

Developer Guide

iPlanet Trustbase Transaction Manager

Version 2.2.1

March 2001

Copyright © 2000 Sun Microsystems, Inc. All rights reserved.

Sun, Sun Microsystems, the Sun logo, Java, iPlanet, JDK, JVM, EJB, JavaBeans, HotJava, JavaScript, Java Naming and Directory Interface, Solaris, Trustbase and JDBC are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Federal Acquisitions: Commercial Software -- Government Users Subject to Standard License Terms and Conditions

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of the product or this document may be reproduced in any form by any means without prior written authorization of the Sun Microsystems, Inc. and its licensors, if any.

THIS DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright © 2000 Sun Microsystems, Inc. Tous droits réservés.

Sun, Sun Microsystems, the Sun logo, Java, iPlanet, JDK, JVM, EJB, JavaBeans, HotJava, JavaScript, Java Naming and Directory Interface, Solaris, Trustbase et JDBC logos sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et d'autre pays.

L'utilisation de ce produit est soumise à des conditions de licence. Le produit décrit dans ce document est distribué selon des conditions de licence qui en restreignent l'utilisation, la copie, la distribution et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable écrite de Sun, et de ses bailleurs de licence, s'il y en a.

CETTE DOCUMENTATION EST FOURNIE "EN L'ÉTAT", ET TOUTES CONDITIONS EXPRESSES OU IMPLICITES, TOUTES REPRÉSENTATIONS ET TOUTES GARANTIES, Y COMPRIS TOUTE GARANTIE IMPLICITE D'APTITUDE À LA VENTE, OU À UN BUT PARTICULIER OU DE NON CONTREFAÇON SONT EXCLUES, EXCEPTÉ DANS LA MESURE OÙ DE TELLES EXCLUSIONS SERAIENT CONTRAIRES À LA LOI.

Contents

Introduction	7
Overall Layout.....	8
Related Documents	9
Solaris 8 and Java Development Kit 1.2	9
iPlanet Application Server 4.1	9
iPlanet Web Server 6.0	9
iPlanet Certificate Management System.....	9
Oracle 8i Installation and Configuration Guides.....	9
Hardware Security nCipher KeySafe 1.0 and CAFast	9
Identrus Message Specifications	9
Introduction	10
iPlanet Trustbase Transaction Manager Platform.....	10
Identrus Transaction Coordinator.....	10
Product Features	12
iPlanet Trustbase Transaction Manager Architecture	13
Overview	14
External interfaces	15
Transport protocols	15
Enterprise connectivity	16
Server to server connectivity	16
Routing	17
Authorisation	17
Services.....	18
Presentation logic	19
Overview	20
Protocol handlers.....	21
Identrus protocol handler	21
Message Readers.....	23
Default Message Reader, HTTP Reader	23
Using the default HTTP Reader	23
Default Message Reader, Identrus Reader	24
Default Message Reader, Identrus Error Reader	24
Message Writers	25
Default HTML Message Writer.....	25
Script Tags	26
Using the ScriptWriter tags	29
Default Identrus Message Writer.....	34
Default Identrus Error Writer	34
Connection Manager.....	35
Protocol Map Manager.....	37
URL Connection Implementation	39
Routing	40
Messages.....	41
Message Attributes.....	41
Identrus Message Attributes	42

Router Architecture	43
Authentication and Authorisation	45
Authentication	45
Authorisation	45
Default routing	47
Router Rules	47
Routing to service	47
Return path	47
Advanced Routing	48
Routing Rulesets	48
Router Rule Syntax	49
Complete Router Rule DTD	53
Configuration Management	54
Configuration Objects	55
Configuration Manager	55
Configuration Store	56
Configuration Services	56
Standard Services	57
Overview	58
Error and Audit Logging	59
Overview	60
Audit logs	61
Audit Logging an Event	61
Defining New Audit Types	62
Error handling and logging	64
Error Logging	64
Defining a New Error	66
Exception Handling	67
Identrus logging	69
Overview	70
Data definitions	70
Connection information	70
Raw log tables	72
Billing records	73
Building Identrus solutions	74
Methodology	75
Development process	75
Class generation	75
Service development	78
Service Building	79
Service Deployment	79
Ping Example	82
Create DTD Definitions	83
API	85
PingService Source Code	86
Creating the Identrus Service JAR	87
Deploying ping.jar within iPlanet Trustbase Transaction Manager	88
References and Glossary	91
Software Platform	91
Solaris 8 and JDK	91
Java	91
iPlanet Application Server 4.1	91
iPlanet Web Server 6.0	91
Oracle 8i	91

