



**EUROPEAN COMMISSION**  
DIRECTORATE-GENERAL  
INFORMATICS  
Information systems Directorate

# **European Commission**

## **e-TrustEx**

### **Software Architecture Document**

Date:	04/12/2013
Version:	2.01
Authors:	Sandro D'Orazio, Andras Imre, Olivier Derveau
Revised by:	Didier Thunus
Approved by:	Chrysanthi Giortsou, Tanya Chetcuti
Public:	
Reference Number:	

## TABLE OF CONTENTS

<b>1. INTRODUCTION .....</b>	<b>4</b>
1.1. Purpose .....	4
1.2. Scope .....	4
1.3. References .....	4
1.4. Definitions .....	5
1.5. Document Content Overview.....	6
<b>2. ARCHITECTURAL REPRESENTATION .....</b>	<b>7</b>
<b>3. ARCHITECTURAL GOALS AND CONSTRAINTS .....</b>	<b>7</b>
<b>4. SECURITY.....</b>	<b>8</b>
4.1.1. Introduction .....	8
4.1.2. Confidentiality .....	8
4.1.3. Authentication.....	8
4.1.4. Authorisation .....	9
4.1.5. Integrity .....	9
4.1.6. Validity .....	9
4.1.7. Auditing .....	9
4.1.8. Non Repudiation.....	9
4.1.9. Storage Security.....	10
<b>5. USE-CASE VIEW.....</b>	<b>11</b>
5.1.1. Selection Rationale .....	11
<b>6. LOGICAL VIEW .....</b>	<b>12</b>
6.1.1. Overview.....	12
6.1.2. Architecturally Significant Design Packages .....	12
6.1.3. Use-Case Realizations .....	13
<b>7. DEPLOYMENT VIEW.....</b>	<b>20</b>
<b>8. IMPLEMENTATION VIEW .....</b>	<b>21</b>
8.1. Overview.....	21
<b>9. DATA VIEW.....</b>	<b>23</b>
9.1. Data Model .....	23
9.2. State Machines.....	24
9.2.1. Introduction .....	24
9.2.2. Generic State Machines .....	25
9.2.3. The document bundle state machine .....	25
<b>10. SIZE AND PERFORMANCE .....</b>	<b>26</b>
10.1. Size .....	26

10.2. Performance .....	26
<b>11. QUALITY.....</b>	<b>26</b>
11.1. Extensibility .....	26
11.2. Reliability .....	26
11.3. Portability .....	27

## Document History

Version	Date	Comment	Modified Pages
2.01	04/12/2013	Version 2.01 (final)	All
2.00	15/11/2013	Submission of version 2.00 for review	All

## 1. INTRODUCTION

### 1.1. Purpose

This document provides a comprehensive architectural overview of the system, using a number of different architectural views to depict different aspects of the system. It is intended to capture and convey the significant architectural decisions that have been made on the system.

### 1.2. Scope

The architecture described in this document concerns the system e-TrustEx developed and used by the European Commission as part of the ISA programme.

### 1.3. References

#	Document	Contents outline
[REF1]	<b>GIP</b> of 07/02/2008 by DIGIT/B4	Global Implementation Plan written for the IDABC PEGCSO, and its annexes. Annex I: <b>Feasibility Study</b> ; Annex II: <b>Requirements Catalogue</b> .
[REF2]	e-PRIOR EDI validation by DIGIT/B4	Validation of the e-PRIOR system against the e-Invoicing and e-Archiving requirements, as mentioned in the EU Council Directive 2001/115/EC, now incorporated in Council directive 2006/112/EC.
[REF3]	e-PRIOR Interface Control Document v1.000 by DIGIT/B4 of 06/02/2009	Interface control document for front-offices and external systems
[REF4]	Test Management Plan by DIGIT/B4	Description of the test strategy and implementation of e-PRIOR.
[REF5]	<a href="#">Enterprise Integration Patterns</a>	Patterns and best practices for Enterprise Integration by Gregor Hohpe
[REF6]	<a href="#">Spring Integration</a>	Extension of the Spring programming model to support the well-known Enterprise Integration Patterns
[REF7]	<a href="#">Spring Web Services</a>	Spring Web Services is a product of the Spring community focused on creating document-driven Web services
[REF8]	<a href="#">Schematron</a>	Rule-based validation language for making assertions about the presence or absence of patterns in XML trees
[REF9]	<a href="#">JBoss</a> Application Server	Java EE certified open source application server
[REF10]	<a href="#">EJB</a>	Enterprise Java Beans
[REF11]	<a href="#">REST</a>	Representational State Transfer
[REF12]	<a href="#">SOAP</a>	Simple Object Access Protocol
[REF13]	<a href="#">WS-Transfer</a>	Web Services Transfer : This specification describes a general SOAP-based protocol for

		accessing XML representations of Web service-based resources
[REF14]	<a href="#">WS-Security</a>	Standard set of SOAP [SOAP11, SOAP12] extensions that can be used when building secure Web services to implement message content integrity and confidentiality
[REF15]	<a href="#">WS-Policy</a>	The Web Services Policy Framework (WS-Policy) provides a general purpose model and corresponding syntax to describe the policies of a Web Service
[REF16]	<a href="#">WS-SecurityPolicy</a>	WS-Policy defines a framework for allowing web services to express their constraints and requirements. These are security related policies
[REF17]	<a href="#">X.509</a>	Public-Key Infrastructure (X.509)
[REF18]	<a href="#">JAX-WS</a>	JAX-WS Web-service API definition
[REF19]	<a href="#">Metro</a>	The Metro Web-service stack (JAX-WS implementation)
[REF20]	<a href="#">CEN-BII</a>	CENBII specification is a meant to facilitate effective public procurement solution with focus on cross-border interoperability.

#### 1.4. Definitions

#	Document	Contents outline
[REF21]	Document	A Document in the scope of the e-TrustEx platform is a piece of information structured using XML
[REF22]	MTOM	<i>Message Transmission Optimization Mechanism</i>
[REF23]	Transaction	A Transaction in the scope of the e-TrustEx platform represents an operation on a document. The submission of a document bundle through the system is an e-TrustEx transaction
[REF24]	Profile	A document exchange profile aggregates a set of transactions.
[REF25]	Party	A Party is an entity exchanging documents through the platform.
[REF26]	Bundle	Set of multiple documents exchanged by the Parties playing a role in a data exchange scenario.
[REF27]	Document wrapper	A Document Wrapper is an entity composed of a binary file and its metadata, which can be exchanged through the platform
[REF28]	Interchange agreement	Contract between two or more Parties on the use of e-TrustEx for the electronic exchange of information, in the context of a one or more Profile.

## **1.5. Document Content Overview**

After summarizing the architectural representation, goals and constraints, this document describes the system using several architectural views (Use Case, logical, process, deployment, implementation and data) and then concludes with size, performance and quality considerations.

## 2. ARCHITECTURAL REPRESENTATION

The next two sections of the document will describe architectural goals and constraints because they have important consequences on the system architecture.

