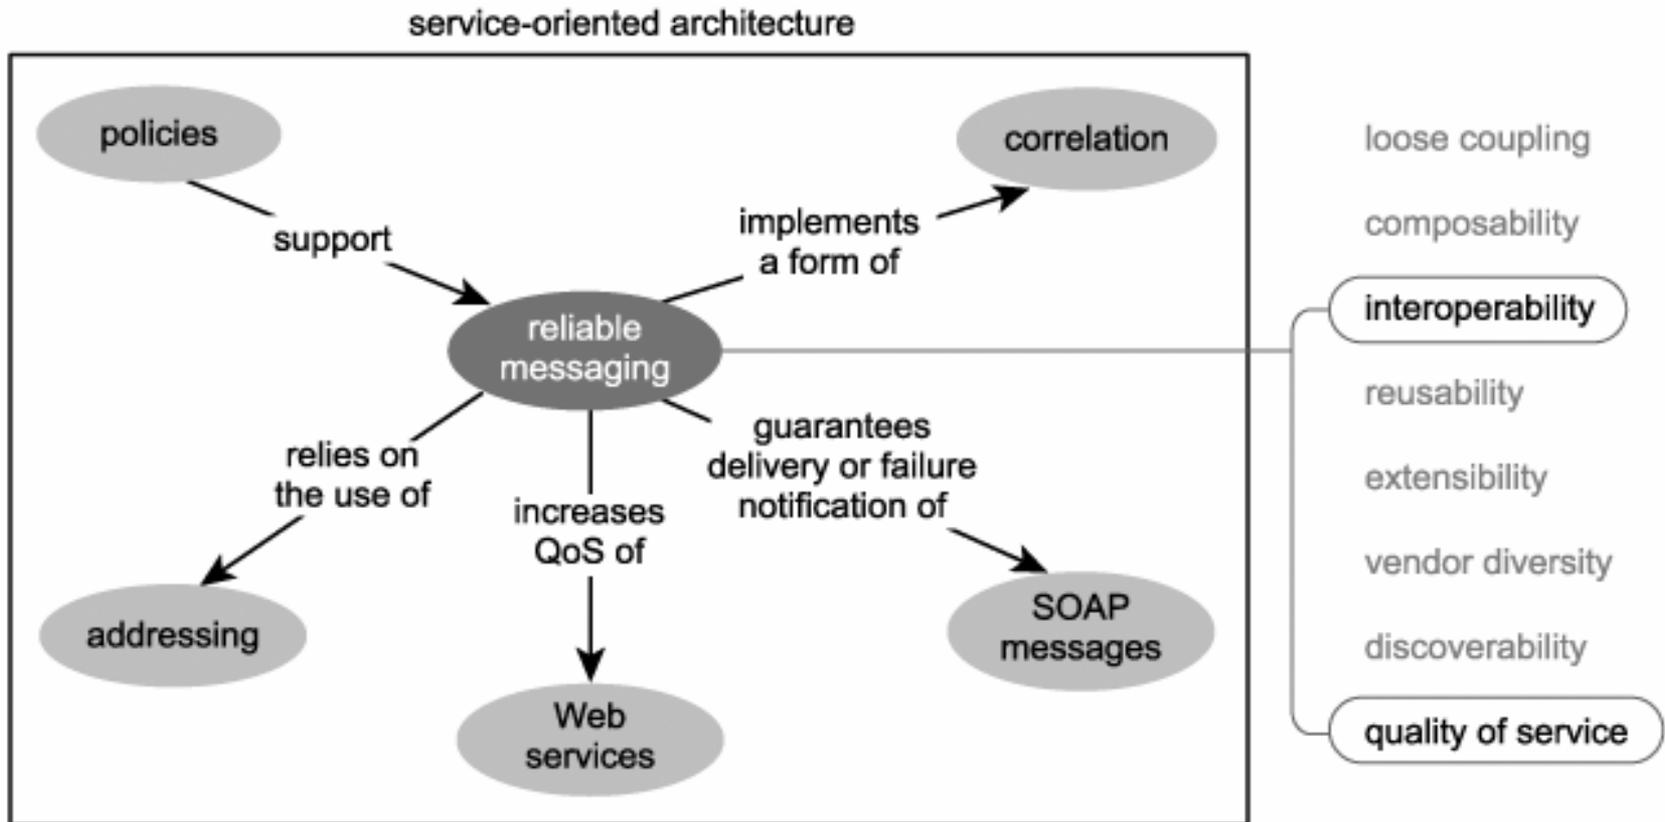
# Reliable messaging and addressing

- WS-Addressing is closely tied to the WS-ReliableMessaging framework.
- In fact, it's interesting to note that the rules around the use of the WS-Addressing message id header were altered specifically to accommodate the WS-ReliableMessaging specification.
- Originally, message id values always had to be unique, regardless of the circumstance. However, the delivery assurances supported by WS-ReliableMessaging required the ability for services to retransmit identical messages in response to communication errors.
- The subsequent release of WS-Addressing, therefore, allowed retransmissions to use the same message ID.

# Reliable messaging and SOA

- Reliable messaging brings to service-oriented solutions a tangible quality of service. It introduces a flexible system that guarantees the delivery of message sequences supported by comprehensive fault reporting.
- This elevates the robustness of SOAP messaging implementations and eliminates the reliability concerns most often associated with any messaging frameworks.
- By increasing the delivery quality of SOAP messages, reliable messaging increases the quality of cross-application communication channels as well.
- The limitations of a messaging framework no longer inhibit the potential of establishing enterprise-level integration.

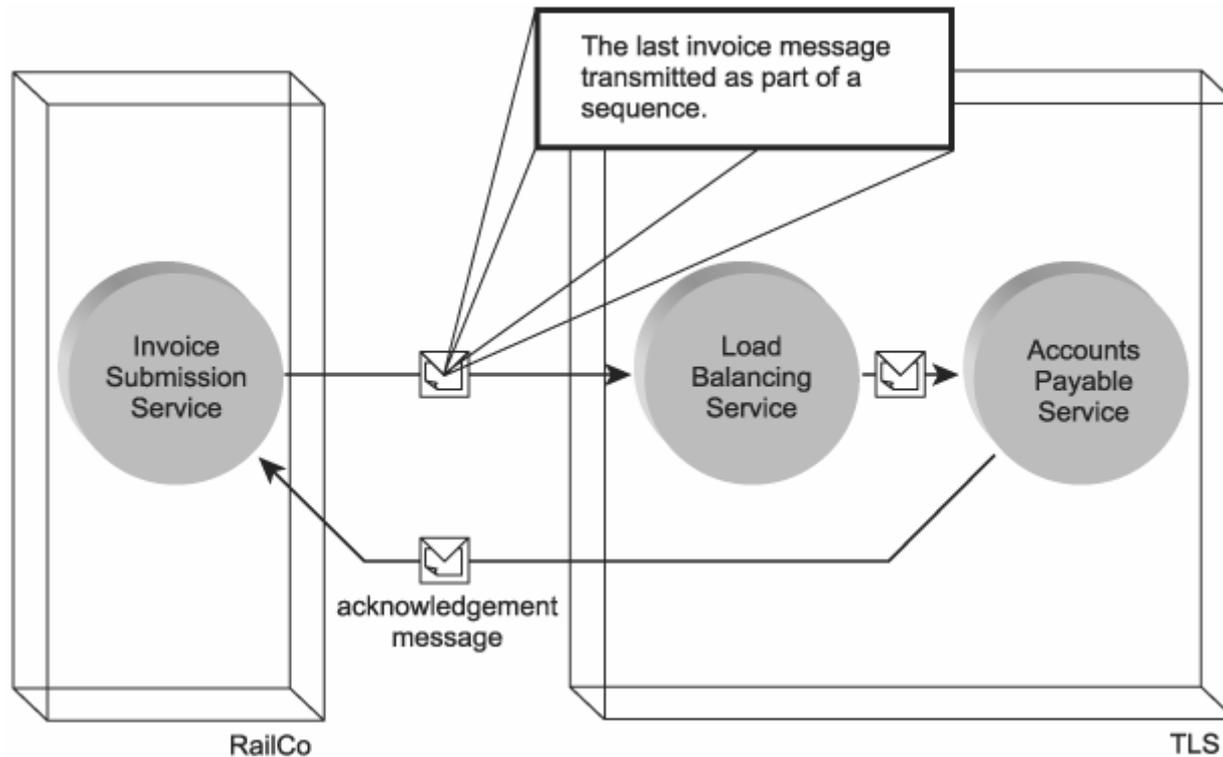
# Reliable messaging relating to other parts of SOA



# Case Study

- To accommodate their existing accounting practices, RailCo sometimes prefers to issue bulk, month-end invoice submissions. TLS has other vendors that require this option and therefore accepts these forms of bulk submissions but under the condition that they must be transmitted as part of the same sequence. This gives TLS the ability to issue an acknowledgement that communicates which of the invoice messages were successfully received.
- RailCo complies with this requirement and enhances its existing Invoice Submission Service to package invoices in SOAP messages that support reliable messaging extensions

After transmitting a series of invoice messages, the last message within the sequence triggers the issuance of an acknowledgement message by TLS



# Case Study

- The first submitted batch consists of 15 invoices. Much to RailCo's dismay, upon transmitting the last message in the sequence, TLS issues an acknowledgement message indicating that only 11 of the 15 invoice messages were actually received.
- In preparation for subsequent bulk submissions, RailCo extends its Invoice Submission Service to issue an acknowledgement request message after every second invoice message sent as part of a sequence. This allows RailCo to better monitor and respond to failed delivery attempts.

# Summary of Key Points

- WS-ReliableMessaging establishes a framework that guarantees the delivery of a SOAP message or the reporting of a failure condition.
- The key parts of this framework are a notification system based on the delivery of acknowledgement messages and a series of delivery assurances that provide policies comprised of reliability rules.
- WS-ReliableMessaging is closely associated with the WS-Addressing and WS-Policy specifications.
- Reliable messaging significantly increases SOA's quality of service level and broadens its interoperability potential.

# Correlation

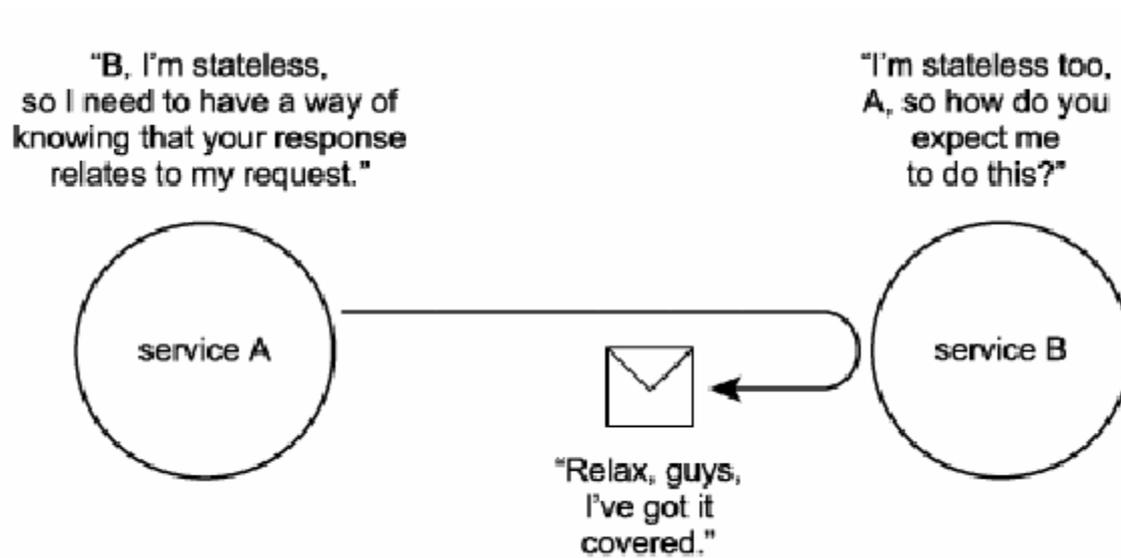
# Correlation

- One of the fundamental requirements for exchanging information via Web services is the ability to persist context and state across the delivery of multiple messages.
- Because a service-oriented communications framework is inherently loosely coupled, there is no intrinsic mechanism for associating messages exchanged under a common context or as part of a common activity.
- Even the execution of a simple request-response message exchange pattern provides no built-in means of automatically associating the response message with the original request.

# Correlation

- Correlation addresses this issue by requiring that related messages contain some common value that services can identify to establish their relationship with each other or with the overall task they are participating in.
- The specifications that realize this simple concept provide different manners of implementation.
- We therefore dedicate the following section to explaining what correlation is and comparing how it is implemented by some of the WS-\* extensions

Correlation places the association of message exchanges into the hands of the message itself



# Correlation in abstract

- To establish a neutral point of reference, let's start with a very basic description of correlation without any reference to its implementation. In tightly bound communication frameworks the issue of correlated units of communication (individual transmissions) rarely arose.
- The technology that enabled tightly bound communication between components, databases, and legacy applications typically established an active connection that persisted for the duration of a given business activity (or longer).
- Because the connection remained active, context was inherently present, and correlation between individual transmissions of data was intrinsically managed by the technology protocol itself.

# Correlation in abstract

- Things change dramatically when you fiddle with the coupling, however. When one stateless service sends a message to another, it loses control of the message and preserves no context of the activity in which the message is participating.
- It is up to the message to introduce the concept of correlation to provide services with the ability to associate a message with others.
- This is achieved by embedding a value in the message that is propagated to all related messages. When a service processes a message and locates this value, it can establish a form of context, in that it can be used to associate this message with others.
- The nature of the context can vary. For example, a message could be part of a simple exchange activity, an atomic transaction, or a long running orchestration.

# Correlation in MEPs and activities

- Because they are generic and non-business-specific in nature, MEPs and activities have no predefined notion of correlation.
- They are simple, conceptual building blocks incorporated and assembled by either custom-developed solutions that employ custom correlation identifiers and related processing logic or by specifications that impose proprietary forms of correlation (as described next).

# Correlation in coordination

- The context management framework provided by WS-Coordination establishes a sophisticated mechanism for propagating identifiers and context information between services.
- A separate activation service is responsible for creating new activities and subsequently generating and distributing corresponding context data. Services can forward this information to others that can use it to register for participation in the activity.
- While context management uses a correlation-type identifier to uniquely represent a specific activity context, it goes well beyond correlation features to provide a comprehensive context management framework that can be leveraged through activity protocols, such as those supplied by the WS-AtomicTransaction and WS-BusinessActivity extensions.

# Correlation in orchestration

- WS-BPEL orchestrations need to concern themselves with the correlation of messages between process and partner services.
- This involves the added complexity of representing specific process instances within the correlation data. Further complicating this scenario is the fact that a single message may participate in multiple contexts, each identified by a separate correlation value.
- To facilitate these requirements, the WS-BPEL specification defines specific syntax that allows for the creation of extensible correlation sets. These message properties can be dynamically added, deleted, and altered to reflect a wide variety of message exchange scenarios and environments.

# Correlation in addressing

- WS-Addressing's message id and relationship MI headers provide inherent correlation abilities, which can be leveraged by many composition and messaging extensions.

# Correlation in reliable messaging

- Every message that participates in a WS-ReliableMessaging sequence carries sequence information with it.
- This data consists of a sequence identifier that represents the series of messages required to follow the messaging rules of the sequence, along with a message identifier that identifies how the current message fits into the overall sequence.
- As a whole, this information can be considered correlation-related. However, its primary purpose is to support the enforcement of reliability rules.

# Correlation and SOA

- Correlation is a key contributor to preserving service autonomy and statelessness.
- Though simple by nature, the ability to tie messages together without requiring that services somehow manage this association is an important function of correlation, primarily because of how common message associations are in enterprise SOAs.