Hardware Security nCipher KeySafe 1.0 and CAFast	91
Transport Protocols.....	91
HTTP	91
SMTP RFC821	91
Security Related Protocols.....	91
S/MIME Version 2 Message Specification	91
DOMHASH	91
OCSP	92
Certificate requests and responses.....	92
Trading protocols	92
Identrus	92
Message Protocols	92
DOM	92
DTD	92
XML	92
XML Syntax Processing specification	92
HTML.....	92
Security related terms	93
Java Related terms	95
Server Definitions.....	98
Index	102

Table of Figures

Figure 1 Product features	12
Figure 2 Three Tier Architecture	14
Figure 3 iPlanet Trustbase Transaction Manager Interfaces	15
Figure 4 Presentation layer components	20
Figure 5 XML Repeat Iterators	27
Figure 6 Connection Manager Architecture	35
Figure 7 Router Architecture	43
Figure 8 Default routing rules	47
Figure 9 Rule Sets	48
Figure 10 DTD Rule Body	50
Figure 11 Component replication	58
Figure 12 iPlanet Trustbase Transaction Manager log manager	60
Figure 13 Error Severity Types	64
Figure 14 Error_codes Table	65
Figure 15 Error Table	65
Figure 16 Error Support Table	66
Figure 17 iPlanet™ Trustbase Transaction Manager Exception Hierarchy	67
Figure 18 Exceptions	68
Figure 19 SSL Connection Table	70
Figure 20 SMIME Transport Table	70
Figure 21 SMTP Connection Table	71
Figure 22 SMTP Message Table	71
Figure 23 OCSP table	71
Figure 24 Raw log Table	72
Figure 25 Certdata Table	72
Figure 26 Bill data Table	73
Figure 27 Development process	75
Figure 28 Message path processing	78
Figure 29 Deploying PingService within iPlanet Trustbase Transaction Manager	80
Figure 30 Assigning a role to a service	81
Figure 31 Classgen Options	84
Figure 32 Deploying PingService within iPlanet Trustbase Transaction Manager	88
Figure 33 Assigning a role to a service	90

Chapter 1

Introduction

This document forms one of the iPlanet™ Trustbase Transaction Manager framework documentation set. This document is aimed at designers and developers looking to produce applications that utilise the iPlanet™ Trustbase Transaction Manager framework. The guide is divided into three main sections:

- Introduction to the environment iPlanet™ Trustbase Transaction Manager has been designed for
- An overview of the Framework and how the components interact
- Development of a iPlanet™ Trustbase Transaction Manager application

Overall Layout

The complete documentation set comprises of:

- iTTM2.2-Utility-Guide.pdf that provides some tools for helping with PKI Certificate Management.
- iTTM2.2-Install-Configuration-Guide.pdf is designed for operators looking to produce applications that utilise the iPlanet Trustbase Transaction Manager framework. It is designed to provide information for operators looking to install the iPlanet Trustbase Transaction Manager platform. This guide identifies hardware and software required prior to installation, how to install iPlanet Trustbase Transaction Manager from CD-ROM.
- iTTM2.2-Developer-Guide.pdf (this Document) that indicates how to build and deploy your own services.
- API reference <install_dir>/Trustbase/TTM/V2.2/apidocs is a softcopy Java documentation set that is provided as part of the iPlanet™ Trustbase Transaction Manager installation. It is designed to provide application developers with the information to utilise the framework and tools provided within the iPlanet™ Trustbase Transaction Manager framework.

This manual assumes the reader is familiar with Java standards described in <http://www.javasoft.com> and XML described in <http://www.w3.org/TR/REC-xml>.