Architecturally relevant Use Cases will be presented with a Use Case diagram and a short explanation detailing its impact on the architecture. The following views will also be provided:

- A logical view will provide the "big picture" view of the platform presenting the structure of the system through its components and their interactions.
- An implementation view will describe the software layers decomposition and the main software components. A component diagram will be used in this view.
- A deployment view will provide a description of the hardware components and how they are linked together. This view gives a technical description of protocols and hardware nodes used.
- A data view will provide information about data persistency. A class diagram will be used to model main system data.

UML diagrams will be used systematically to represent the different views of the system.

## 3. ARCHITECTURAL GOALS AND CONSTRAINTS

The following non-functional requirements that can impact the architectural solution have been identified:

Non-functional requirement	Description
Support of large files	The application shall allow the transfer of large files (up to 100 MB each) of any type
Support of groups of files	The application shall allow the transfer of groups of linked files (up to 500 files)
Interoperability	The application shall allow the connection of heterogeneous systems
Confidentiality	The system shall ensure that transferred documents are not viewed by anybody else than the sender and the final recipient
Integrity	The system shall guarantee integrity of the transferred files
File storage	A dedicated file system shall be used to store binaries



## 4. SECURITY

### 4.1.1. Introduction

The e-TrustEx platform supports requirements for authentication, integrity and validity of documents both during transmission and storage. Besides those, it also includes measures for authorization, confidentiality, auditing and non-repudiation.

The security requirements supported by e-TrustEx are described as follows:

- *Confidentiality* is needed to prevent third parties from eavesdropping on information that is being transmitted.
- *Authentication* ensures that the parties involved in communication are really who they say they are.
- *Authorization* ensures that users only have access to the resources they allowed to.
- *Integrity* guarantees that a message is not modified during its transmission.
- *Validity* certifies a stored message does not lose its legal attributes.
- *Auditing* controls enable relevant parties like authorized bodies to inspect transactions afterwards.
- *Non-repudiation* measures prevent users from denying actions they have undertaken.

To provide an answer to these requirements the platform provides the means to link Confidentiality, Integrity and Availability (C.I.A.) levels to a document interchange agreement between parties. The Retrieve Interchange agreement service allows the sender to retrieve the C.I.A. information before sending a document.

The following chapters of the document describe how the security requirements are implemented in the platform.

### 4.1.2. Confidentiality

The connection to the platform web services must be done using HTTPS. HTTPS is a secure version of the HTTP protocol and is being used to protect data transactions for commercial and financial purposes. It uses a Secure Socket Layer (SSL) and digital certificates to encrypt a data transfer session over an otherwise insecure HTTP connection.

The use HTTPS guarantees transport level confidentiality however there might be a need to also ensure it at message level. The platform provides support for end-to-end encryption acting as public key repository (used to encrypt the message). It will also provide the sender the information about the Confidentiality level of the message exchange through the Retrieve ICA service.

### 4.1.3. Authentication

By default, the system requires basic authentication over the HTTPS connection. The fact that HTTPS is used to set up the initial secure connection makes it difficult for someone to steal the passwords by listening on the communication as it is encrypted

If stronger authentication is required, implementers can decide to impose the usage of two-way SSL connections to access the platform web-services. In that case, the client will have to provide a certificate trusted by the server to establish the SSL connection and thus authenticate to the server using this certificate. This configuration is external to the platform as it depends on the implementer's IT infrastructure. This explains why e-TrustEx cannot provide out of the box solution for two-way SSL.

#### 4.1.4. Authorisation

The platform comes with a complete authorisation implementation based on the concept of interchange agreements.

- A **Document exchange profile** is a set of *transactions* that define operations (e.g. submission, retrieval, etc.) on specific documents.
- An **Interchange Agreement** models the contract between parties in the context of a specific document exchange profile.

Every time a party tries to access the system, e-TrustEx checks if there is an interchange agreement authorising the caller to do so.

#### 4.1.5. Integrity

The integrity of message exchange is guaranteed by the platform through the use of XML digital signatures. The messages returned by the platform are signed and parties can be required to sign their calls to the platform web-services. The use of XML digital signatures is an optional feature and can be activated through configuration of the system. The use of XML signature is not mandatory as it adds complexity to the solution and depending on the business case HTTPS encryption might be judged good enough to ensure the required level of integrity.

This integrity is only valid for the transport and the platform's signatures are purely technical and have no "business" value. The platform does not provide out of the box support for end to end signature however if such business signatures are required the platform will provide the information to sender parties through the retrieve ICA services.

#### 4.1.6. Validity

The platform supports two types of message validation, the standard XSD validation and the schematron validation, a rule-based validation language for making assertions about the presence or absence of patterns in XML trees. This guarantees the business validity of the messages stored in the system.

#### 4.1.7. Auditing

All the calls to the platform are logged and transactions are stored in the system database. Auditors can be given temporary access to the platform database and log files in case of control.

#### 4.1.8. Non Repudiation

When a party submits a message, e-TrustEx generates a signed acknowledgement containing information about the submitted message. This signed acknowledgement can be used as a proof of submission ensuring the non-repudiation of (submission of) messages sent through the platform. The party submitting the message is responsible for storing the acknowledgement generated by e-TrustEx.

#### *4.1.9. Storage Security*

The databases that contains business data, logging information and audit trails will be kept in a secure environment where only authorized users have access. Audit trails are kept to ensure that the data is not tampered with. Additionally, it is assumed that regular backups of the databases are generated and securely archived.

The large attachments files are encrypted before being stored on the file system. In addition to the file encryption proper segregation of duties policies shall be applied to ensure that people having access to the encryption key cannot access the file system and vice versa.

## 5. USE-CASE VIEW

This section provides a representation of the architecturally significant use cases.