# Case Study

- The PO Submission Process we described earlier consists of a complex activity involving the TLS Purchase Order and the RailCo Order Fulfillment Services. In our previous examples we explained how the path of the PO message can be determined and extended dynamically at runtime and how it spans multiple services.
- For each service that receives the PO message to understand the context under which it should process the message contents, it needs to be able to differentiate the message from others. It accomplishes this by associating a unique value with the message. In this case, the message identifier used is a value partially auto-generated and partially derived from the message PO number.

# Summary of Key Points

- Correlation is a required part of any SOA, as it enables the persistence of activity context across multiple message exchanges, while preserving the loosely coupled nature of service-oriented solutions.
- WS-\* specifications implement correlation in different ways, but many specifications increasingly are relying on WS-Addressing for a form of standardized correlation.
- Even though values from a message's content can be used for correlation purposes, SOAP headers are the most common location for correlation identifiers.
- Correlation is an essential part of messaging within SOA, as it preserves service statelessness and (to an extent) supports message autonomy.

# Policies

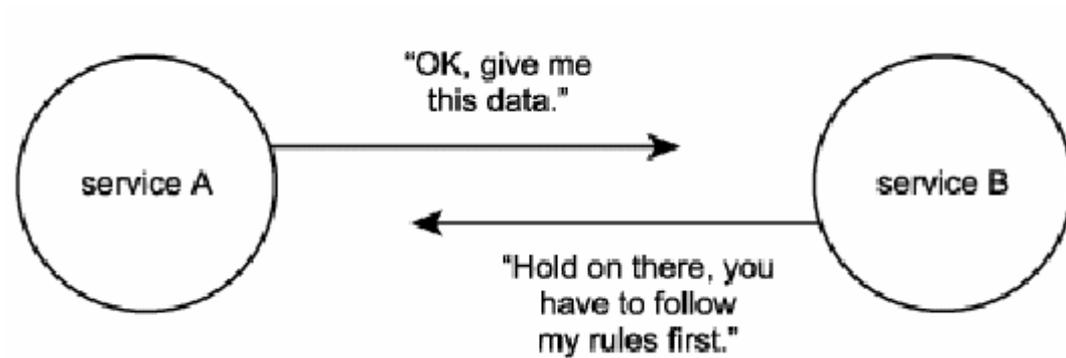
# Policies

- Every automated business task is subject to rules and constraints. These characteristics trickle down to govern the behavior of the underlying services that automate the task.
- The source of these restrictions could be:
  - actual business-level requirements
  - the nature of the data being exchanged
  - organizational security measures

# Policies

- Further, every service and message has unique characteristics that may be of interest to other services that cross its path.
- Examples include:
  - behavioral characteristics
  - preferences
  - technical limitations
  - quality of service (QoS) characteristics
- Services can be outfitted with publicly accessible metadata that describes properties such as the ones listed here. This information is housed in a policy

Policies can express a variety of service properties, including rules



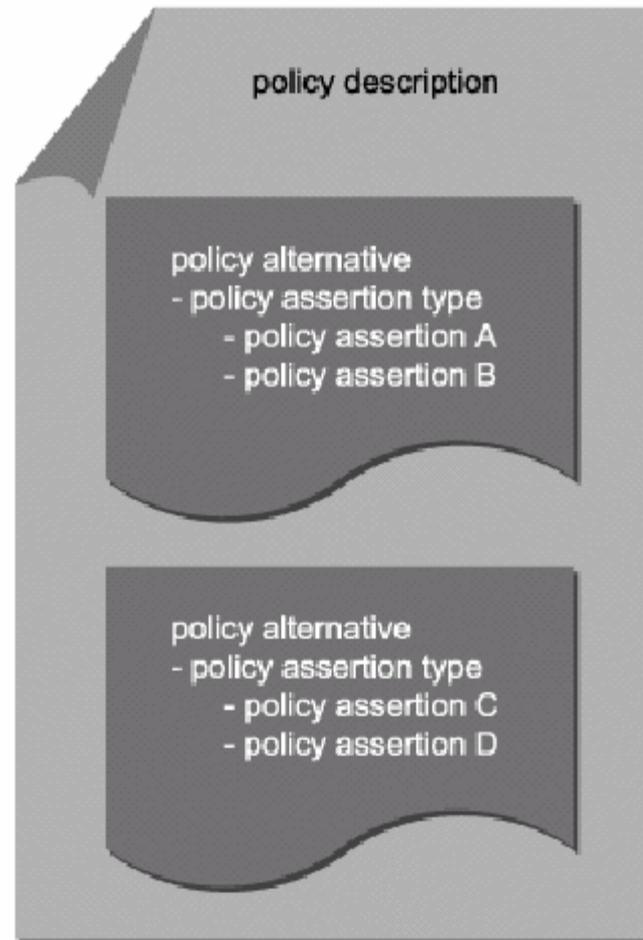
# Policies

- The use of policies allows a service to express various characteristics and preferences and keeps it from having to implement and enforce rules and constraints in a custom manner.
- It adds an important layer of abstraction that allows service properties to be independently managed.

# The WS-Policy framework

- The WS-Policy framework establishes extensions that govern the assembly and structure of policy description documents, as well as the association of policies to Web resources.
- This framework is comprised of the following three specifications:
  - WS-Policy
  - WS-PolicyAttachments
  - WS-PolicyAssertions

# The basic structure of a policy description



# The WS-Policy framework

- Note also that the WS-Policy framework forms part of the WS-Security framework. Specifically, the WS-SecurityPolicy specification defines a set of policy assertions intended for use with WS-Security
- Policies can be programmatically accessed to provide service requestors with an understanding of the requirements and restrictions of service providers at runtime. Alternatively, policies can be studied by humans at design time to develop service requestors designed to interact with specific service providers.
- Recent revisions to the WS-Policy framework have extended the structure of a policy description and its associated terminology.

# Policy assertions and policy alternatives

- The service properties expressed by a policy description are represented individually by policy assertions.
- A policy description therefore is comprised of one or more policy assertions. Examples of policy assertions include service characteristics, preferences, capabilities, requirements, and rules. Each assertion can be marked as optional or required.
- Policy assertions can be grouped into policy alternatives. Each policy alternative represents one acceptable (or allowable) combination of policy assertions.
- This gives a service provider the ability to offer service requestors a choice of policies.

# Policy assertion types and policy vocabularies

- Policy assertions can be further categorized through policy assertion types.
- Policy assertion types associate policy assertions with specific XSD schemas.
- In the same manner as XML vocabularies are defined in XSD schemas, policy vocabularies simply represent the collection of policy types within a given policy.
- Similarly, a policy alternative vocabulary refers to the policy types contained within a specific policy alternative.

# Policy subjects and policy scopes

- A policy can be associated with a Web service, a message, or another resource.
- Whatever a policy is intended for is called a policy subject.
- Because a single policy can have more than one subject, the collection of a policy's subjects is referred to as the policy scope.

## Policy expressions and policy attachments

- Policy assertions that are physically implemented using the WS-Policy language are referred to as policy expressions.
- In other words, a policy expression is simply the XML statement used to express a policy assertion in a manner so that it can be programmatically processed.
- Policy expressions are physically bound to policy scopes using policy attachments.

# Policies in coordination

- When the WS-Coordination context coordination service generates context information for participating services, it can make the distribution of context data subject to the validation of security credentials and other forms of policy information.
- To enforce these requirements, WS-Coordination can incorporate rules established in policies.

# Policies in orchestration and choreography

- Policies can be applied to just about any subjects that are part of orchestrations or choreographies.
- For example, a policy can establish various requirements for orchestration partner services and choreography participants to interact.

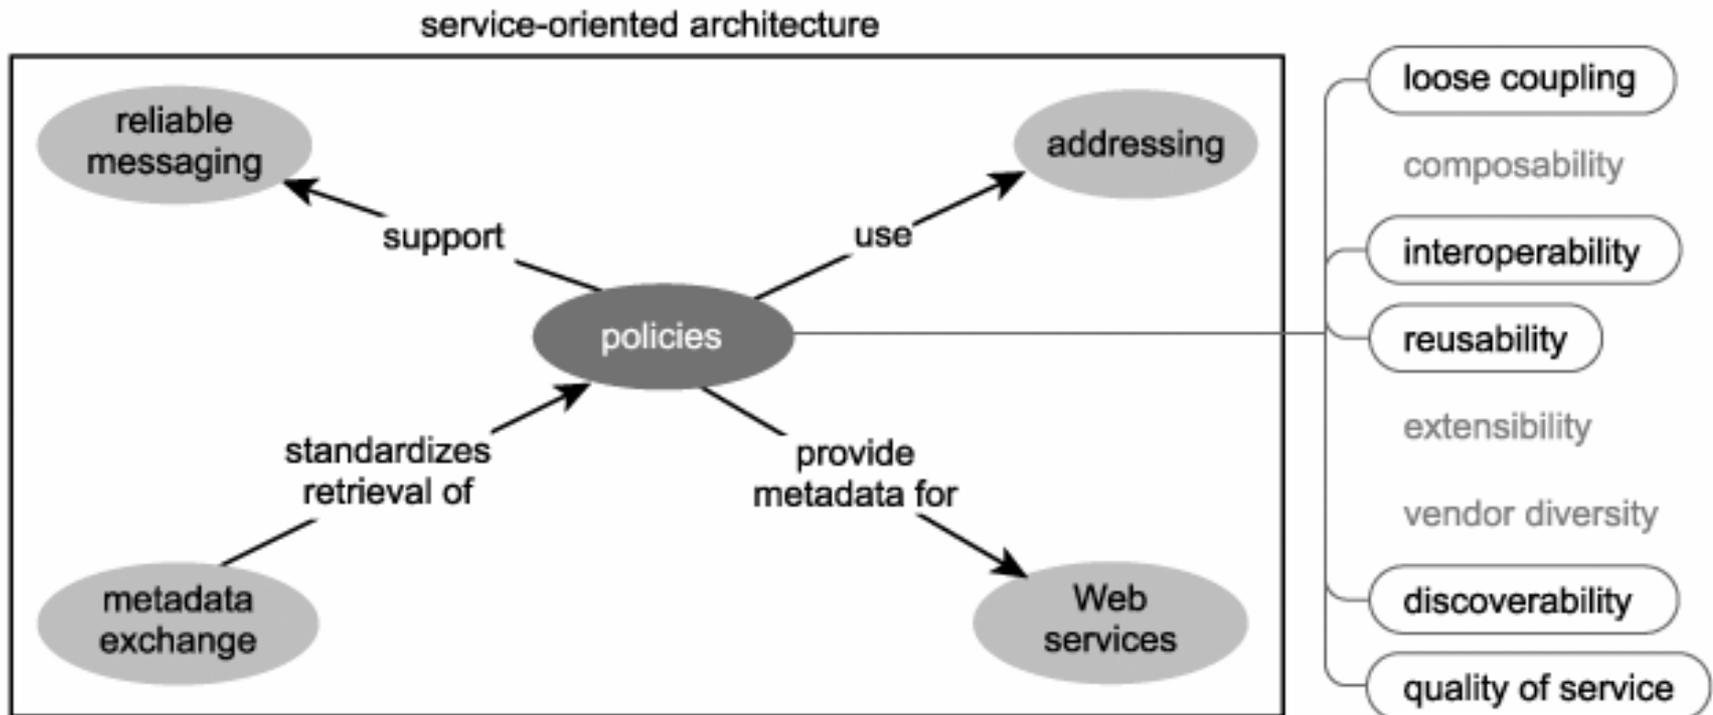
# Policies in reliable messaging

- The WS-ReliableMessaging specification depends on the use of the WS-Policy framework to enable some of its most fundamental features.
- Policies are used to implement delivery assurances through the attachment of policy assurances to the messages that take part in reliable messaging exchanges.
- A further set of policy assertions is provided to add various supplemental rules, constraints and reliability requirements.

# Policies and SOA

- If an SOA is a city, then policies are certainly the laws, regulations, and guidelines that exist to maintain order among inhabitants.
- Policies are a necessary requirement to building enterprise-level service-oriented environments, as they provide a means of communicating constraints, rules, and guidelines for just about any facet of service interaction.
- As a result, they improve the overall quality of the loosely coupled arrangement services are required to maintain.

# Policies relating to other parts of SOA



# Policies and SOA

- Policies allow services to express so much more about themselves beyond the fundamental data format and message exchange requirements established by WSDL definitions. And policies enable services to broaden the range of available metadata while still allowing them to retain their respective independence.
- The use of policies increases SOA's quality of service level by restricting valid message transmissions to those that conform to policy rules and requirements. A side benefit of inserting endpoint level constraints is that the application logic underlying services is not required to perform as much custom exception handling to deal with invalid message submissions.

# Policies and SOA

- Policies naturally improve the ability for services to achieve better levels of interoperation because so much more information about service endpoints can be expressed and published.
- Finally, because they increase the richness of service contracts, they open the door to dynamic discovery and binding.

# Case Study

- TLS recently upgraded some of its middleware, which now provides support for the most recent version of the WS-ReliableMessaging specification. TLS wants to utilize this support but realizes that its partners still may need to continue using the previous version of WS-ReliableMessaging for some time. As a result, it chooses to support both versions by issuing a policy document containing a policy alternative.
- This policy alternative states that its Vendor Profile Service will accept invoice submission sequence headers that conform to both versions of WS-ReliableMessaging, but it also expresses the fact that the newer version is preferred by TLS.
- Later, TLS expands these policy assertions to include a requirement for a specific message encoding type. Regardless of which alternative is chosen by a service requestor, the same text encoding format is required.

# Summary of Key Points

- The WS-Policy framework provides a means of attaching properties (such as rules, behaviors, requirements, and preferences) to Web resources, most notably Web services.
- Individual properties are represented by policy assertions, which can be marked as optional or required. This allows a service to communicate non-negotiable and preferred policies.
- WS-Policy can be incorporated within the majority of WS-\* extensions.
- Policies add an important layer of metadata to SOAs that increases the interoperability and discovery potential for services, while also elevating the overall quality of messaging within SOA.

# Metadata exchange

# Metadata exchange

- The sole requirement for a service requestor to interact with a service provider acting as the ultimate receiver is that the service requestor be in possession of the service provider's service description.
- The WSDL definition, along with any associated XSD schemas, provides the basic set of metadata required to send valid SOAP messages for consumption by the service provider.

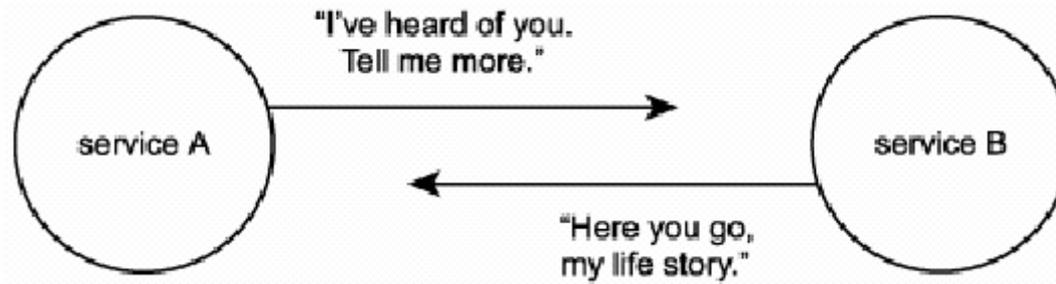
# Metadata exchange

- It is clear that policies add another important layer to the metadata stack. Using policies, our service requestor now can send SOAP messages that comply with both the WSDL interface requirements and the associated policy assertions.
- Again, though, regardless of how much metadata a service makes available, the fact is that we still need to retrieve this information by either:
  - manually locating it by searching for published documents
  - manually requesting it by contacting the service provider entity (the service owner)
  - programmatically retrieving it via a public service registry
  - programmatically retrieving it by interacting with proprietary interfaces made available by the service provider entity

# Metadata exchange

- With the exception of using the public service registry, none of these options are particularly attractive or efficient.
- It would be ideal if we could simply send a standardized request such as, "give me all of the information I need to evaluate and interact with your service provider."