The manual also assumes that the reader has attended the iPlanet™ Trustbase Transaction Manager Developer Training course.

Related Documents

Solaris 8 and Java Development Kit 1.2

<http://docs.sun.com>

<http://java.sun.com/products/jdk/1.1/docs/index.html>

iPlanet Application Server 4.1

<http://docs.iplanet.com/docs/manuals/ias.html>

iPlanet Web Server 6.0

<http://docs.iplanet.com/docs/manuals/enterprise.html>

iPlanet Certificate Management System

<http://docs.iplanet.com/docs/manuals/cms.html>

Oracle 8i Installation and Configuration Guides

<http://www.oracle.com>

Hardware Security nCipher KeySafe 1.0 and CAFast

<http://www.ncipher.com>

Identrus Message Specifications

<http://www.identrus.com>

Transaction Coordinator requirements (IT-TCFUNC)

Core messaging specification (IT-TCMPD)

Certificate Status Check Messaging specification (IT-TCCSC)

Introduction

The iPlanet™ Trustbase Transaction Manager platform provides a message oriented middleware platform capable of supporting a variety of banking and trade facilitation applications. This platform is specifically designed to enable Financial Institutions to make use of the Identrus Network by providing all of the function required to process messages that conform to the Identrus messaging specification.

iPlanet Trustbase Transaction Manager Platform

The iPlanet™ Trustbase Transaction Manager Platform operates at a high level as follows:

- Incoming messages are received from dedicated security services that support secure channel communication - including HTTP over SSLv3 and SMIMEv2 over SMTP.
- Requests are passed through a message parsing and verification engine that verifies that the incoming message is requesting a service that is offered by the platform, is complete and correctly formatted and builds the service request message request for processing through the system. The platform supports a range of standards for message coding including XML and HTML. A naming service identifies the application logic that the message object will be routed to for processing.
- Application Logic. The platform can support a range of application logic and can be extended by both the deploying and third party organisations.
- Third party services. Applications executed on the iPlanet™ Trustbase Transaction Manager platform may apply to third party systems during the decision making process. These include an organisation's operational systems and third party services.
- Response Management. A response to the relying party is constructed and returned to the requesting party based on the policy defined for each particular transaction supported by the platform.

All stages of the transaction processing process can be recorded in the repository's audit facilities.

Identrus Transaction Coordinator

The iPlanet™ Trustbase Transaction Manager provides an implementation of the Identrus Transaction Coordinator as specified in the following Identrus documents:

Title	Description and Document Reference
Transaction Coordinator	Transaction Coordinator requirements (IT-TCFUNC)
Transaction Coordinator Messaging Protocol Definition	Core messaging specification (IT-TCMPD)
Transaction Coordinator Certificate Status Check Protocol Definition	Certificate Status Check Messaging specification (IT-TCCSC)

The iPlanet™ Trustbase Transaction Manager is designed to be an extensible platform that not only performs the core CSC function of an Identrus Transaction Coordinator, but also allows the developer to produce applications that conform to the Identrus messaging specification.

The iPlanet™ Trustbase Transaction Manager makes producing Identrus compliant applications simple for the Developer by providing the following functionality:

- Standard presentation of messages independent of the transport mechanism.
- Standard Authentication and Authorisation of requests.
- Common validation and error handling of all Identrus Network layer information within messages.
- Tools for generation of new application message types that conform to the Identrus specifications.

The early sections of this guide identify these common components and the functions supported by each. The later sections of the guide show how an Identrus application may be produced and deployed on the iPlanet™ Trustbase Transaction Manager infrastructure.