### 5.1.1. Selection Rationale

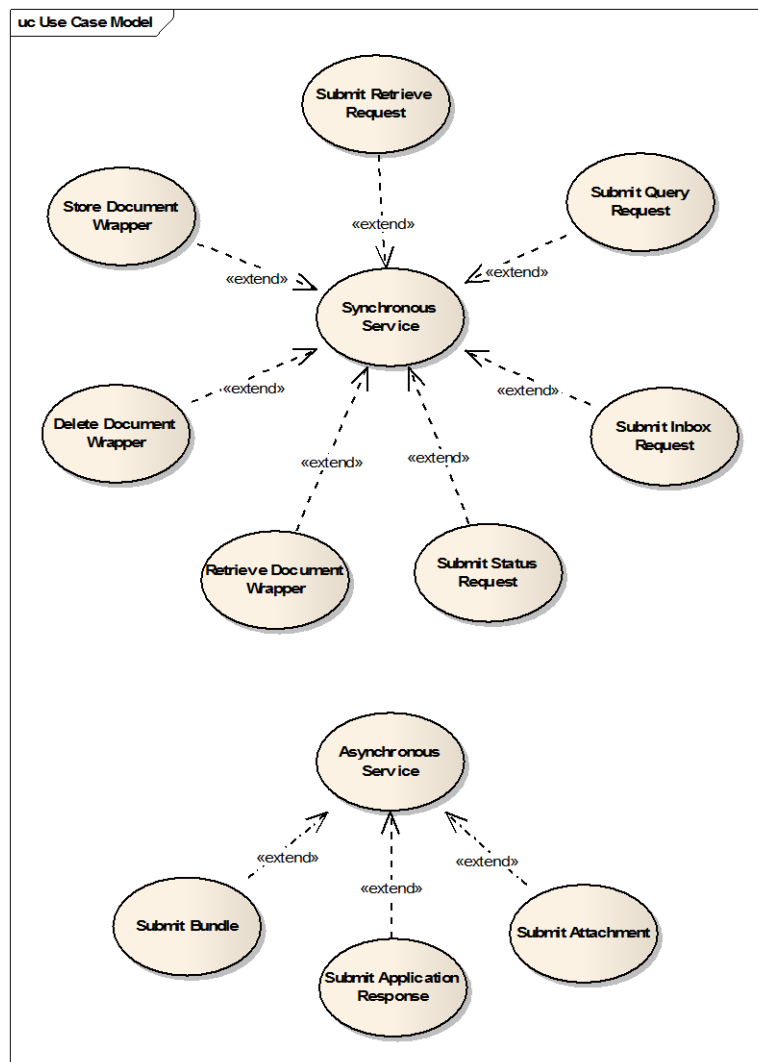
The architecturally significant use cases have been selected based on the following criteria:

- Use cases affecting multiple components of the system, thereby providing a cross-cutting view of the system architecture;
- Use cases representing critical parts of the architecture, thereby addressing the technical risks of the project at an earlier stage.

The following use cases have been selected.

- Use case addressing the synchronous services offered by the platform
- Use case addressing the submission of (potentially large) binary files (UC Store Document Wrapper)
- Use case addressing the asynchronous services offered by the platform

The following Use Cases diagram displays the e-TrustEx use cases.



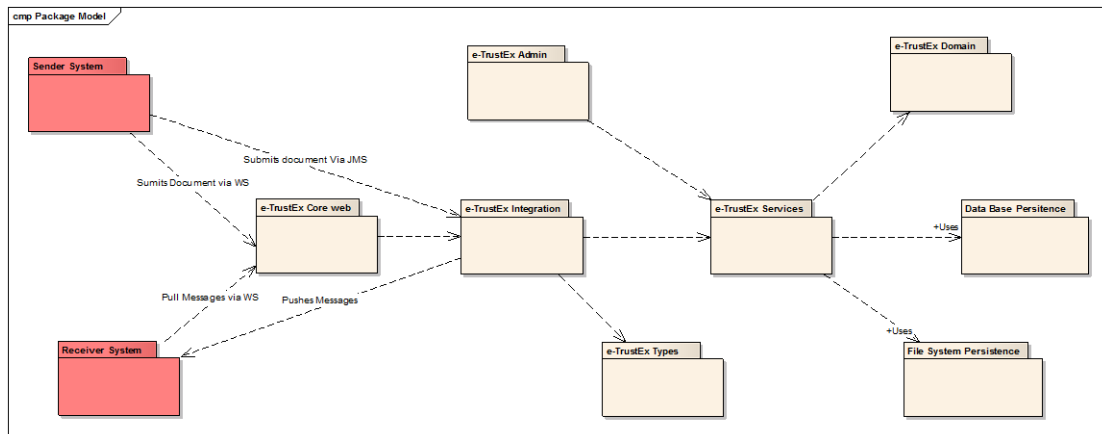
## 6. LOGICAL VIEW

### 6.1.1. Overview

This chapter describes the main application modules, how they interact and how they implement system use cases.

### 6.1.2. Architecturally Significant Design Packages

The following diagram provides a high-level view of the main packages composing the system. The Database persistence and file system persistence are logical packages representing the physical data storages used by the platform. The others represents different application layers and give an overview of the organisation of the platform's code as each of them translates into a separate maven project.



### Sender and Receiver systems

Sender systems can submit documents using the platform web-services or by pushing them into JMS queues.

Receiver systems can either retrieve their messages polling the platform read services or receive them via the e-TrustEx store and forward mechanisms. E-TrustEx can forward messages via JMS or by calling a "call-back" web service implemented by the receiver system.

### e-TrustEx core web layer

This module is a Java J2EE web application offering web services interfaces to the clients of the platform. It offers a generic SOAP endpoint using the Spring WS framework and specific JAX-WS web service dealing with the large attachments (Using MTOM and Http Chunking).

### e-TrustEx Integration

This is the core of the message-processing platform. This module is a lightweight ESB based on the Spring integration framework that provides common message processing building blocks for checking authorisation, storing and validating the incoming documents and finally route them to the receiver system.

### e-TrustEx types

This package contains all the java classes generated from the different WSDLs and XSDs used by the system. These JAXB annotated objects are used by the integration layer to parse incoming requests and to build responses.

### e-TrustEx Domain

This is the platform domain layer; it contains EJB3 entities implementing the domain objects. This layer is further described in the domain model description chapter of this document.

### e-TrustEx Services

This is the platform services layer that provides access to the database and to the file system to the integration, the core web layer and the CIPAdmin component. Services are POJOs that are configured in the dependent components using the Spring Framework Dependency Injection mechanism.

### e-TrustEx Storage (file system and database)

The platform uses a database to store its configuration and the incoming messages. The file storage is used to store the binary files submitted to the platform.

### e-TrustEx Admin (a.k.a. CIPAdmin)

The administration web console of the platform meant to facilitate the configuration of the system. This is a standard java J2EE web application using Spring Model View Controller pattern implementation and JQuery for the presentation layer.

#### 6.1.3. Use-Case Realizations

The following use cases have been chosen to describe how software packages behave:

- A sender system stores a document wrapper (binary file)
- A Party calls a Synchronous service
- A Party calls an Asynchronous service

There are two kinds of services implemented in the platform:

- the asynchronous services that are used by a party to send a document (e.g. submit bundle service),
- the synchronous services that are used to handle large binaries or to query and retrieve documents like store document wrapper or retrieve request.

As any use case of the system derives from the generic synchronous and asynchronous use cases, they will be detailed in this chapter of the document. Storing of a large binary by the system is also relevant for this chapter as, even if it is also a synchronous service, it slightly deviates from the generic one.

As mentioned in the previous chapter e-TrustEx uses the Spring integration framework. This framework is an implementation of Gregor Hohpe Enterprise Integration Patterns (EIP). The following chapters will use EIP notation diagrams to describe the message processing flows.

Here are the basic concepts for messaging systems as define on the [EIP website](#) that will come in handy to understand how the platform works:

***Channels*** — *Messaging applications transmit data through a Message Channel, a virtual pipe that connects a sender to a receiver. A newly installed messaging system doesn't contain any channels; you must determine how your applications need to communicate and then create the channels to facilitate it.*

***Messages*** — *A Message is an atomic packet of data that can be transmitted on a channel. Thus to transmit data, an application must break the data into one or more packets, wrap each packet as a message, and then send the message on a channel. Likewise, a receiver application receives a message and must extract the data from the message to process it. The*

message system will try repeatedly to deliver the message (e.g., transmit it from the sender to the receiver) until it succeeds.