Metadata exchanges let service requestors ask what they want to know about service providers



# The WS-MetadataExchange specification

- This specification essentially allows for a service requestor to issue a standardized request message that asks for some or all of the meta information relating to a specific endpoint address.
- In other words, if a service requestor knows where a service provider is located, it can use metadata exchange to request all service description documents that comprise the service contract by sending a metadata exchange request to the service provider.
- Originally the WS-MetadataExchange specification specified the following three types of request messages:
  - Get WSDL
  - Get Schema
  - Get Policy

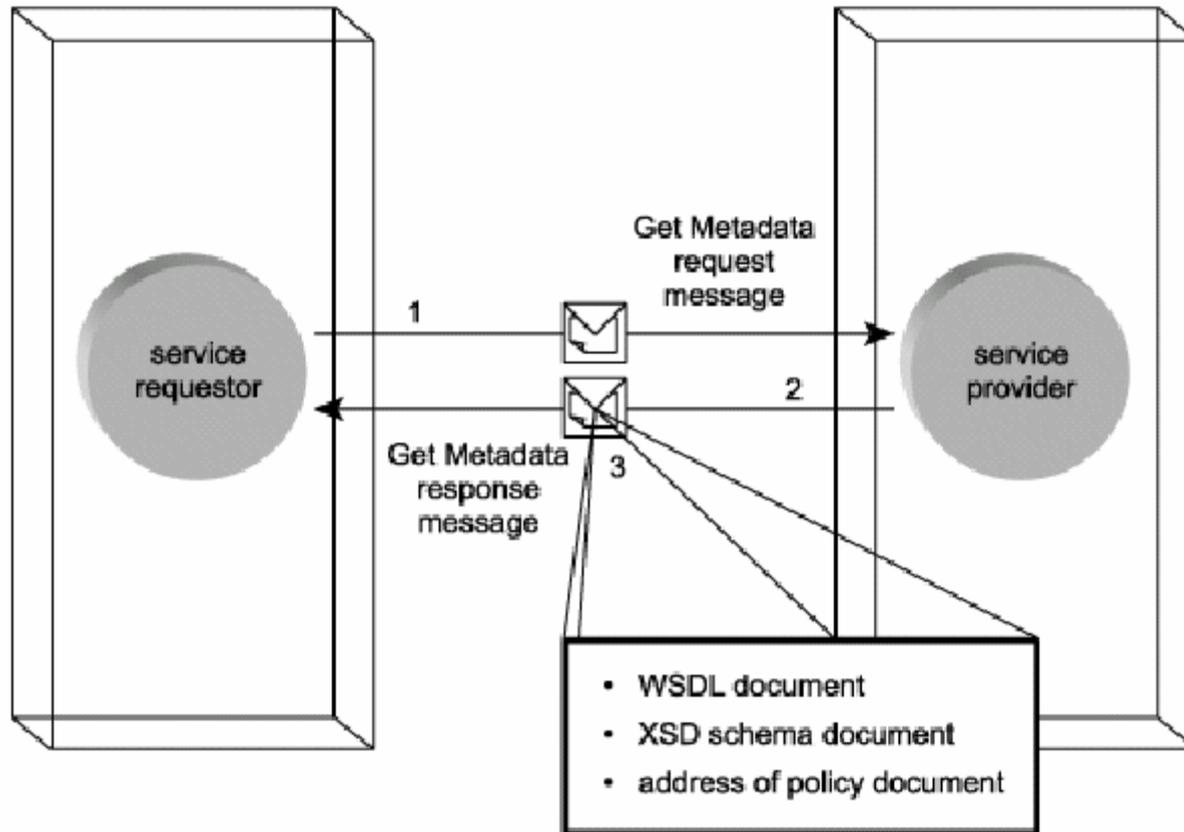
# The WS-MetadataExchange specification

- Even though these represent the three most common types of meta information currently attached to Web services, the specification authors realized that future metadata documents would likely emerge.
- A subsequent revision therefore resulted in a single type of request message:
- Get Metadata
- This message is further supplemented by the Get request message.

# Get Metadata request and response messages

- As previously mentioned, a service requestor can use metadata exchange to programmatically request available metadata documents associated with a Web service.
- To do so, it must issue a Get Metadata request message. This kicks off a standardized request and response MEP resulting in the delivery of a Get Metadata response message.
- Here's what happens for a metadata retrieval activity to successfully complete:
  1. A service requestor issues the Get Metadata request message. This message can request a specific type of service description document (WSDL, XSD schema, policy), or it can simply request that all available metadata be delivered.
  2. The Get Metadata request message is received at the endpoint to which it is delivered. The requested meta information is documented in a Get Metadata response message.
  3. The Get Metadata response message is delivered to the service requestor. The contents of this message can consist of the actual metadata documents, address references to the documents, or a combination of both.

# Contents of a sample Get Metadata response message



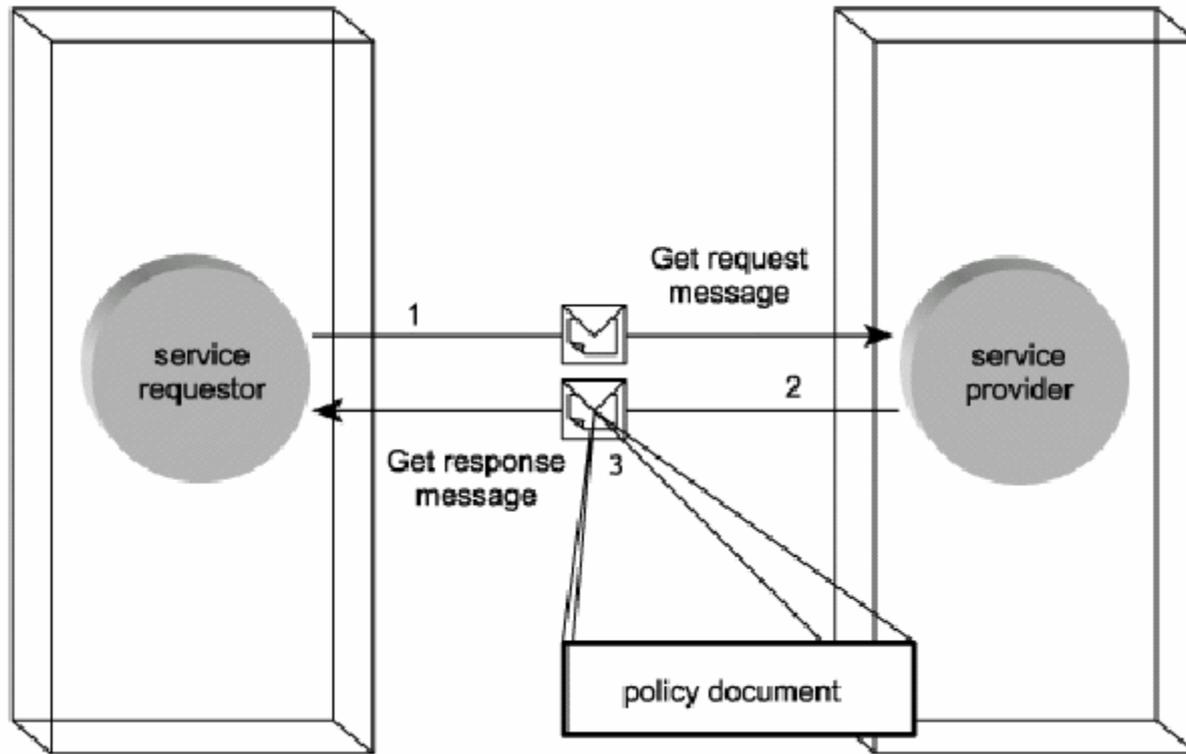
# Get request and response messages

- In Step 3 of the preceding scenario, we explained how the Get Metadata response message does not need to actually contain all of the requested metadata. It can simply provide a list of URIs that point to the separate documents.
- To allow the retrieval of all meta information to be fully automated, the WS-MetadataExchange specification provides a means for the service requestor to explicitly request the document content for any references that were provided as part of the original Get Metadata response message.
- It achieves this through the use of the Get request and Get response messages.

# Get request and response messages

- Here's a brief description of the steps involved in this sub-process:
  1. Upon receiving the Get Metadata response message, the service requestor determines that it would like to receive the actual content of the metadata documents for which it only received references. As a result, the service requestor issues a Get request message indicating which metadata information it would like retrieved.
  2. The Get request message is received at the endpoint to which it was delivered. The requested data is placed into a Get response message.
  3. The Get response message is delivered to the service requestor.
- The next figure shows the execution sequence of these steps, which should provide the service requestor with all the information it needs (and therefore concludes the metadata exchange process).

# Contents of a sample Get response message



# Selective retrieval of metadata

- Meta documents describing services with comprehensive interfaces and features can be large in size, especially when assembled into one mega-description.
- Use of the selective Get request message type therefore reduces the chances of unnecessary information being transported.
- The Get Metadata response message first sends along what is considered the essential piece of service meta information.
- It is then up to the service requestor to determine what further metadata it requires. (Note that the endpoint to which a Get Metadata request message is sent can represent multiple WSDL, XSD schema, and policy documents.)

# Metadata exchange and service description discovery

- It also is important to note that metadata exchange does not really help service requestors discover service providers.
- Service registries, such as those implemented using the UDDI standard, can be used to discover service descriptions that meet certain search criteria.
- While service registries also provide location information for the actual WSDL definition of a service, they can be used in conjunction with metadata exchange messages.

# Metadata exchange and service description discovery

- Essentially, a service requestor could first query a public registry to retrieve the endpoint addresses of any Web service candidates that appear to provide the sought-after features.
- The same requestor could then employ metadata exchange to contact each candidate and request associated metadata documents. This would give the service requestor more information to better assess which service provider it should be working with.
- It also would provide the service requestor with all of the details it needs to begin interacting with the chosen service.
- So while it may not further the cause of attaining discoverable services, it does support discovery by rounding out the overall dynamic discovery process.

# Metadata exchange and version control

- So far we've focused on the ability of metadata exchange to enable service requestors to retrieve any necessary meta information for them to begin interacting with service providers.
- Another important aspect of this WS-\* extension is its potential to automate the administration of service contracts.
- As services evolve, the nature and scope of the functionality they deliver can undergo alterations. This can result in changes trickling down to the service meta layer, which, in turn, can lead to new versions of a service's WSDL, XSD schema, or policy documents.

# Metadata exchange and version control

- This raises the age-old version control challenges. Service requestors already interacting with a service provider either need to be notified ahead of time of upcoming changes, or they need to be supported with an outdated service description.
- Some services-based solutions have dealt with this problem by building custom operations that can be used to retrieve the latest service description (metadata) information.
- While the same functionality is essentially provided by metadata exchange, the main benefit of its use is that it implements this feature in a standardized manner. Now any service-oriented application that supports metadata exchange can allow service requestors to retrieve the latest service contract as often as they like.

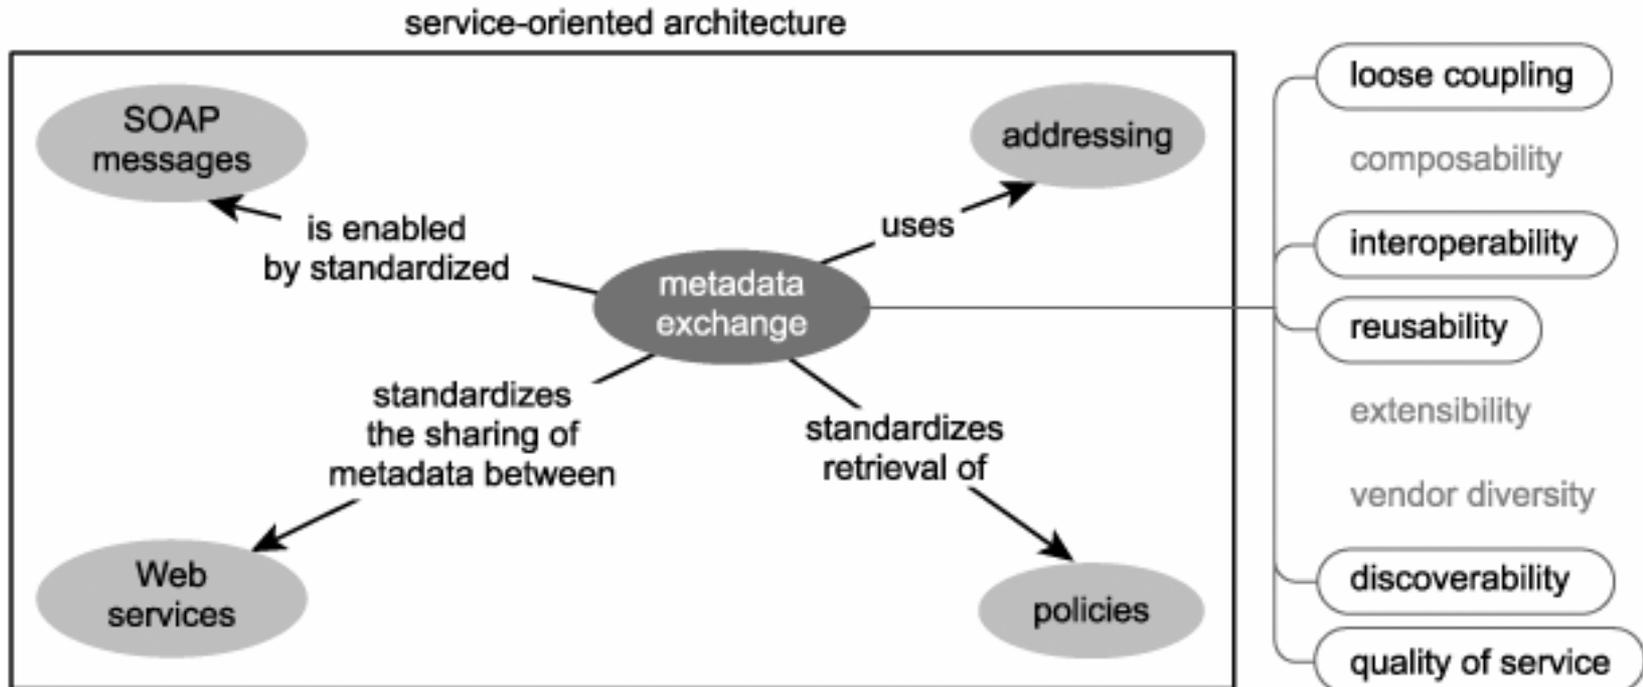
# Metadata exchange and version control

- When changes to meta information are expected to occur frequently, a service requestor could be programmed to periodically retrieve available metadata documents to compare them to the documents already in use.
- In fact, service requestors could even build metadata exchange features into their exception handling.
- If a standard SOAP request is rejected by the service provider as a result of an interface, schema, or policy incompatibility error, the service requestor's exception handling routine could respond by retrieving and checking the latest metadata documents.

# Metadata exchange and SOA

- The simple concepts behind metadata exchange support some key aspects of SOA.
- Its ability to automate the retrieval of meta information reinforces loose coupling between services, and increases the ability for service requestors to learn more about available service providers.
- By standardizing access to and retrieval of metadata, service requestors can programmatically query a multitude of candidate providers.
- Because enough service provider metadata can more easily be retrieved and evaluated, the overall discovery process is improved, and the likelihood for services to be reused is increased.

# Metadata exchange relating to other parts of SOA



# Metadata exchange and SOA

- By establishing a standardized means of service description exchange, this extension can vastly improve interoperability when broadly applied to volatile environments.
- By being able to query service providers prior to attempting access, requestors can verify that the correct metadata is in fact being used for their planned message exchanges.
- This can increase the QoS factor of SOA, as it tends to avoid a multitude of maintenance problems associated with service contract changes.

# Metadata exchange and SOA

- It is also worth mentioning that metadata exchange reduces the need for developers to attain meta information at design time and eliminates the need for custom-developed metadata retrieval extensions.
- Finally, the dynamic exchange of service descriptions can lead to the potential of automating version control and other metadata-related functions.