Product Features

iPlanet™ Trustbase Transaction Manager provides a platform for the delivery of business to business E-Commerce solutions. The principal challenges to be addressed are as follows:

Figure 1 Product features

Identity	<p>The solution recognises different types of credential and uses them as the basis for authentication requests, sometimes to remote and/or third party services.</p> <p>It supports the establishment of multi-party trust relationships, including those based on iPlanet™ Trustbase Transaction Manager Third Parties and other external CA services.</p>
Entitlement	<p>Once the identity of trading partners has been established, the platform provides a means for a business to determine which services it should provide to a potential customer.</p> <p>The e-commerce platform specifies and enforces terms for the provision of services based on the user type, credential type, access channel, credit worthiness, and so on.</p>
Dispute Resolution	<p>Using certificate-based technology, the platform provides services for transaction signing, audit logging and non-repudiation.</p>
Integration	<p>iPlanet™ Trustbase Transaction Manager integrates with existing back and mid-office solutions such as Customer Relationship Management and Cash Management solutions.</p> <p>It supports a flexible e-commerce model, catering for '2-to-n'-party transactions.</p>
Availability	<p>It provides Business to business commerce on a global basis that allows 24x7x365 availability.</p>
Scalability	<p>It services an increasing customer base catering for unpredictable performance requirements.</p>
Security	<p>E-commerce solutions must provide total security for all parties. This includes identification and authentication services, data confidentiality and integrity, transactional security, and security management services.</p>

Chapter 2

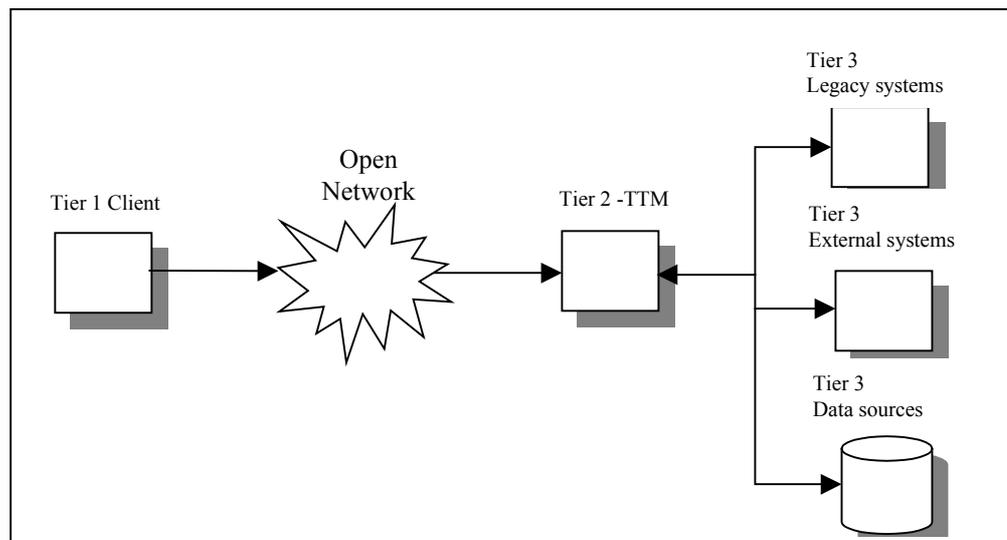
iPlanet Trustbase Transaction Manager Architecture

The iPlanet™ Trustbase Transaction Manager is designed to fulfil the need for Identrus enabled message oriented middleware for Financial Institutions. The platform provides a means of offering Financial Services Applications over the Internet that are consistent and re-useable.

Overview

Within a multi-tiered architecture, the iPlanet™ Trustbase Transaction Manager component is the middle tier infrastructure that provides the ability to offer new function, but shields first tier clients from the complexity of the existing enterprise infrastructure (see below)

Figure 2 Three Tier Architecture



In order to achieve this, the iPlanet™ Trustbase Transaction Manager is designed to be:

- Highly available
- Secure
- Reliable and scalable

It achieves this by using server side standards such as Servlets and EJB and leverages the existing reliability and scalability functions of the iPlanet Application Server.

The iPlanet™ Trustbase Transaction Manager Platform extends the iPlanet application server function by providing a message-handling pipeline that may be extended to process specific message types and formats. This message-handling pipeline contains four major components:

- Transport listeners
- Presentation and formatting handlers
- Routing and authorisation management
- Business logic plug ins

The message pipeline is populated with a set of components that provide all of the necessary pre-processing to support applications using the Identrus messaging protocol and HTML based communication.