**Multi-step delivery** — In the simplest case, the message system delivers a message directly from the sender's computer to the receiver's computer. However, actions often need to be performed on the message after it is sent by its original sender but before it is received by its final receiver. For example, the message may have to be validated or transformed because the receiver expects a different message format than the sender. A Pipes and Filters architecture describes how multiple processing steps can be chained together using channels.

**Routing** — In a large enterprise with numerous applications and channels to connect them, a message may have to go through several channels to reach its final destination. The route a message must follow may be so complex that the original sender does not know what channel will get the message to the final receiver. Instead, the original sender sends the message to a Message Router, an application component and filter in the pipes-and-filters architecture, which will determine how to navigate the channel topology and direct the message to the final receiver, or at least to the next router.

**Transformation** — Various applications may not agree on the format for the same conceptual data; the sender formats the message one way, yet the receiver expects it to be formatted another way. To reconcile this, the message must go through an intermediate filter, a Message Translator, that converts the message from one format to another.

**Endpoints** — An application does not have some built-in capability to interface with a messaging system. Rather, it must contain a layer of code that knows both how the application works and how the messaging system works, bridging the two so that they work together. This bridge code is a set of coordinated Message Endpoints that enable the application to send and receive messages."

The format of the Messages exchanged through the e-TrustEx platform is XML.

#### 6.1.3.1. Common message processing components

Some "message processing" steps are common to all types of messages processed by the platform. In order to avoid code duplication these common behaviours have been implemented using reusable components. These components need to access the platform service layer and are thus implemented as EIP Service Activators.

The **Authorisation service activator** is responsible for assessing if a party is allowed to call a specific service. In order to do that it uses the credentials provided to the system through HTTP basic authentication and retrieves the Party linked to these credentials (see Section 9.1), the *CallerParty*. The system then tries to extract the *SenderParty* from the incoming message. If this required piece of information is missing, or the specified Sender Party is not known by the system, the platform will return an unauthorised access error to the caller.

The e-TrustEx platform offers support for third parties, meaning that one party can delegate to another one the right to perform transactions on the system. This is why the service activator checks if the *CallerParty* and the *SenderParty* are the same and, if they are different and no delegation exists between them, the platform will return an unauthorised access error to the caller.

If a *ReceiverParty* is specified in the incoming message, the system will query its database to determine if there is an interchange agreement involving the two parties pointing to a profile containing the transaction related to the occurring service call. If no *Interchange Agreement* is found the platform will return an unauthorised access error to the caller. Specifying a *ReceiverParty* is mandatory for all the asynchronous service.

If the *ReceiverParty* is not present in the incoming message, the system checks if there is an *Interchange Agreement* involving the *SenderParty* pointing to a profile containing the transaction related to the occurring service call without taking into account the receiver and will again return an unauthorised access error to the caller if nothing can be found.

The ***XSD Validation Service Activator*** is a configurable component that applies XSD validation on the incoming XML message. The XSD is linked to a document or a transaction in the system metadata. The platform allows the user to either provide an URL pointing to the XSD or to store it directly in the system database. The XSD validation failure will either trigger a SOAP fault if executed synchronously or lead to the generation of an Application Response error message if executed asynchronously.

An *Application response* is an XML document generated by the system to notify the caller asynchronously of errors that occurred during the message processing or, more generally speaking, to notify changes in the message state (e.g. a message moving to the ERROR state). This message is also the means provided by the platform to parties willing to modify the state of message using the *Submit Application Response* service. Readers can find more information on the topic in the data view chapter of this document (Section 9.1).

The ***Schematron Service Activator*** :

*"Schematron is a rule-based validation language for making assertions about the presence or absence of patterns in XML trees. It is a structural schema language expressed in XML using a small number of elements and XPath .In a typical implementation, the Schematron schema XML is processed into normal XSLT code for deployment anywhere that XSLT can be used"*

This service activator applies the schematron XSLT transformation to the incoming XML document and handles possible errors the same way than the XSD Validation service activator does. The XSL templates are uploaded into the system database and linked to the platform Documents.

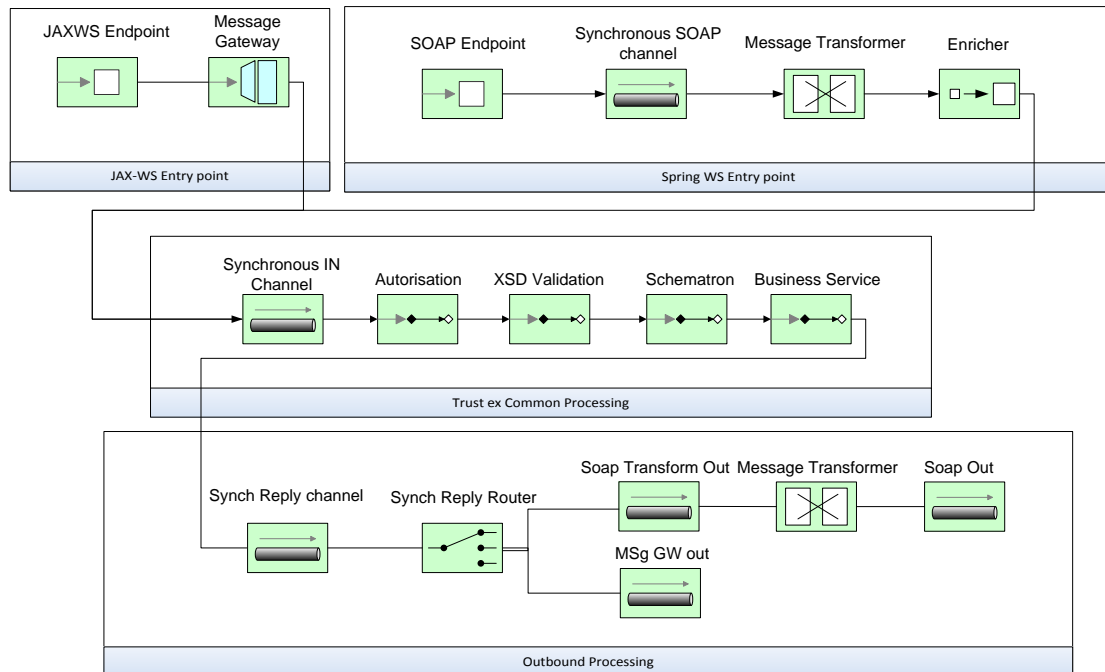
Schematron is used to implement more business validation rules on incoming XML messages that are not supported by standard XSD validation.

The ***Business Service Activator*** is responsible for calling the java object implementing the specific business logic for the service. These must be configured in the platform Spring application context and implement specific platform java interfaces. The Business service activator retrieves the business service using either a naming convention or using metadata that can be stored in the platform database.



### 6.1.3.2. Synchronous services

This is the EIP representation of the message flow for synchronous services:



The platform offers two entry points to provide synchronous service. They are both SOAP based but use different web services implementation.