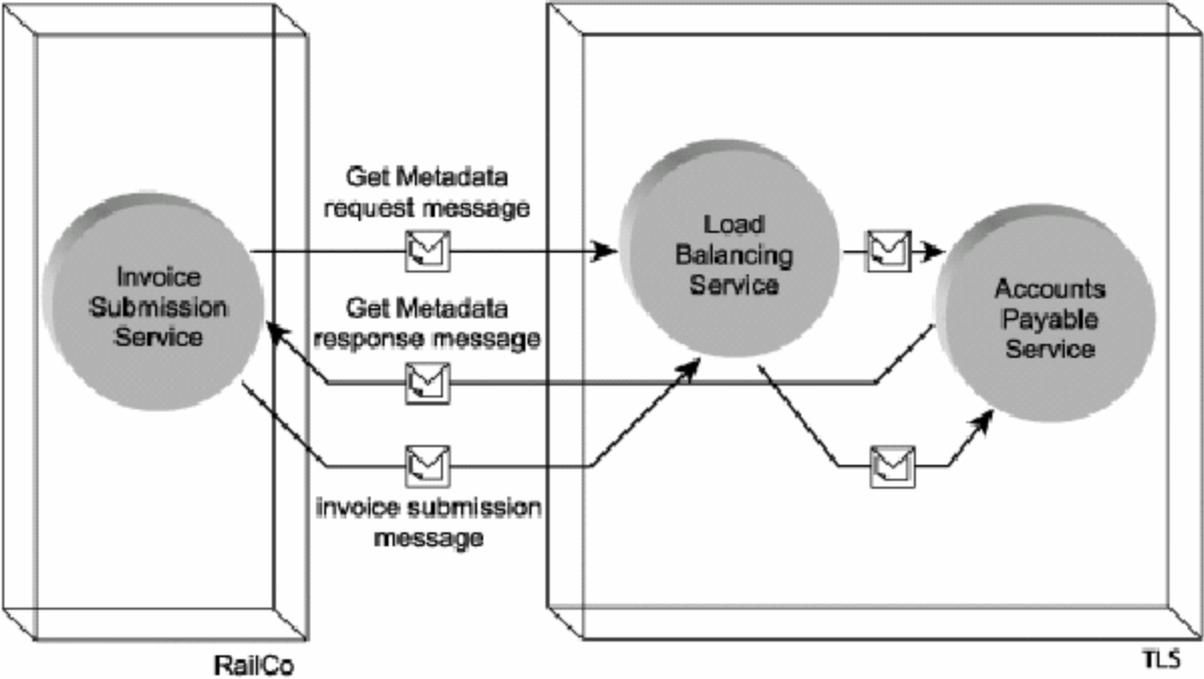
# Case Study

- As TLS continues to evolve its B2B solution, new features are added and some existing functionality is modified. This can, occasionally, result in changes to the WSDL interface definitions of TLS services, as well as revisions to service policies. Any of these changes can obviously affect the online partners that regularly connect to TLS.
- Therefore, all public TLS services support the processing of WS-MetadataExchange requests. At the onset, partners who register for the TLS B2B solution are strongly encouraged to issue Get Metadata request messages frequently to receive the latest service contracts.
- RailCo learned about this the hard way. To date they never bothered incorporating metadata exchange functionality within their services, as they were not required to do so. After a change to the TLS Accounts Payable Service WSDL, though, the RailCo Invoice Submission Service submitted an invoice message that was rejected by TLS.

# Case Study

- The resulting error description was unclear, and exception handling logic within the RailCo service assumed this condition was the result of the TLS service being unavailable. It was therefore designed to periodically retry the message submission on a daily basis. Only after three days did someone at RailCo notice that an acknowledgement had not been received from TLS.
- A lengthy investigation led to the eventual discovery that the failed submissions were the result of a change to the TLS WSDL definition.
- As a result of this experience, RailCo revised their Invoice Submission Service to interact with the metadata exchange functionality offered by TLS. The service now issues a periodic Get Metadata message to the TLS Accounts Payable Service.

The revised RailCo Invoice Submission Process now includes a periodic metadata exchange with TLS



**Security**

# Security

- Security requirements for automation solutions are nothing new to the world of IT.
- Similarly, service-oriented applications need to be outfitted to handle many of the traditional security demands of protecting information and ensuring that access to logic is only granted to those permitted.
- However, the SOAP messaging communications framework, upon which contemporary SOA is built, emphasizes particular aspects of security that need to be accommodated by a security framework designed specifically for Web services.

# Security

- A list of security specifications that may be used as part of SOA. For more information regarding these specifications, visit: [www.specifications.ws](http://www.specifications.ws).
  - WS-Security
  - WS-SecurityPolicy
  - WS-Trust
  - WS-SecureConversation
  - WS-Federation
  - Extensible Access Control Markup Language (XACML)
  - Extensible Rights Markup Language (XrML)
  - XML Key Management (XKMS)
  - XML-Signature
  - XML-Encryption
  - Security Assertion Markup Language (SAML)
  - .NET Passport
  - Secure Sockets Layer (SSL)
  - WS-I Basic Security Profile

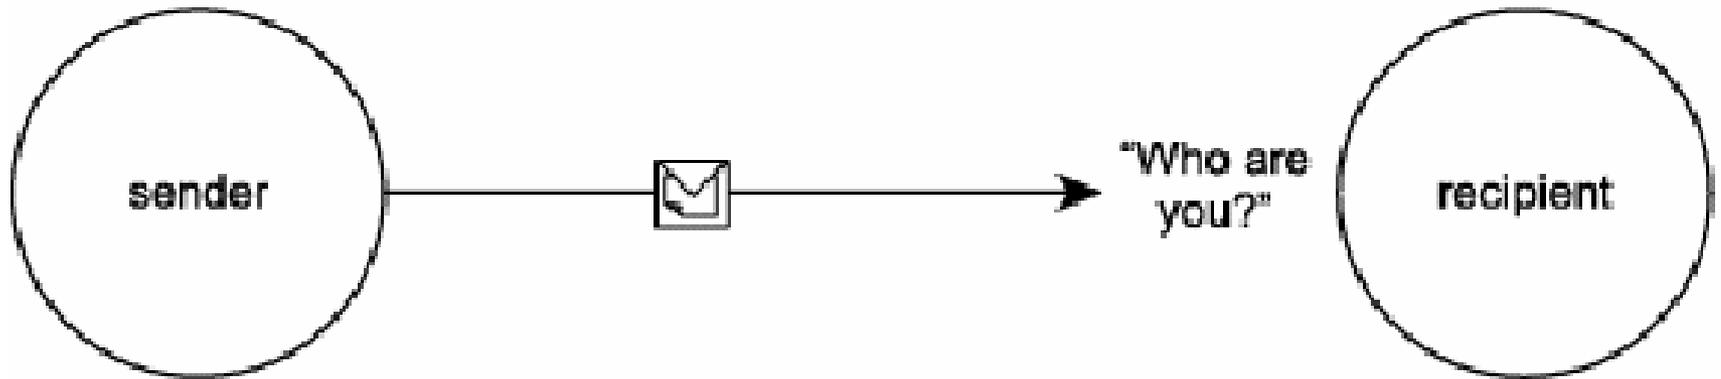
# Security

- A family of security extensions parented by the WS-Security specification comprise such a framework, further broadened by a series of supplementary specifications with specialized feature sets.
- A list of current security-related specifications is given in the previous slide.
- While we clearly cannot discuss concepts for all of them, it is worth spending some time looking at the basic functions performed by the following three core specifications:
  - WS-Security
  - XML-Signature
  - XML-Encryption

# Identification, authentication, and authorization

- For a service requestor to access a secured service provider, it must first provide information that expresses its origin or owner.
- This is referred to as making a claim.
- Claims are represented by identification information stored in the SOAP header.
- WS-Security establishes a standardized header block that stores this information, at which point it is referred to as a token.

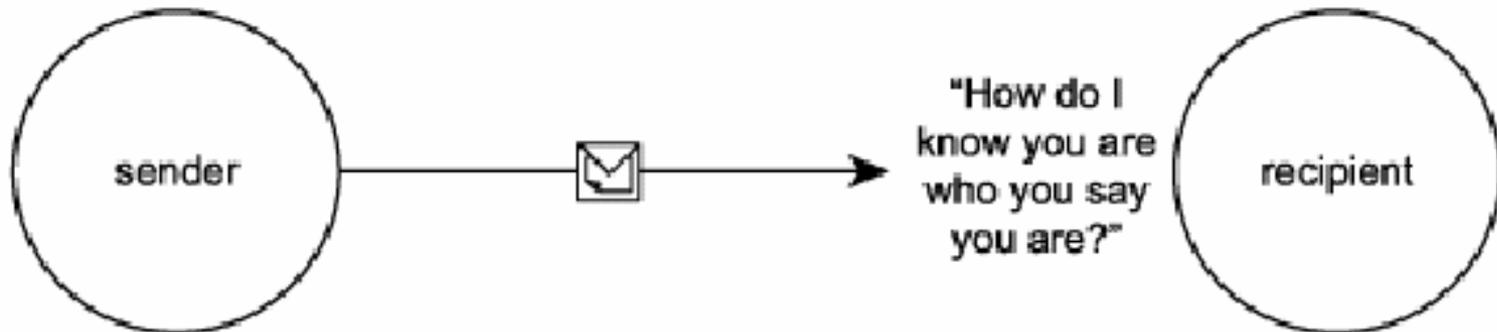
An identity is a claim made regarding the origin of a message



# Identification, authentication, and authorization

- Authentication requires that a message being delivered to a recipient prove that the message is in fact from the sender that it claims to be.
- In other words, the service must provide proof that its claimed identity is true.

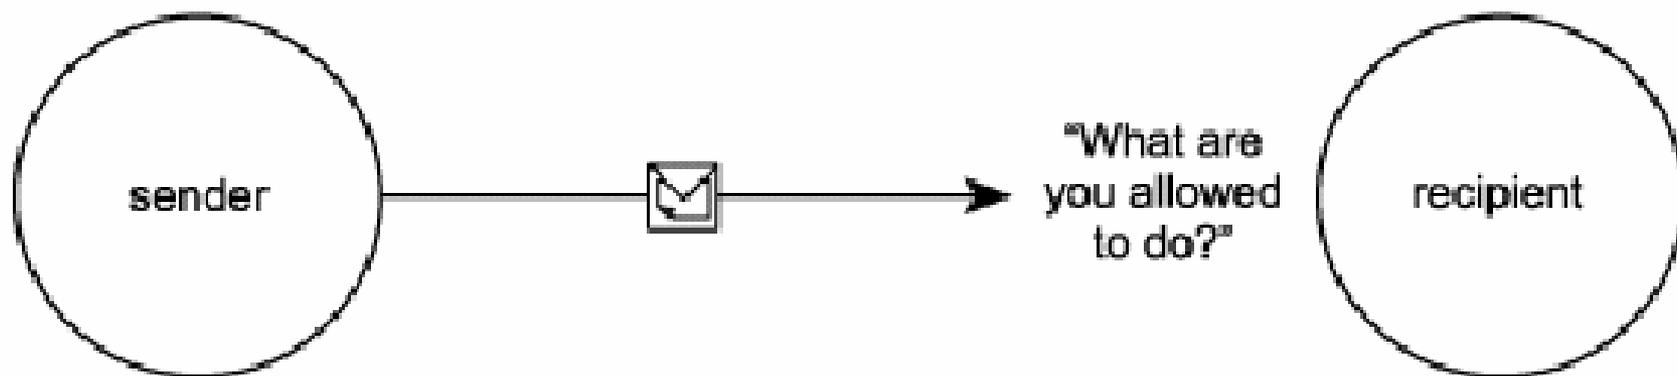
# Authentication means proving an identity



# Identification, authentication, and authorization

- Once authenticated, the recipient of a message may need to determine what the requestor is allowed to do.
- This is called authorization.

Authorization means determining to what extent authentication applies



# Single sign-on

- A challenge facing the enablement of authentication and authorization within SOA is propagating the authentication and authorization information for a service requestor across multiple services behind the initial service provider.
- Because services are autonomous and independent from each other, a mechanism is required to persist the security context established after a requestor has been authenticated.
- Otherwise, the requestor would need to re-authenticate itself with every subsequent request.

# Single sign-on

- The concept of single sign-on addresses this issue.
- The use of a single sign-on technology allows a service requestor to be authenticated once and then have its security context information shared with other services that the requestor may then access without further authentication.
- There are three primary extensions that support the implementation of the single sign-on concept:
  - SAML (Security Assertion Markup Language)
  - .NET Passport
  - XACML (XML Access Control Markup Language)

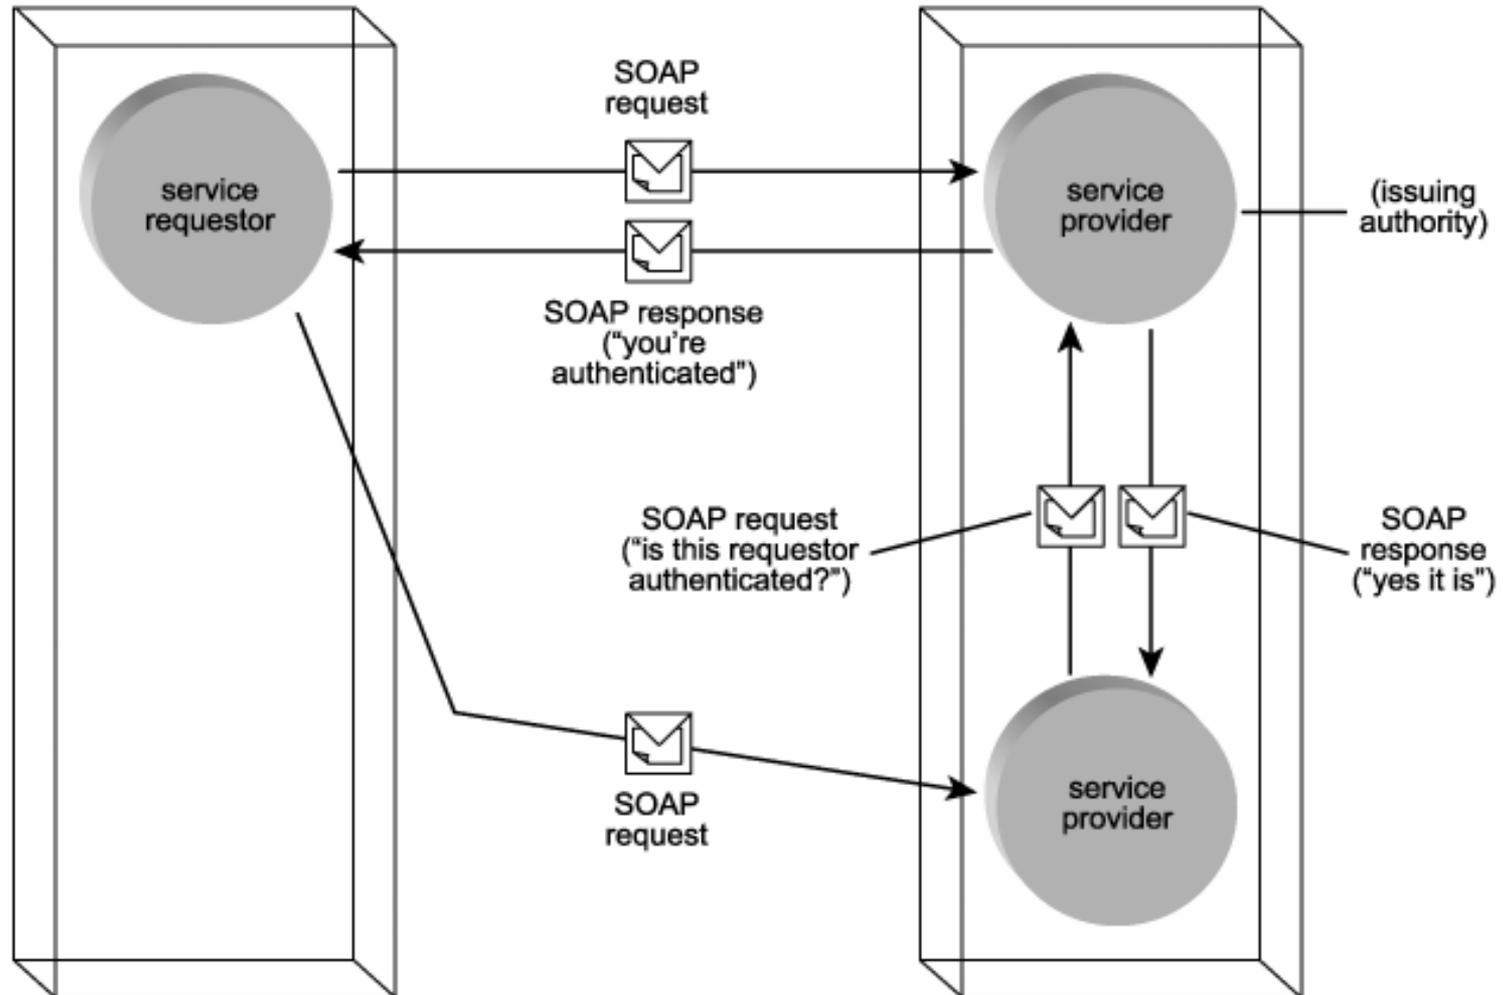
# Single sign-on

- As an example of a single sign-on technology that supports centralized authentication and authorization, let's briefly discuss some fundamental concepts provided by SAML.
- SAML implements a single sign-on system in which the point of contact for a service requestor can also act as an issuing authority.
- This permits the underlying logic of that service not only to authenticate and authorize the service requestor, but also to assure the other services that the service requestor has attained this level of clearance.