The Spring WS entry point uses Spring WS framework to provide a generic SOAP endpoint for synchronous services. It routes incoming messages to the appropriate channel based on the message content (content based routing). Incoming messages are routed to the Synchronous SOAP channel where they are transformed into e-TrustEx platform messages, a specialisation of the Generic Spring integration Message by the Message Transformer component.

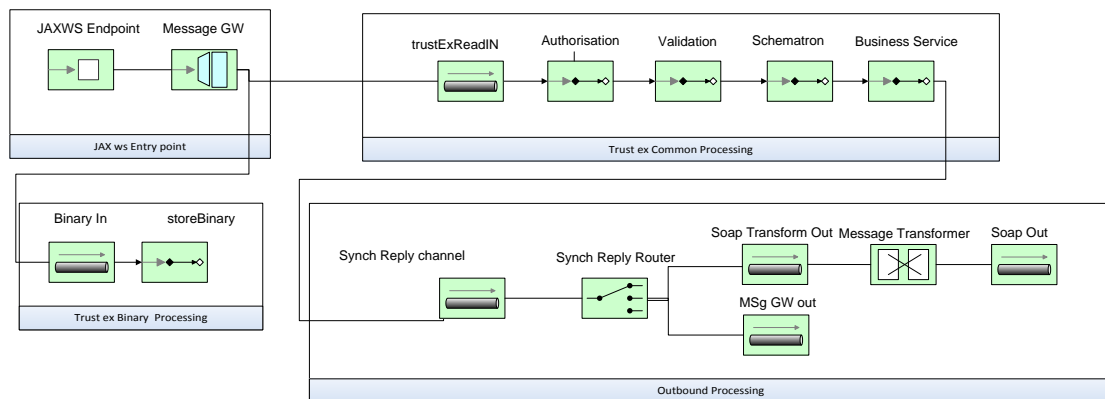
This entry point directly deals with the XML avoiding to bind it to java objects. The Message Transformer runs a set of configurable XQueries on the incoming XML to extract relevant information like the message ID, the sender and receiver party IDs and performs some basic integrity checks. The extracted information set in the platform message header.

Some routing metadata is the added to the header by the Header enricher component. This metadata is used by the platform to determine to which channel the response must be submitted.

The incoming message, after the initial transformation into a platform message is routed to the e-TrustEx common processing message handling chain where it is validated and processed by the business service linked with the service being called

### 6.1.3.3. Store document wrapper

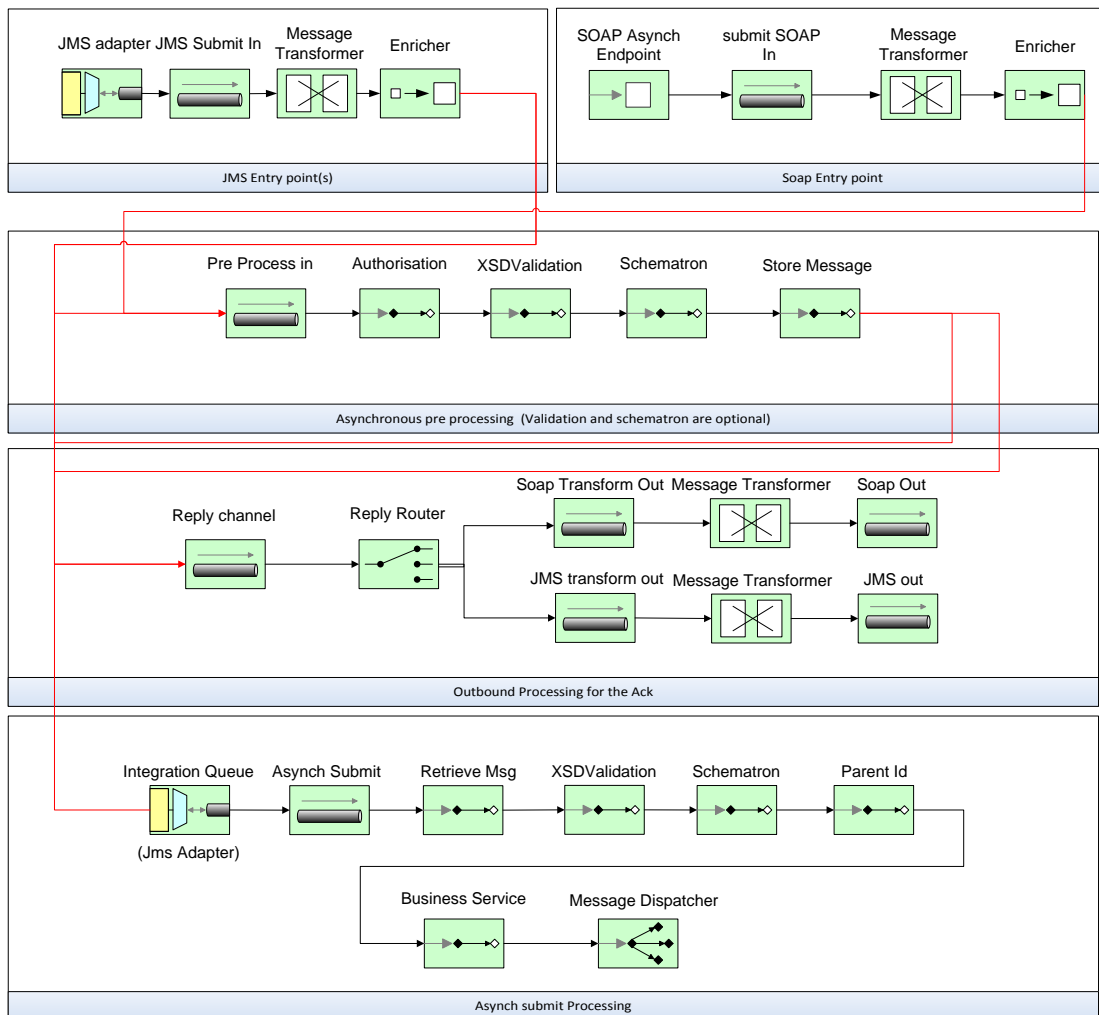
This service slightly deviates from the other synchronous services offered by the platform because the Document wrappers containing potentially large binary files are processed in two steps. The binary attachment is streamed to the file system and removed from the incoming XML message to avoid high memory consumption. The rest of the XML message is then sent to the common synchronous message handling chain of the system. Here is the EIP representation of the Store Document Wrapper message flow:



The *Store Document Wrapper* web service interface uses MTOM to optimize the size of the transferred encoded binary size and http chunking to allow the streaming of the binary file to the disk.

### 6.1.3.4. Asynchronous services

The following diagram represents the message flow for the asynchronous services provided by the system.



The platform provides asynchronous services via SOAP or JMS interfaces. The SOAP entry point for asynchronous services uses the same components than the one for synchronous services.

The JMS entry point is composed of one JMS queue and other message processing components. The message adapter that retrieves messages from the JMS queue and sends them to the JMS submit channel where the message is transformed to a platform internal message and enriched. The security configuration of the JMS resources is application server dependent and will not be detailed in this document.

Asynchronous submission of messages to the platform is a two-step process. The first step is the pre-processing of the message. During this step the platform performs the authorisation verification and some basic integrity checks on the incoming messages. This pre-processing is done synchronously for SOAP calls and asynchronously for message submitted via JMS. It is also possible to configure the platform to perform the XSD and Schematron validation of the message in the pre-processing chain if the message is coming from the SOAP entry point. At the end of this synchronous processing, the message is stored in the system database and routed to the 'outbound processing for the Ack' message handling chain and the asynchronous processing channel via the system internal integration queue.

The second step of the processing is the Asynchronous processing of the message by the platform. In order to avoid sending big XML messages in the internal JMS queue. The pre-processing chain just sends the internal id of the message that needs to be processed. The first step of the Asynchronous processing is thus the retrieval of the message. The message is then validated against the configured XSD and schematron files.

The system supports parent child relationship between messages. It is possible to configure in the platform metadata an XQuery to retrieve the parent id of a specific document. This ID is extracted from the message by the message transformer and carried through the message processing via the e-TrustEx message header. The parent id service activator is responsible of linking the two messages in the system database. The system allows configuring its behaviour in case the system cannot find the parent document. The system can be configured to wait for the parent a given time period, to simply ignore the absence of the parent or to generate an error in case of missing parent.

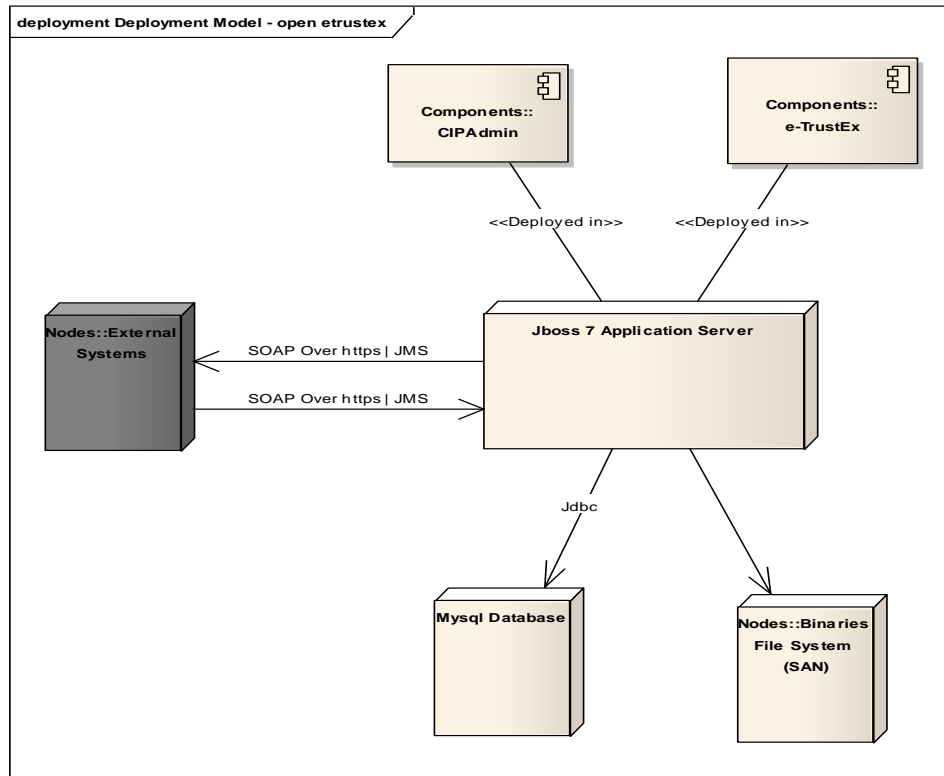
The message is then processes by the business service as described previously in this document and is dispatched to the ReceiverParty messaging endpoint by the Message dispatcher. Currently the platform support message dispatching to JMS queue or to a web service endpoint. In the case of the web-service call message dispatching the ReceiverParty system must implement the exact same interface than the platform (submit bundle for instance). Message dispatching is also configurable in the system database.

Finally, the ‘outbound processing for the Ack’ message handling chain builds and signs an XML acknowledgement message that it returns to the caller.

## 7. DEPLOYMENT VIEW

The following is a description of the hardware nodes running the Execution Environment for the system.

The following diagram, a UML deployment diagram, provides a view of hardware components involved in this project.



It is important to note that not all physical nodes are represented on this diagram. Indeed Database Servers and Application Servers could be duplicated for scalability and availability reasons. Furthermore, security mechanisms like firewalls are not shown.

The following are the identified hardware nodes.

- External System node is a system that can connect to e-TrustEx services that are HTTP SOAP web services. External System can also submit message via JMS. As we are considering large binary files, HTTP chunking and MTOM encoding are used;
- JBoss application server 7.1.1 is responsible of the document management and so the document routing, validation and persistence;
- MySQL 5.5.x is responsible of the persistence messages and system configuration.
- The binary file system potentially shared is responsible of the storage of the large binary file.

The platform has also been tested on the Oracle Weblogic application server and the Oracle 11g database.

## **8. IMPLEMENTATION VIEW**

### **8.1. Overview**

The following diagram describes the software layers of the system and their components coupled with an explanation of each element.

External systems access the platform through the web layer if they are using the web services interfaces or through the integration layer if they are using JMS. The web layer is also responsible of the authentication check. This layer also holds the CIPAdmin component. Web services endpoints are implemented using the Spring WS framework and the JAX-WS framework. The CIPAdmin is a web application using Spring MVC and JQuery.

The integration layer uses the Spring integration framework and is responsible for the Authorisation, the message validation and more generally speaking for the whole message processing. The components included in this layer are detailed in the chapter 5.1.3 of this document.

The types layer contains all the java objects generated from the XSDs and WSDLs used by the platform. These are JAX-B generated objects