# Single sign-on

- Other services that the service requestor contacts, therefore, do not need to perform authentication and authorization steps.
- Instead, upon receiving a request, they simply contact the issuing authority to ask for the authentication and authorization clearance it originally obtained.
- The issuing authority provides this information in the form of assertions that communicate the security details.
- (The two types of assertions that contain authentication and authorization information are simply called authentication assertions and authorization assertions.)

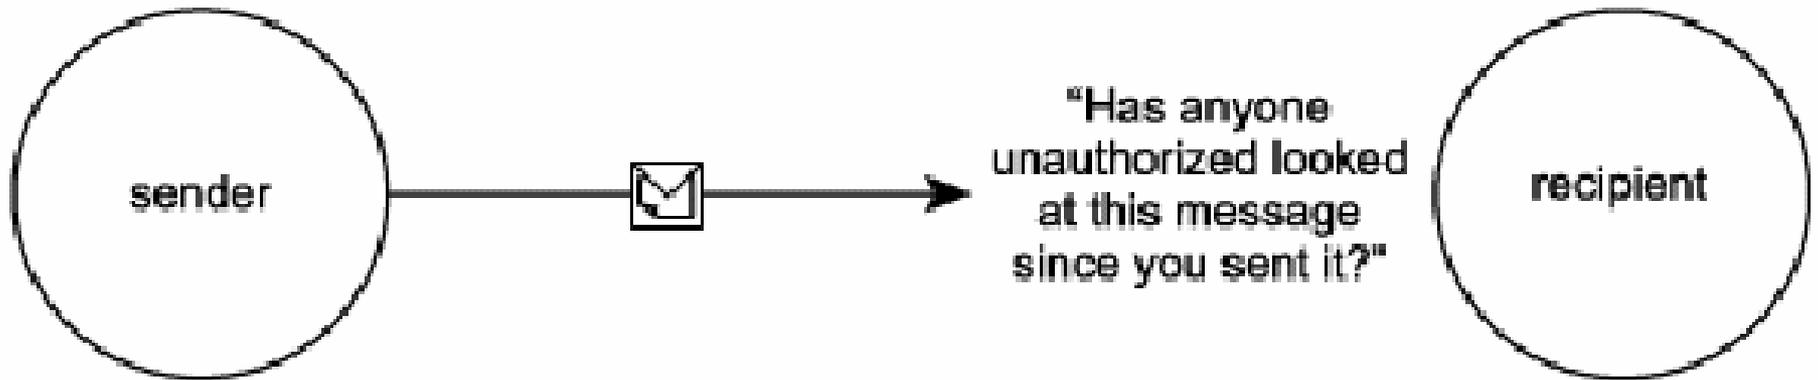
A basic message exchange demonstrating single sign-on (in this case, as implemented by SAML).



# Confidentiality and integrity

- Confidentiality is concerned with protecting the privacy of the message contents.
- A message is considered to have remained confidential if no service or agent in its message path not authorized to do so viewed its contents.

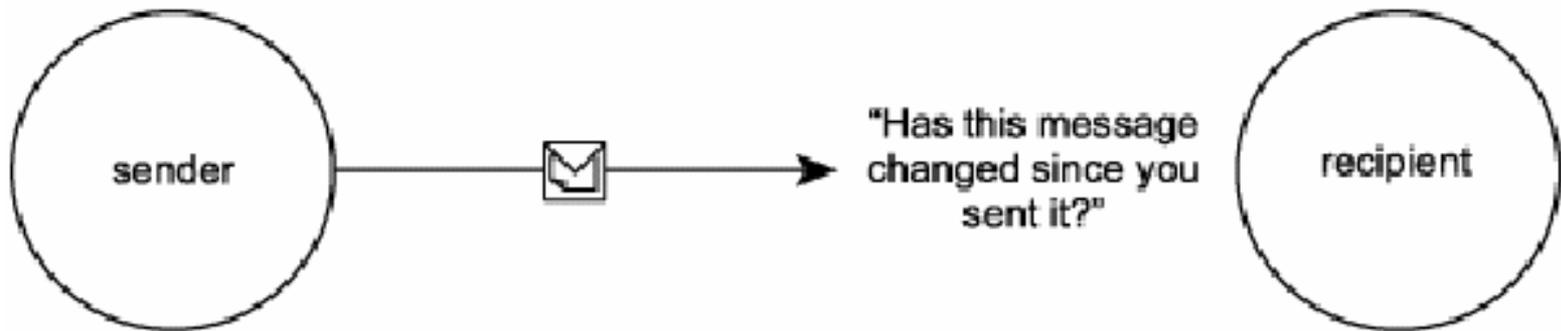
Confidentiality means that the privacy of the message has been protected throughout its message path



# Confidentiality and integrity

- Integrity ensures that a message has not been altered since its departure from the original sender
- This guarantees that the state of the message contents remained intact from the time of transmission to the point of delivery.

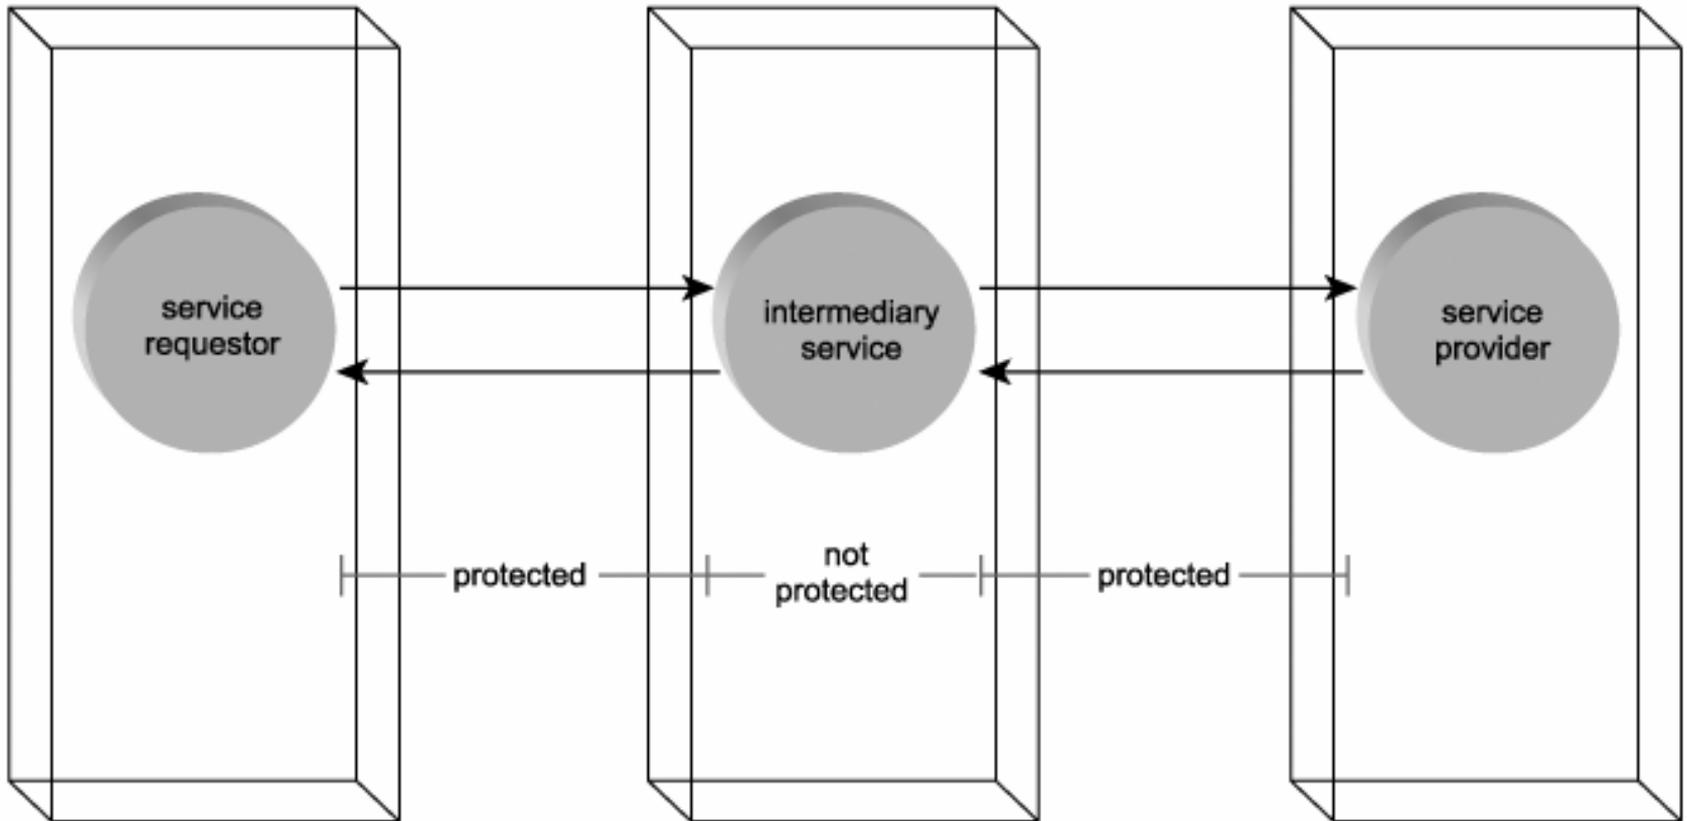
Integrity means ensuring that a message's contents have not changed during transmission



# Transport-level security and message-level security

- The type of technology used to protect a message determines the extent to which the message remains protected while making its way through its message path.
- Secure Sockets Layer (SSL), for example, is a very popular means of securing the HTTP channel upon which requests and responses are transmitted.
- However, within a Web services-based communications framework, it can only protect a message during the transmission between service endpoints.
- Hence, SSL only affords us transport-level security.

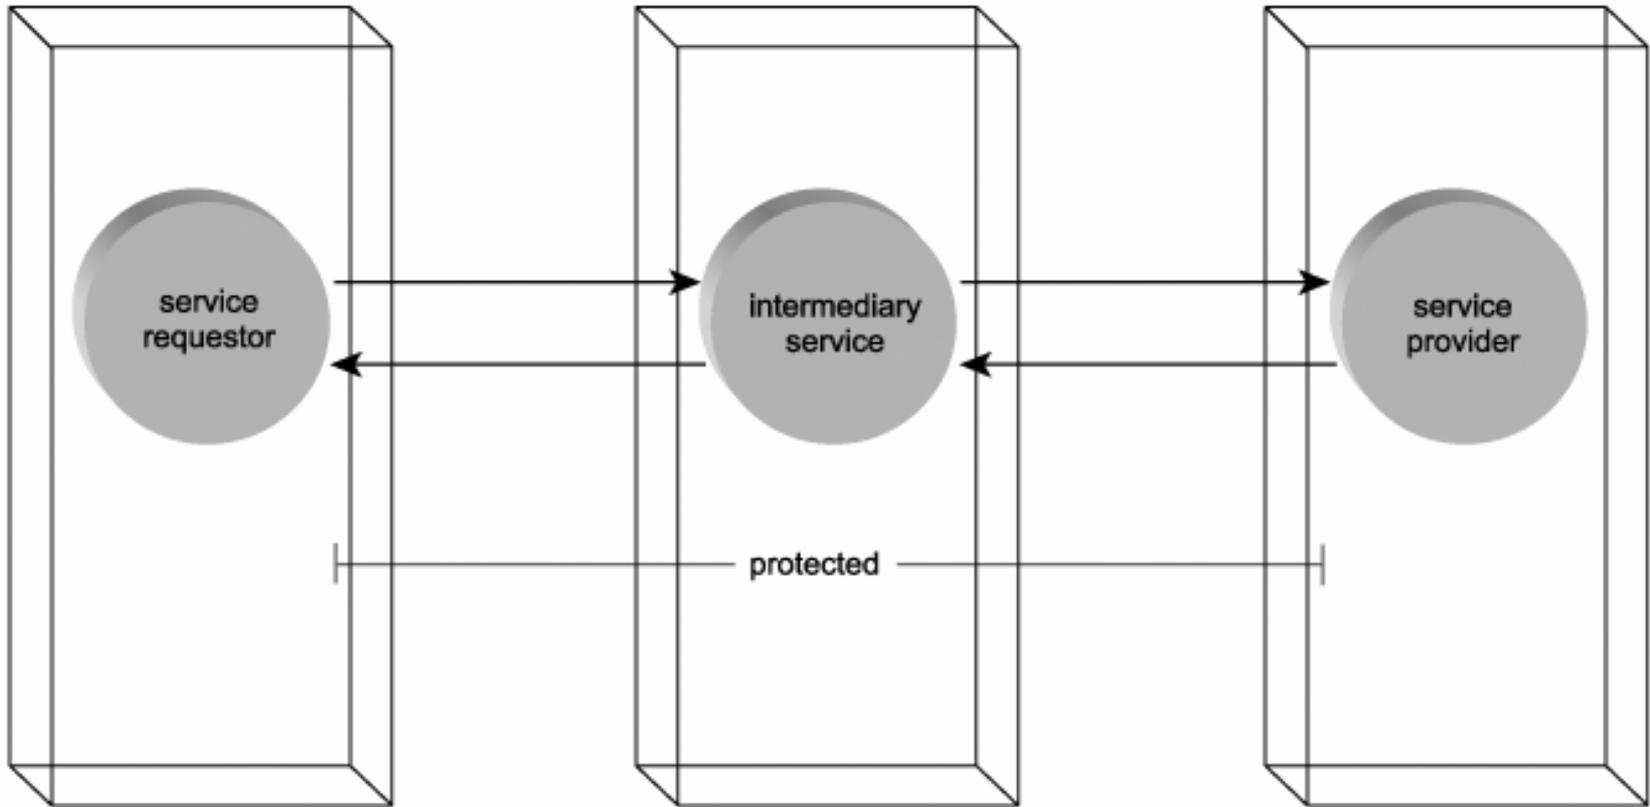
Transport-level security only protects the message during transit  
between service endpoints



# Transport-level security and message-level security

- If, for example, a service intermediary takes possession of a message, it still may have the ability to alter its contents. To ensure that a message is fully protected along its entire message path, message-level security is required.
- In this case, security measures are applied to the message itself (not to the transport channel on which the message travels).
- Now, regardless of where the message may travel, the security measures applied go with it.