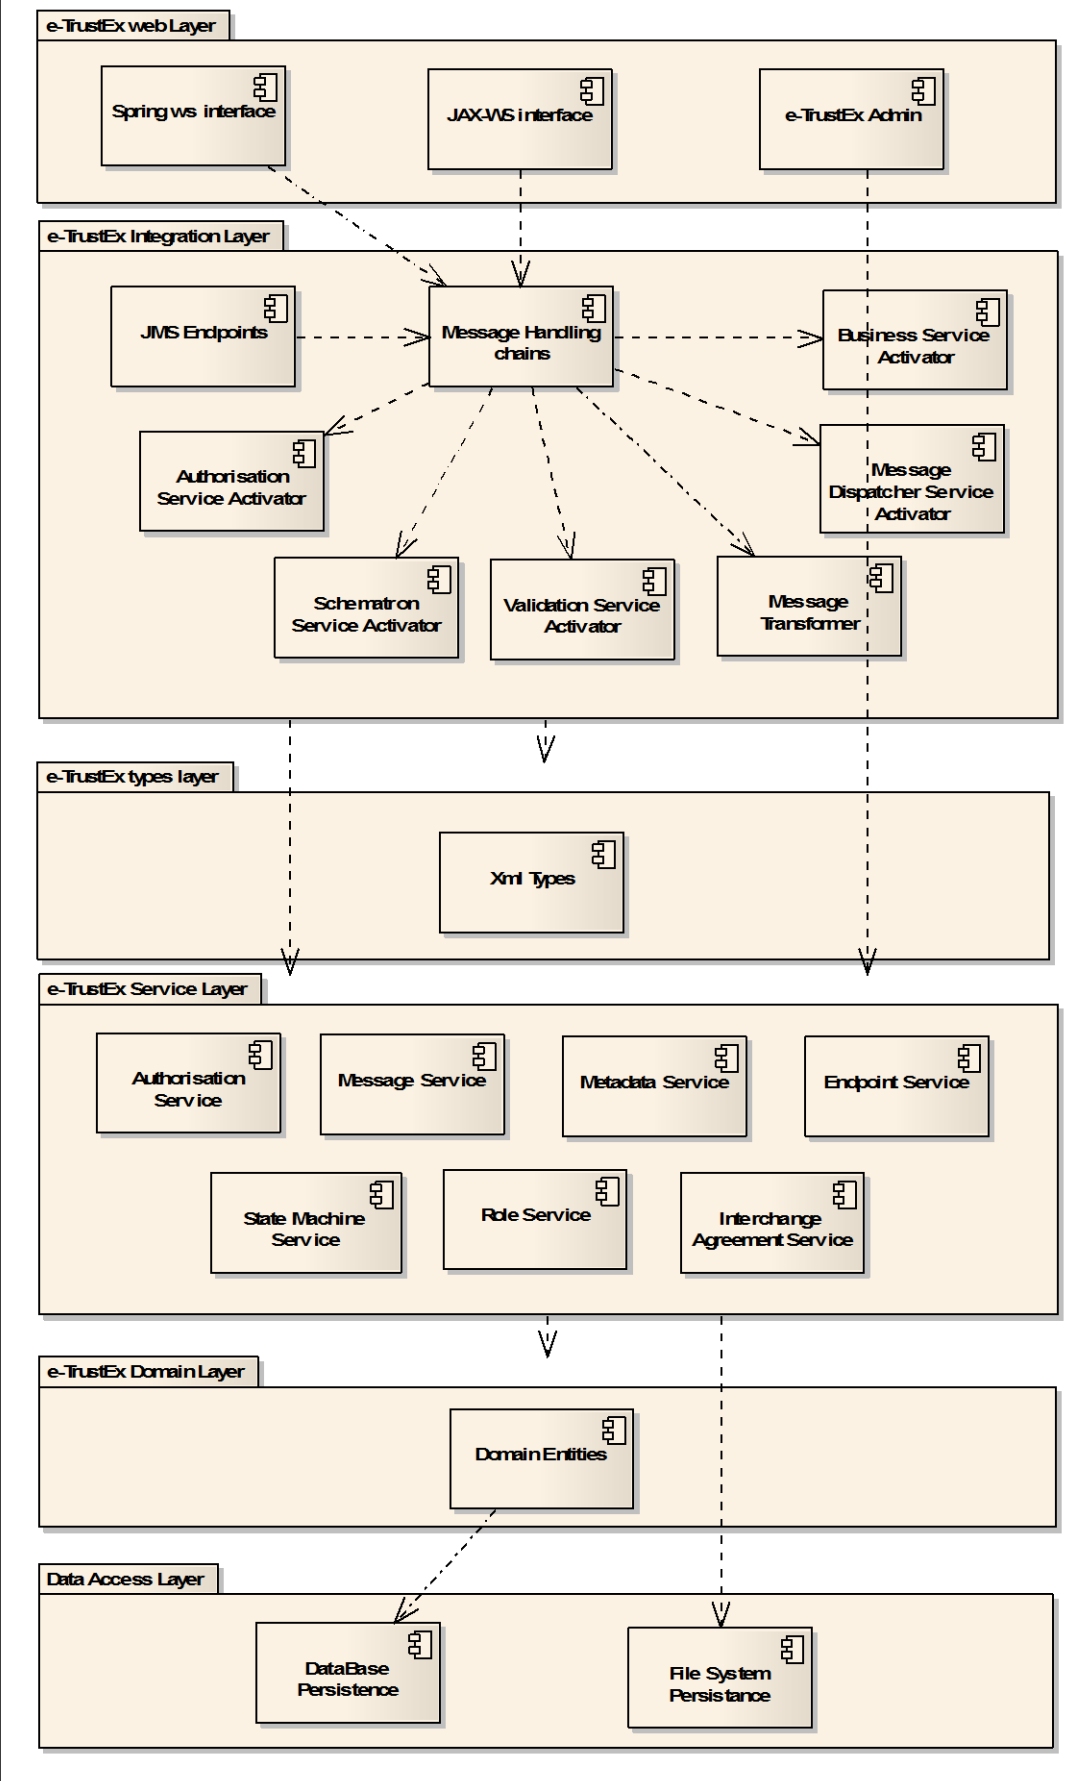
The service layers offers access to the domain objects of the platform as well as to the platform data layer. These services are Plain Old Java Objects relying the Spring framework for dependency injection and for transaction management.

The Domain layer holds all the platform entities. The persistence of these entities is implemented using the Java Persistence API version 2.0.

Finally the data persistence relies on a database and a file system to persist the data. The file system is used to store the large binaries and the database to persist the incoming messages and the system configuration.

All these layers run on a J2EE application server. The platform has been tested on Jboss 7 and Weblogic 12c.

crmp Component Model - etrust ex layers







The **Role** Entity defines roles available in the system. E.g. Customer, Supplier, BundleExchanger.

**PartyRole** One Party can fulfil different roles depending on which types of messages are being exchanged. Message exchange transactions are based on a specified sender and a receiver role. One Party can have more than one role, and these roles can be used independently.

**PartyAgreement:** Based on an agreement between parties, a party can have delegates, who are allowed to act on behalf of the delegating party. This delegate information is stored in PartyAgreement.

The **InterchangeAgreement** specifies an interchange agreement between parties. An interchange agreement has participating parties, and each of them is playing a certain role. An interchange agreement may have a start date from which it is valid from. The interchanges agreement model the document exchange contract between parties. Parties agree on a set of business transactions, a **profile** that reflect the business they are doing together.

**CIALevel** specifies a Confidentiality, Integrity and Availability levels tuple. A CIALevel can be associated to an Interchange Agreement, Profile or Transaction.

A **Message** is an element of conversation between parties. It represents the information sent from a sender party to a receiver party. A message has a status code, and associated transaction.

A **MessageBinary** either contains the raw XML message that is stored by the system, or is a placeholder which refers to a file in the file store.

**StateMachine** defines state machines for Documents, according to the SCXML specification. A StateMachine defines valid states and transitions for a Document.

## 9.2. State Machines

### 9.2.1. Introduction

Each object of the system (including the Bundles) will have its own state machine. A state machine is always initiated by a Sender. It starts with a specific message (e.g. "The Sender sends a Bundle").