# Message-level security guarantees end-to-end message protection



# Encryption and digital signatures

- Message-level confidentiality for an XML-based messaging format, such as SOAP, can be realized through the use of specifications that comprise the WS-Security framework.
- In this section we focus on XML-Encryption and XML-Signature, two of the more important WS-Security extensions that provide security controls that ensure the confidentiality and integrity of a message.
- XML-Encryption, an encryption technology designed for use with XML, is a cornerstone part of the WS-Security framework.
- It provides features with which encryption can be applied to an entire message or only to specific parts of the message (such as the password).

# Encryption and digital signatures

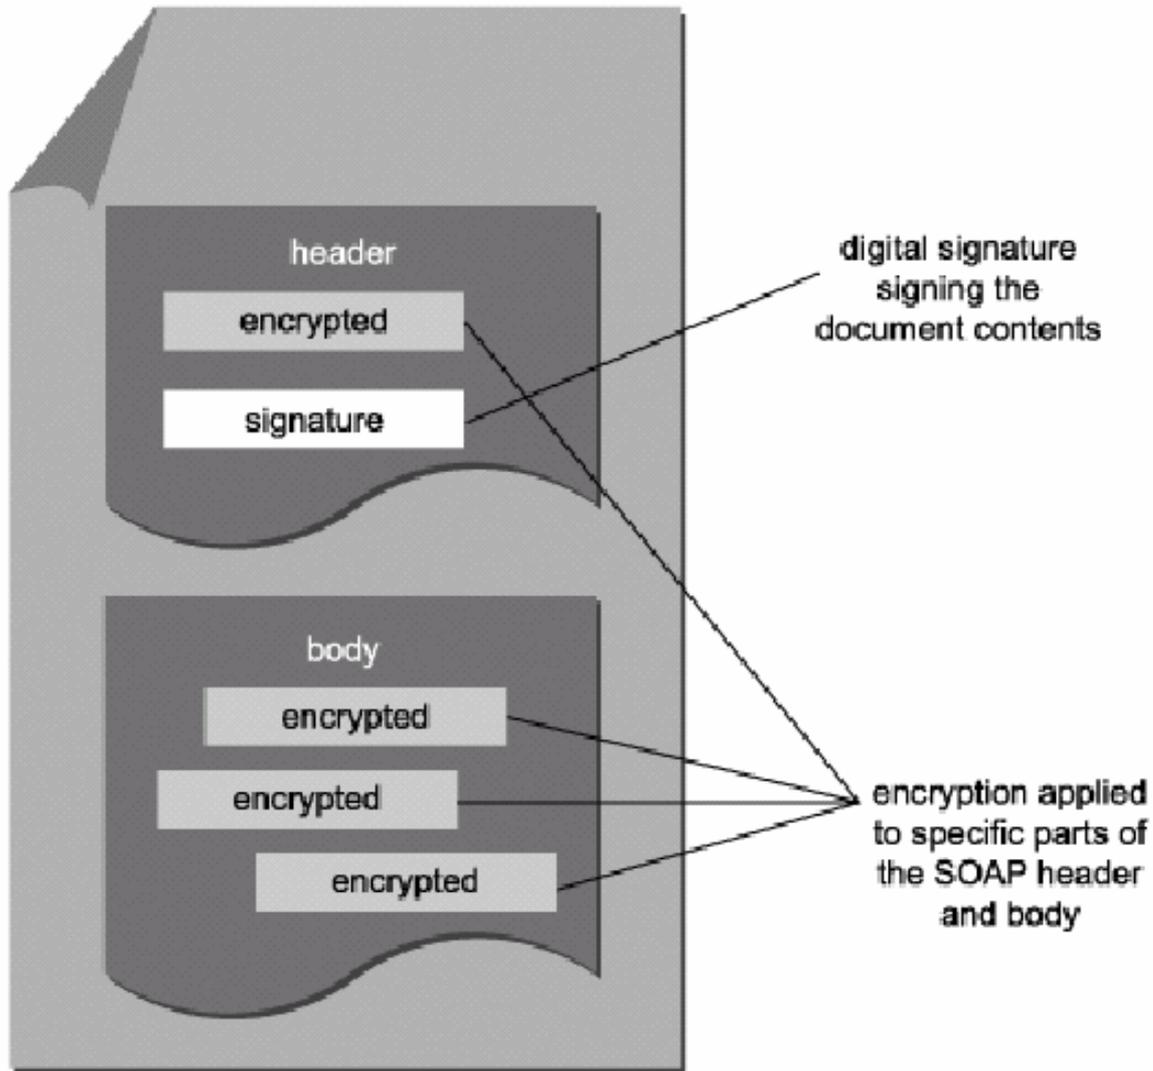
- To ensure message integrity, a technology is required that is capable of verifying that the message received by a service is authentic in that it has not been altered in any manner since it first was sent.
- XML-Signature provides features that allow for an XML document to be accompanied by a special algorithm-driven piece of information that represents a digital signature.
- This signature is tied to the content of the document so that verification of the signature by the receiving service only will succeed if the content has remained unaltered since it first was sent.

# Encryption and digital signatures

## Note

- Digital signatures also support the concept of non-repudiation, which can prove that a message containing a (usually legally binding) document was sent by a specific requestor and delivered to a specific provider.
- As illustrated in the next figure, XML-Encryption can be applied to parts of a SOAP header, as well as the contents of the SOAP body.
- When signing a document, the XML-Signature can reside in the SOAP header.

# A digitally signed SOAP message containing encrypted data



# Encryption and digital signatures

## Note

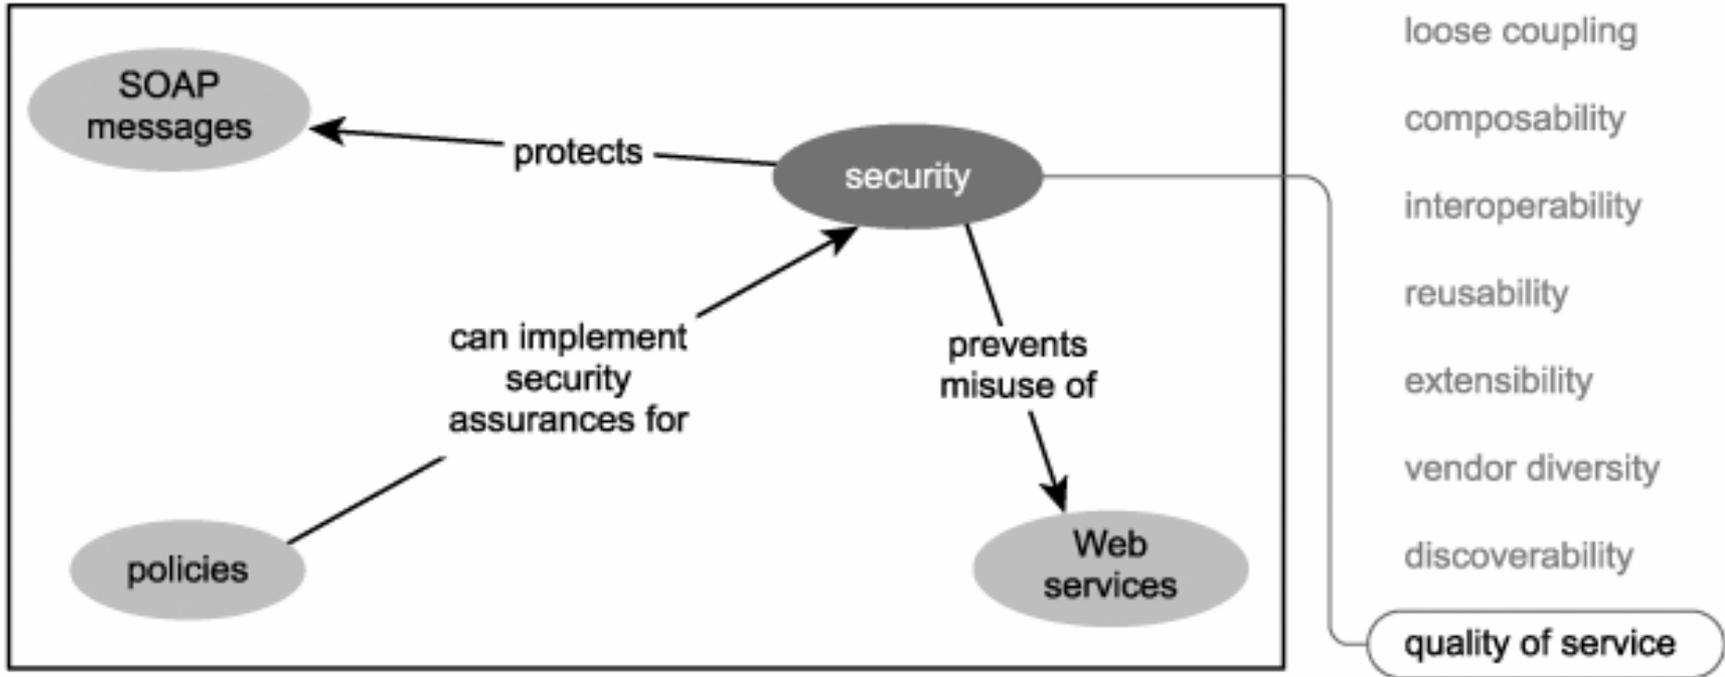
- Both encryption and digital signature technologies rely on the use of keys.
- These are special values used to unlock the algorithm upon which encryption and digital signatures are based.
- Shared keys are typically used by encryption technologies and require that both the sender and receiver of a message use the same key.
- Public/private key pairs are commonly used by digital signature technologies, where the message sender signs the document with a key that is different from the one used by the recipient. (One of the keys is public, but the other is private.)

# Security and SOA

- Message-level security can clearly become a core component of service-oriented solutions.
- Security measures can be layered over any message transmissions to either protect the message content or the message recipient.
- The WS-Security framework and its accompanying specifications therefore fulfill fundamental QoS requirements that enable enterprises to:
  - utilize service-oriented solutions for the processing of sensitive and private data
  - restrict service access as required
- As shown in the next figure, the security framework provided by WS-Security also makes use of the WS-Policy framework explained earlier (a separate specification called WS-SecurityPolicy provides a series of supporting policy assertions).

# Security, as it relates to policies, SOAP messages, and Web services

service-oriented architecture



# Case Study

- TLS has a message-level security policy that applies to any business documents sent to its B2B solution.
- The policy has the following rules:
  - Any values residing in documents sent via SOAP messages must be encrypted.
  - Any invoice submitted to TLS with a total value of over Rs. 10,00,000 must also be digitally signed.
- To comply with this policy, RailCo is required to apply XML-Encryption to the parts of the invoice message sent by the Invoice Submission Service that contain monetary values.
- It further embeds a business rule into the Invoice Submission Service's underlying logic that checks for invoice totals that exceed the Rs. 10,00,000 mark. Those that do, have their corresponding SOAP message documents digitally signed using XML-Signature.

# Summary of Key Points

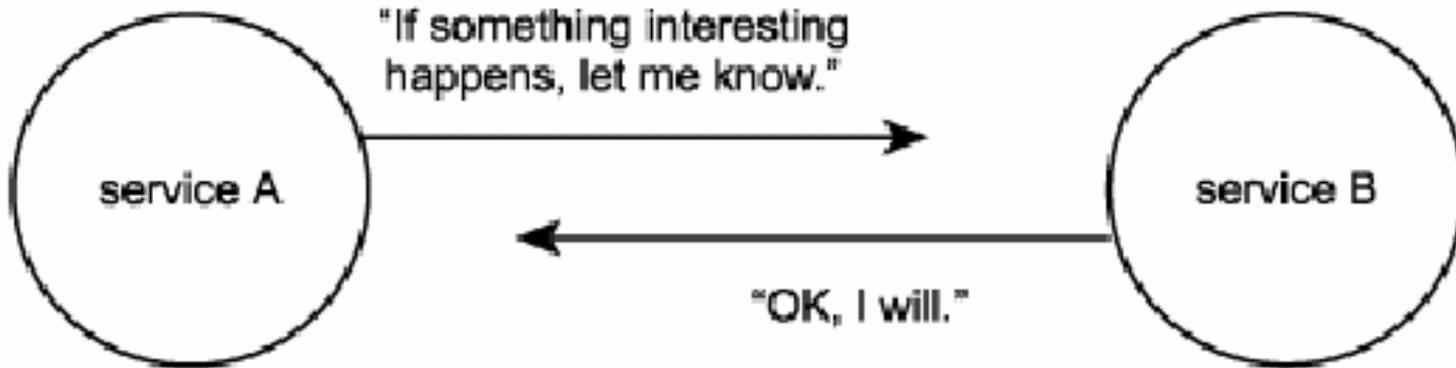
- Security within SOA is a multi-faceted subject matter that encompasses the feature set of numerous specifications. The WS-Security framework governs a subset of these specifications, and establishes a cohesive and composable security architecture.
- The primary aspects of security addressed by these specifications are identification, authentication, authorization, integrity, and confidentiality, as well as non-repudiation.
- Two primary technologies for preserving the integrity and confidentiality of XML documents are XML-Encryption and XML-Signature.

# Notification and eventing

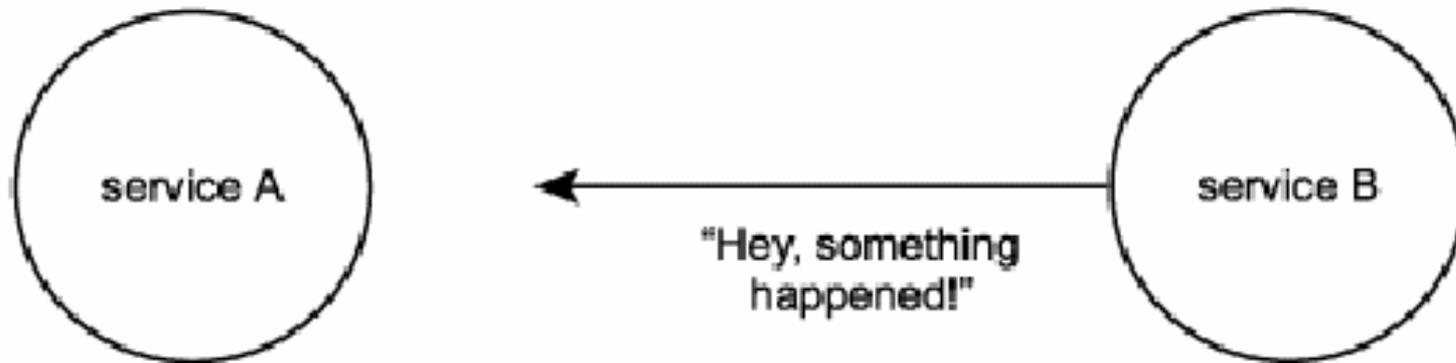
# Notification and eventing

- With its roots in the messaging-oriented middleware era, the publish-and-subscribe MEP introduces a composite messaging model, comprised of primitive MEPs that implement a push delivery pattern.
- It establishes a unique relationship between service providers and service requestors where information is exchanged (often blindly) to achieve a form of dynamic notification.
- While notification itself can be applied to different types of MEPs, the focus of this section is a discussion of notification within the context of the publish-and-subscribe pattern.

Once subscribed, service A is notified of anything service B publishes that is of interest to service A



later ...



# Publish-and-subscribe in abstract

- This messaging pattern can be classified as a complex MEP assembled from a series of primitive MEPs.
- It involves a publisher service that makes information categorized by different topics available to registered subscriber services.
- Subscribers can choose which topics they want to register for, either by interacting with the publisher directly or by communicating with a separate broker service.
- A topic is an item of interest and often is tied to the occurrence of an event.

# Publish-and-subscribe in abstract

- When a new piece of information on a given topic becomes available, a publisher broadcasts this information to all those services that have subscribed to that topic.
- Alternatively, a broker service can be used to perform the broadcast on the publisher's behalf.
- This decouples the publisher from the subscriber, allowing each to act independently and without knowledge of each other.

# One concept, two specifications

- Two major implementations of the publish-and-subscribe pattern exist:
  - The WS-Notification framework
  - The WS-Eventing specification
- Spearheaded by IBM and Microsoft respectively, these use different approaches and terminology to cover much of the same ground.
- It is expected that a single publish-and-subscribe specification eventually will emerge as an industry standard.
- The remainder of this section is dedicated to exploring features of both specifications.

# The WS-Notification Framework

- As with other WS-\* frameworks, what is represented by WS-Notification is a family of related extensions that have been designed with composability in mind.
  - **WS-BaseNotification** Establishes the standardized interfaces used by services involved on either end of a notification exchange.
  - **WS-Topics** Governs the structuring and categorization of topics.
  - **WS-BrokeredNotification** Standardizes the broker intermediary used to send and receive messages on behalf of publishers and subscribers.

# Situations, notification messages, and topics

- The notification process typically is tied to an event that is reported on by the publisher.
- This event is referred to as a situation.
- Situations can result in the generation of one or more notification messages.
- These messages contain information about (or relating to) the situation, and are categorized according to an available set of topics.
- Through this categorization, notification messages can be delivered to services that have subscribed to corresponding topics.

# Notification producers and publishers

- So far we've been using the familiar "publisher" and "subscriber" terms to describe the roles services assume when they participate in the publish-and-subscribe pattern.
- Within WS-Notification, however, these terms have more distinct definitions.
- The term publisher represents the part of the solution that responds to situations and is responsible for generating notification messages.
- However, a publisher is not necessarily required to distribute these messages.
- Distribution of notification messages is the task of the notification producer.
- This service keeps track of subscriptions and corresponds directly with subscribers. It ensures that notification messages are organized by topic and delivered accordingly.

# Notification producers and publishers

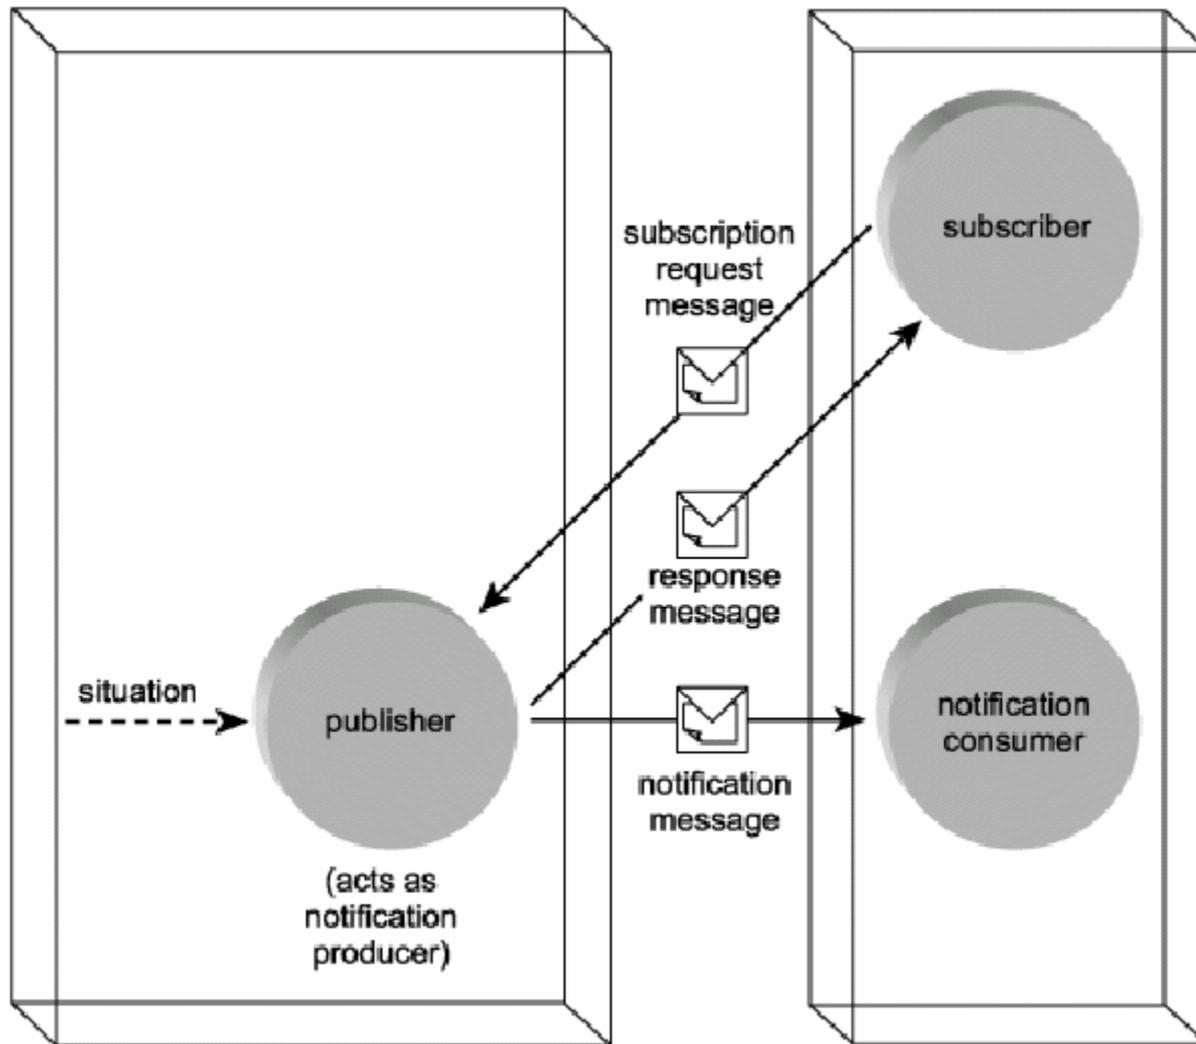
Note that:

- A publisher may or may not be a Web service, whereas the notification producer is always a Web service.
- A single Web service can assume both publisher and notification producer roles.
- The notification producer is considered the service provider.

# Notification consumers and subscribers

- A subscriber is the part of the application that submits the subscribe request message to the notification producer.
- This means that the subscriber is not necessarily the recipient of the notification messages transmitted by the notification producer.
- The recipient is the notification consumer, the service to which the notification messages are delivered

# A basic notification architecture



# Notification consumers and subscribers

Note that:

- A subscriber does not need to exist as a Web service, but the notification consumer is a Web service.
- Both the subscriber and notification consumer roles can be assumed by a single Web service.
- The subscriber is considered the service requestor.

## Notification broker, publisher registration manager, and subscription manager

- To alleviate the need for direct contact between the two groups of services we described in the previous two sections, a set of supplementary services is available.

### **The notification broker**

- A Web service that acts on behalf of the publisher to perform the role of the notification producer.
- This isolates the publisher from any contact with subscribers.
- Note that when a notification broker receives notification messages from the publisher, it temporarily assumes the role of notification consumer.

# Notification consumers and subscribers

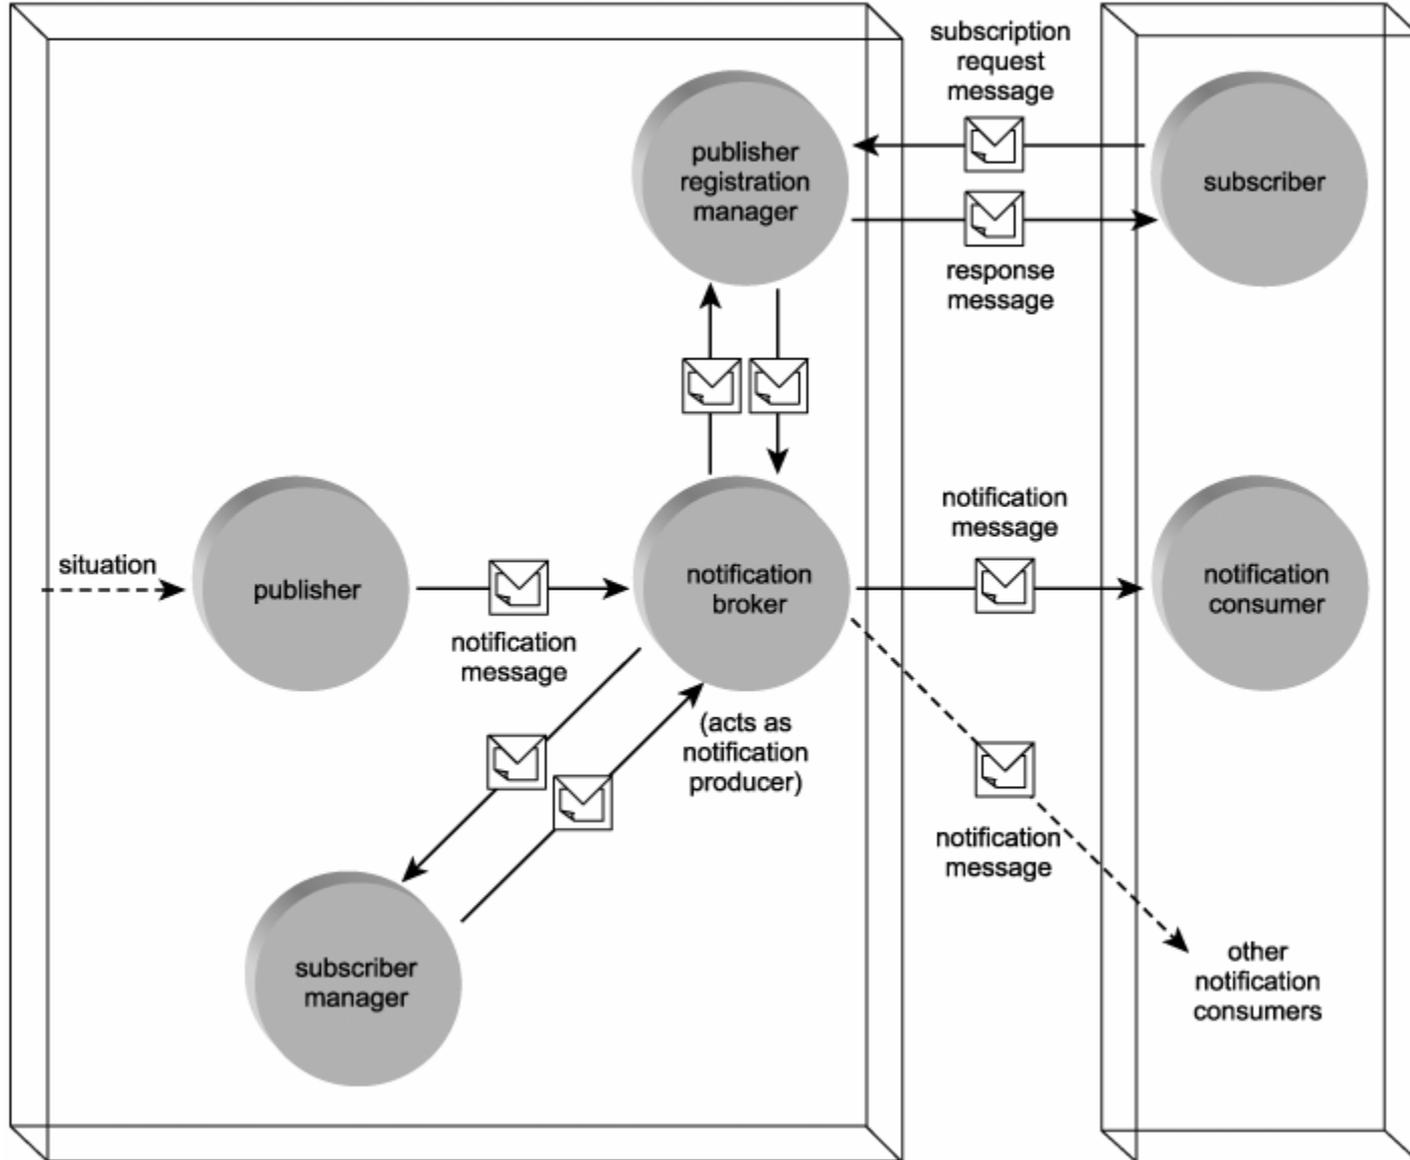
## **The publisher registration manager**

- A Web service that provides an interface for subscribers to search through and locate items available for registration.
- This role may be assumed by the notification broker, or it may be implemented as a separate service to establish a further layer of abstraction.

## **The subscription manager**

- A Web service that allows notification producers to access and retrieve required subscriber information for a given notification message broadcast.
- This role also can be assumed by either the notification producer or a dedicated service.

# A notification architecture including a middle tier



# The WS-Eventing specification

- As its name implies, WS-Eventing addresses publish-and-subscribe requirements by focusing on an event-oriented messaging model.
- When an event related to one Web service occurs, any other services that have expressed interest in the event are subsequently notified.
- Following are brief explanations of the terms and concepts expressed by the WS-Eventing specification.

# The WS-Eventing specification

## **Event sources**

- The term "publisher" is never actually mentioned in the WS-Eventing specification.
- Instead, its role is assumed by a broader-scoped Web service, known as the event source.
- This part of the eventing architecture is responsible for both receiving subscription requests and for issuing corresponding notification messages that report information about occurred events.

# The WS-Eventing specification

## **Event sinks and subscribers**

- On the subscription end of the eventing model, separate Web services manage the processing of notification and subscription messages.
- An event sink is a service designed to consume (receive) notification messages from the event source.
- Subscribers are services capable of issuing various types of subscription requests.

# The WS-Eventing specification

## **Subscription managers**

- An event source, by default, assumes the responsibility of managing subscriptions and transmitting notifications.
- In high volume environments it may be desirable to split these roles into separate services.
- To alleviate the demands on the event source, intermediate services, known as subscription managers, optionally can be used to distribute publisher-side processing duties.

# The WS-Eventing specification

## **Notification messages and subscription end messages**

- When an event occurs, it is reported by the event source via the issuance of a notification message (also called an event message).
- These are standard SOAP messages that contain WS-Eventing-compliant headers to convey event details.
- WS-Eventing allows for an expiry date to be attached to subscriptions.
- This requires that subscribers issue renewal requests for the subscription to continue (as discussed in the next section).
- If a subscription is left to expire, though, it is the event source that often is expected to send a special type of notification to the corresponding event sink, called a subscription end message.