Only a message will be capable of passing an object from a state to the next one. A state machine will never represent the workflow of the object in the receiver back-office. While the object is in the receiver application, the state of the object remains the same in e-TrustEx, until the back-office notifies (via an outbound message) that the workflow is complete.

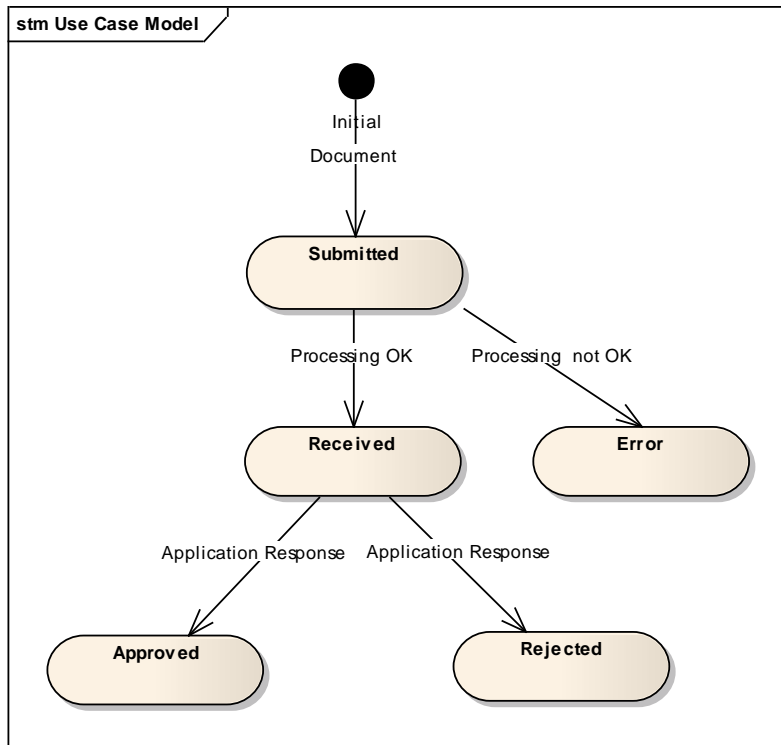
The following sections illustrate the state machines which have to be supported by e-TrustEx.

Note:

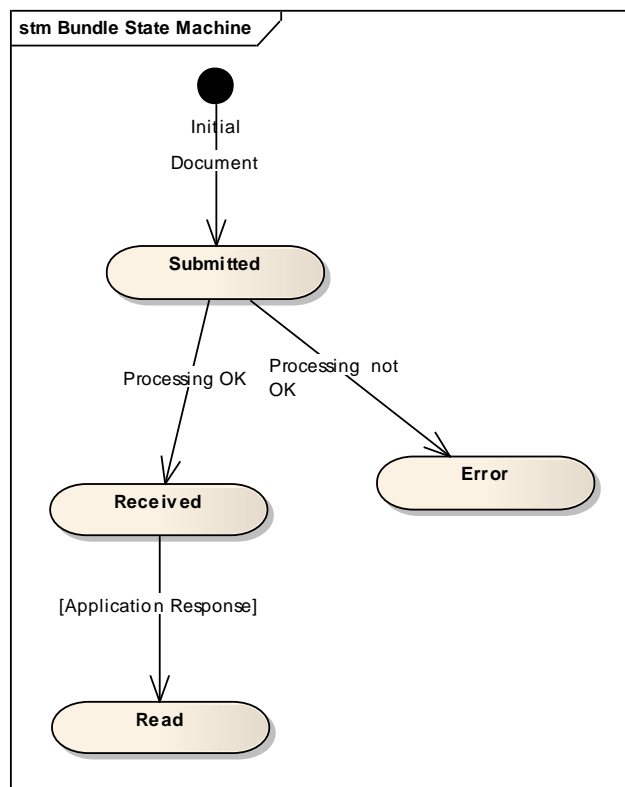
- Once a state machine instance is started, each state transition is performed by the transmission of a message. Each message is subject to logging and monitoring.

### 9.2.2. Generic State Machines

Most of the incoming objects will have the following state machine:



### 9.2.3. The document bundle state machine



## 10. SIZE AND PERFORMANCE

### 10.1. Size

Size restrictions, not on the application or its components themselves, but on the data that is being processed by e-TrustEx have an impact on the architecture and configuration of the system.

To support the exchange of large binary files the MTOM encoding has been chosen to avoid Base 64 encoding that leads to up to 30 % of size overhead.

To be able to control the incoming streams, HTTP chunking is used. The binary stream is read chunk by chunk and if the file size exceeds a given limit, the transmission is automatically ended.

That way of working also avoids loading the entire file into memory. Indeed the binary file chunks are stored in the e-TrustEx file system and kept during a given period (retention policy).

The implementers must determine the file system size, the database size and the maximum size of document wrapper according to their needs. By defaults the maximum size for large binaries is 100 MB and a Document Bundle can contain a maximum of 500 document wrappers.

### 10.2. Performance

An important architectural decision that benefits the performance of e-TrustEx includes the decoupling of the solution into a synchronous and an asynchronous mode of communication. The synchronous mode of communication mode is used by services in which the business response is immediate and which requires limited processing. The asynchronous mode however is aimed at services whose processing involves a series of long-running business actions which often require workflow steps performed by the back-office. In this mode, an incoming message will be stored in a queue where it waits to be read for further processing. If required, the number of asynchronous processing threads can be modified both in the number of total threads running and in the number of threads for a specific service (e.g. Submit Bundle).

## 11. QUALITY

The architecture of e-TrustEx contributes to quality aspects of extensibility, reliability and portability in the following ways.

### 11.1. Extensibility

As indicated in “6.1.2 Architecturally Significant Design Packages” and “□ External System node is a system that can connect to e-TrustEx services that are HTTP SOAP web services. External System can also submit message via JMS. As we are considering large binary files, HTTP chunking and MTOM encoding are used;

- JBoss application server 7.1.1 is responsible of the document management and so the document routing, validation and persistence;
- MySql 5.5.x is responsible of the persistence messages and system configuration.
- The binary file system potentially shared is responsible of the storage of the large binary file.

The platform has also been tested on the Oracle Weblogic application server and the Oracle 11g database.

Implementation View”, e-TrustEx is designed in a layered fashion and consists of multiple interconnected modules. This modular design facilitates upgrades by replacing existing modules and extensions by adding additional modules.

### **11.2. Reliability**

The reliability of e-TrustEx is enhanced through the decoupling of each architectural layer by JMS queues. A store and forward mechanism and automatic retry policy ensures that parts of the system can continue functioning without losing data when an issue occurs in a specific component.

### **11.3. Portability**

As well as being extensible, e-TrustEx is carefully designed in such a way that it is independent of the specific external system that it is serving. The use of generic HTTP SOAP web services leaves the different layers unaffected when an additional external system needs to be supported by e-TrustEx.

The usage of JPA to access the database make it easy for implementers to change the database used to store the platform data.