# The WS-Eventing specification

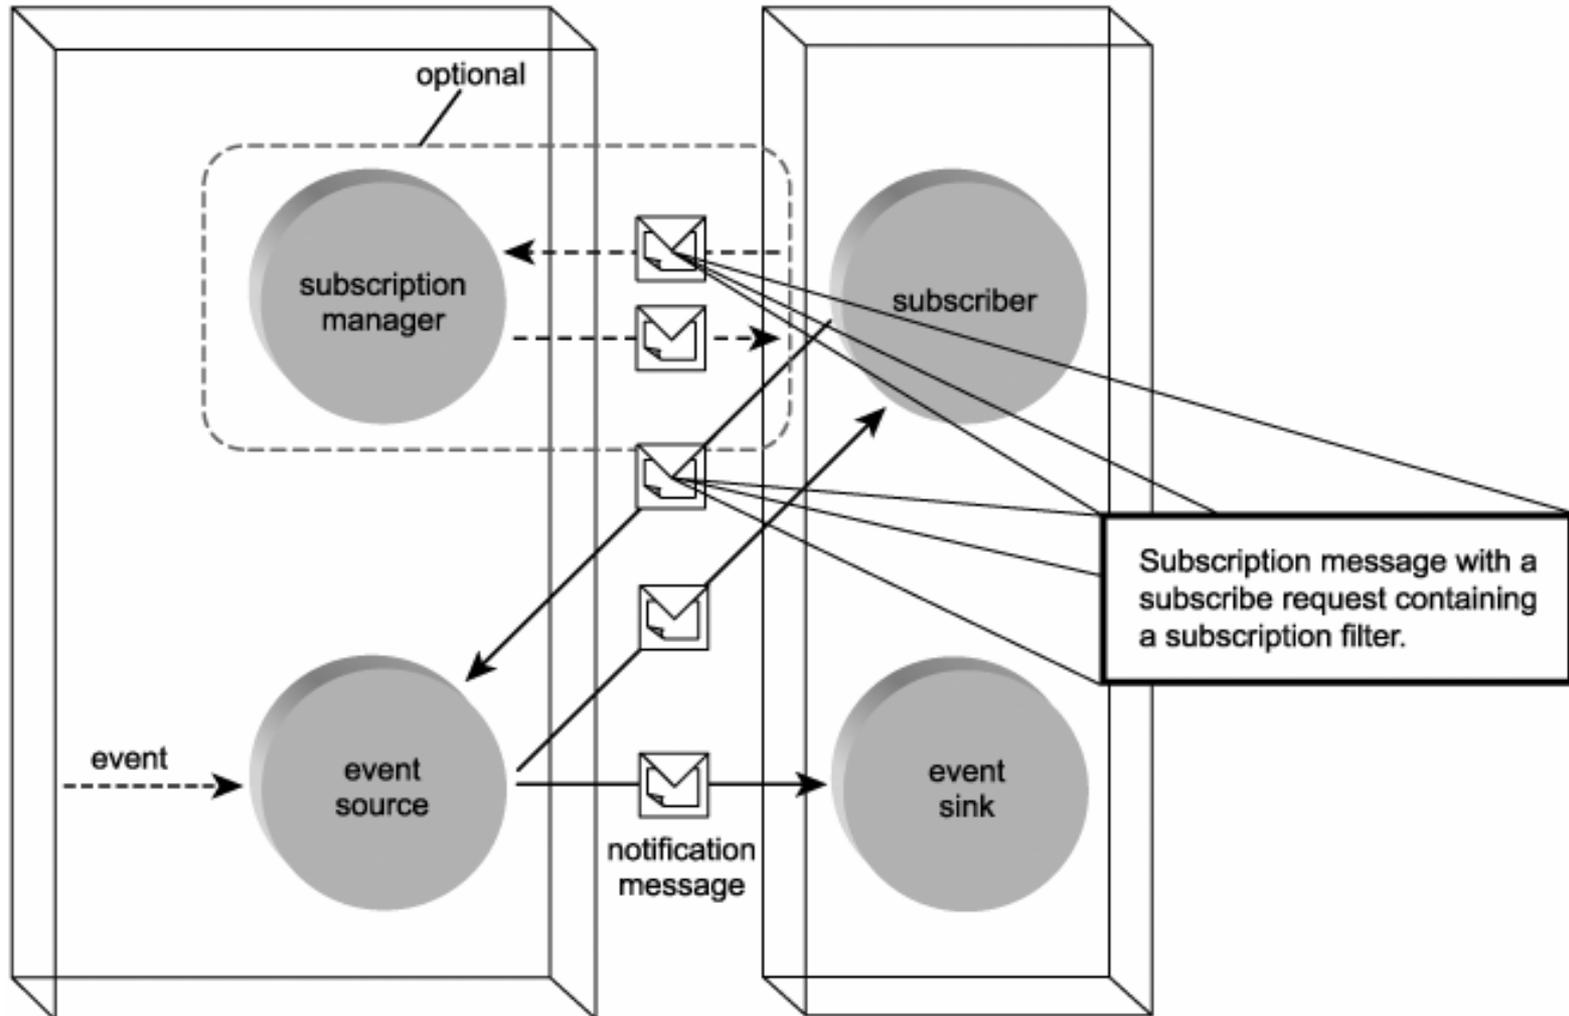
## Subscription messages and subscription filters

- Subscribers issue subscription messages directly to the event source or to an intermediate subscription manager. Different types of subscription-related requests can be transmitted via subscription messages.
- The following specific requests are supported:
  - **Subscribe Requests** that a new subscription be created. (Note that this message also contains the filter details, as well as the endpoint destination to which a subscription end message is to be delivered. Filters are described shortly.)
  - **Unsubscribe Requests** that an existing subscription be canceled.
  - **Renew Requests** that an existing subscription scheduled to expire be renewed.
  - **GetStatus Requests** that the status of a subscription be retrieved.

# The WS-Eventing specification

- For a subscriber to communicate that the event sink (on behalf of whom it is submitting the subscription request) is only interested in certain types of events, it can issue a subscription message containing a subscription filter.
- If the event source does not support filtering (or if it cannot accommodate the requested filter), the subscription request is denied.
- The relationships between the subscription manager, event source, subscriber, and event sink are shown in the next figure

# A basic eventing architecture



# WS-Notification and WS-Eventing

- The fact that these two specifications currently provide overlapping feature sets is no indication that this will remain so in the future.
- It has been speculated that the reason these specifications were created separately was because the individual sponsors had diverging requirements.
- One of IBM's goals is to incorporate WS-Notification with its grid computing initiatives.
- Microsoft, on the other hand, is expected to utilize WS-Eventing within its system administration platform.

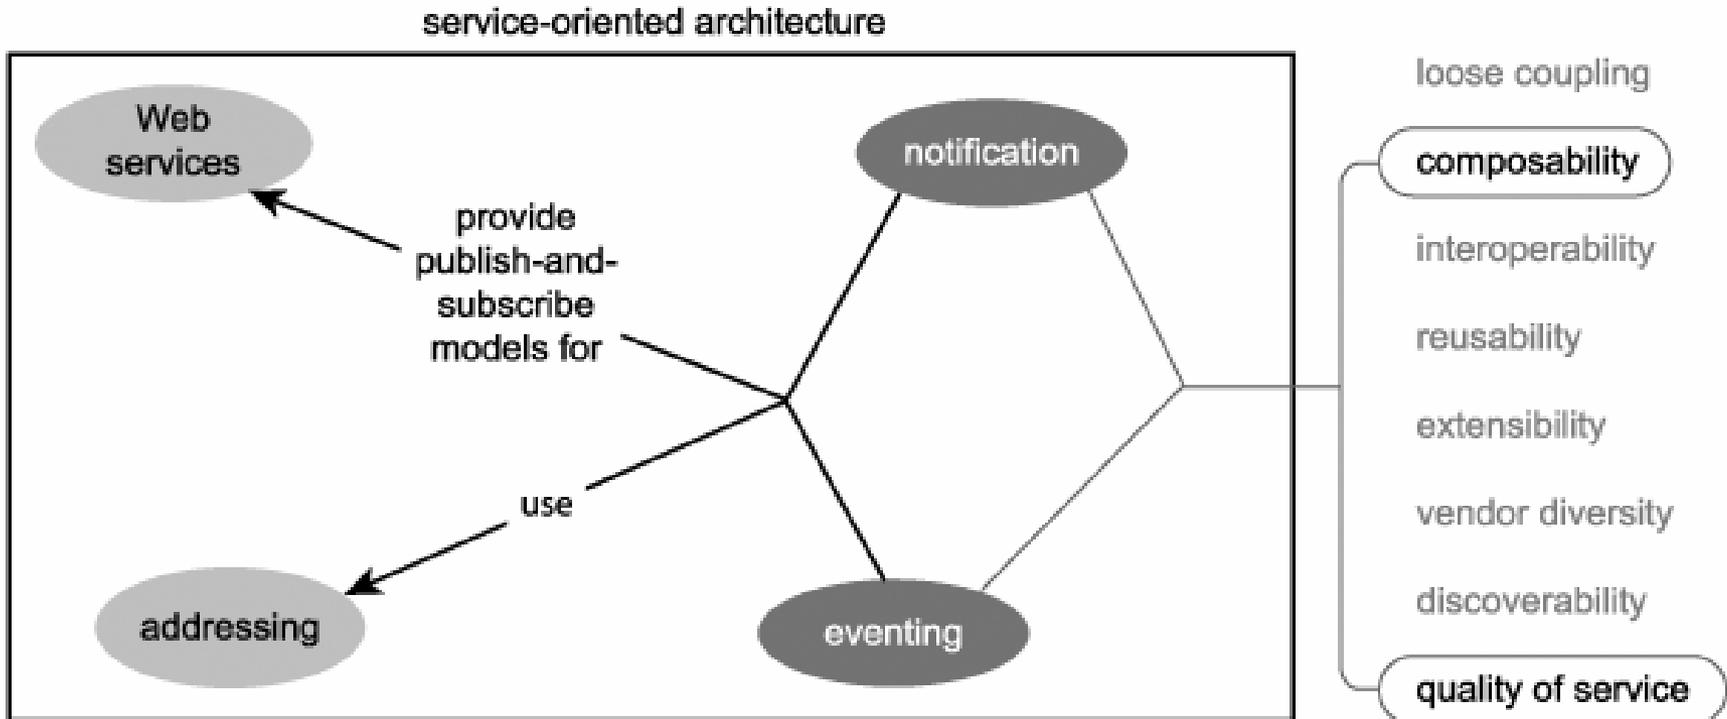
# WS-Notification and WS-Eventing

- In an effort to continue promoting interoperability across proprietary platforms, IBM recently joined the WS-Eventing effort.
- It is entirely within the realm of possibilities that either specification will be modified to align with the other or that the vendors involved will come to an agreement on how to establish a single notification extension that will meet their collective requirements.
- Language descriptions for these two specifications are available at [www.specifications.ws](http://www.specifications.ws).)

# Notification, eventing, and SOA

- By implementing a messaging model capable of supporting traditional publish-and-subscribe functionality, corresponding legacy features now can be fully realized within service-oriented application environments.
- Moreover, the ability to weave a sophisticated notification system into service-oriented solutions can significantly broaden the applicability of this messaging model (as evidenced by the before mentioned plans to incorporate notification with grid computing).

# Notification and eventing establishing standardized publish-and-subscribe models within SOA



# Notification, eventing, and SOA

- Service-oriented solutions can increase QoS characteristics by leveraging notification mechanisms to perform various types of event reporting.
- For example, performance and exception management related events can trigger notification broadcasts to potential service requestors (subscribers), informing them of a variety of conditions.

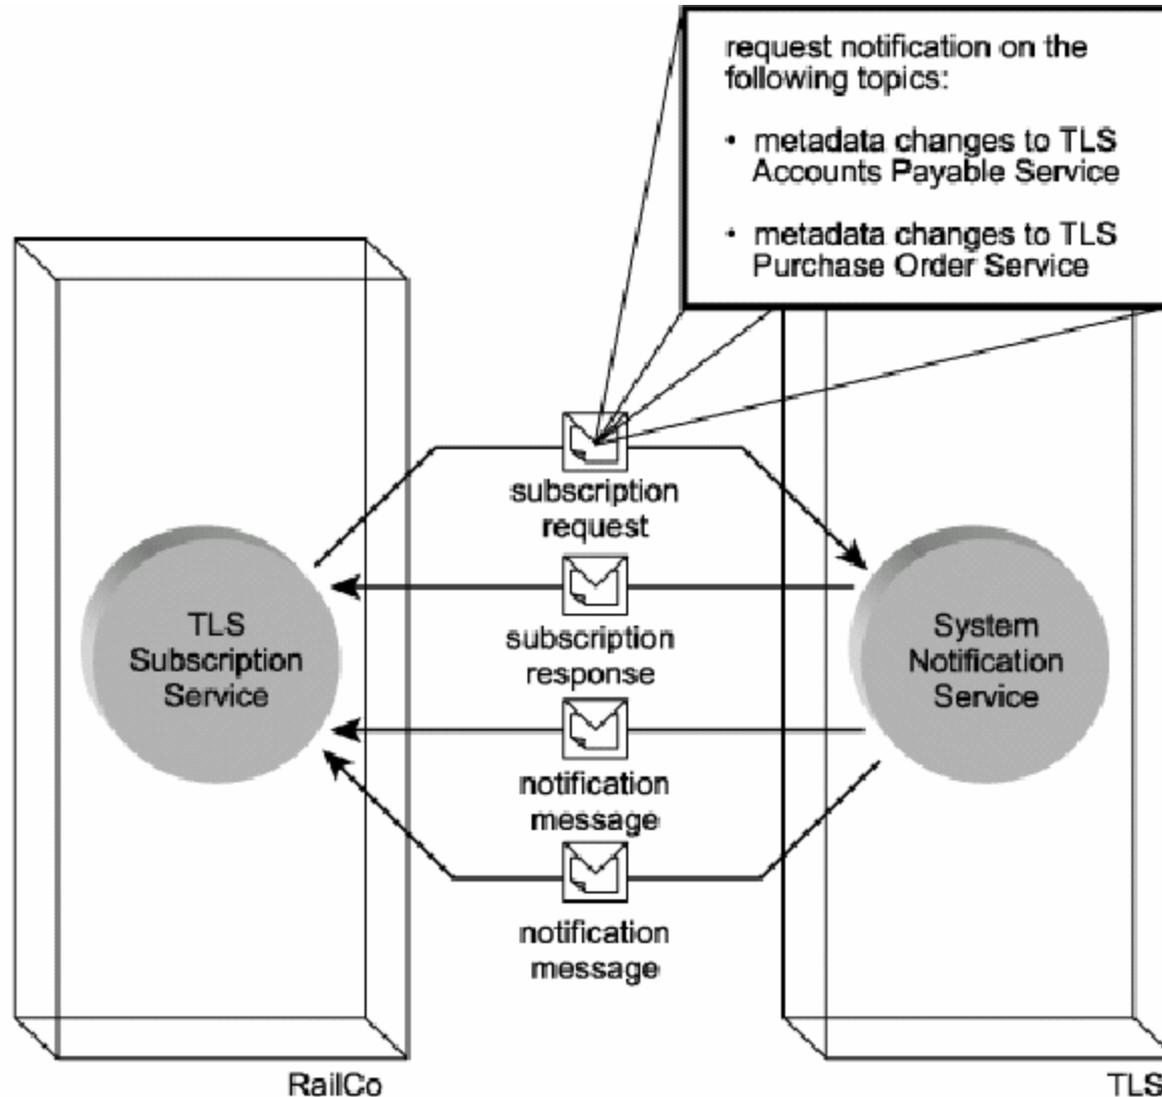
# Case Study

- In response to a series of complaints from vendors who experienced message transmission problems that resulted from changes to TLS service descriptions, TLS has decided to supplement their existing metadata exchange support by implementing a notification system. Now, business partners will be forewarned of any upcoming changes that might impact their systems.
- There are many services that comprise the TLS B2B solution. Each performs a specific function that involves one or more types of partners. Not all partners need to interact with every TLS service. As a result, the notification system is set up in such a manner that partners are able to subscribe to notifications relating to specific TLS services or groups of services.

# Case Study

- For this, TLS has provided a dedicated System Notification Service that acts as the publisher of notification messages. Partners are consequently required to implement their own subscriber services. Each notification message essentially requests that the recipient initiate a WS-MetadataExchange against the provided TLS endpoint(s).
- RailCo creates a separate subscription service to interact with the TLS System Notification Service. Unfortunately, RailCo calls its new service the "TLS Subscription Service," which is sure to lead to confusion in the future. Regardless, RailCo uses its service to subscribe to and receive notifications relating to the two primary services with which it interacts on a regular basis: the TLS Accounts Payable and Purchase Order Services

# The new RailCo subscription service allows RailCo to receive notifications from the TLS System Notification Service



# Summary of Key Points

- The traditional publish-and-subscribe messaging model can be implemented with the WS-Notification framework or the WS-Eventing specification.
- WS-Notification consists of the WS-BaseNotification, WS-Topics, and WS-BrokeredNotification specifications that collectively establish a subscription and notification system.
- The WS-Eventing specification provides similar functionality but is based on a moderately different architecture.
- Notification and eventing realize the popular publish-and-subscribe messaging model within SOA. The sophisticated messaging environment provided by SOA, in turn, introduces new opportunities to leverage these notification mechanisms.