The operational side of the iPlanet™ Trustbase Transaction Manager platform is augmented by a toolkit that allows the developer to generate message classes and deploy the completed application. This is designed to reduce the development effort required to produce applications that conform to the Identrus requirements to an absolute minimum.

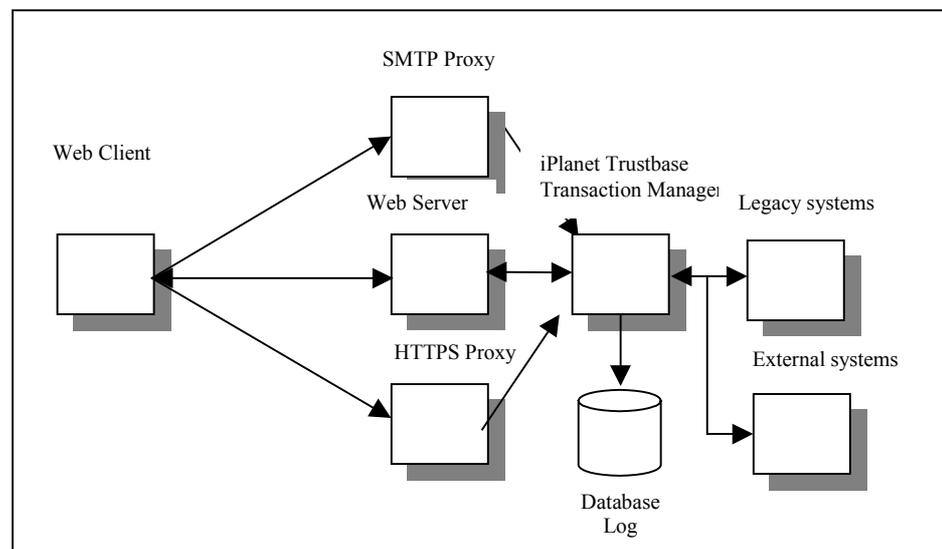
External interfaces

The iPlanet™ Trustbase Transaction Manager is a middleware component that provides a means of accessing and using existing legacy data sources over the Internet. In order to provide this function the platform must be capable of:

- Listening and processing web protocols - HTTP, HTTPS, SMIME
- Accessing existing legacy systems via JDBC, CORBA, RMI
- Using resources provided by other servers on the Web

The figure below shows how these external interfaces relate to the iPlanet™ Trustbase Transaction Manager component.

Figure 3 iPlanet Trustbase Transaction Manager Interfaces



Transport protocols

The iPlanet™ Trustbase Transaction Manager supports three basic transport protocols:

- SMTP - Asynchronous mailed based communication
- HTTP - Synchronous insecure communication
- SSL - Secure synchronous communication

iPlanet™ Trustbase Transaction Manager listeners proxy the SMTP and SSL transport protocols to provide both a means of processing, and a means of logging transport specific data. Any HTTP data is listened for directly by the Web Server.

These three transport protocols provide a means of carrying a variety of application level messaging. The iPlanet™ Trustbase Transaction Manager contains presentation components that deal directly with:

- SMIME wrapped XML or HTML
- HTML
- Identrus compliant XML application messages

The iPlanet™ Trustbase Transaction Manager platform may be extended to support other application messaging protocols as required.

Enterprise connectivity

The iPlanet™ Trustbase Transaction Manager business logic is designed so that the business logic implemented by the developer may use all of the standard connectivity components available within the J2EE platform. Standard J2EE connectivity components include:

- JDBC - Access to relational databases
- RMI and CORBA - Access to remote objects and EJB's
- JNDI - Access to directory and naming services
- JMS - Message oriented interfaces for use with message queues
e.g. MQ Series

The iPlanet™ Trustbase Transaction Manager business logic is also capable of utilising the underlying connectivity components provided by the iPlanet application server for access to internal systems. These enterprise connectors include interfaces for:

- R/3 - Enterprise Resource Planning
- CICS - IBM Mainframe integration
- BEA Tuxedo - Transactional data systems
- Peoplesoft - Enterprise Resource Planning

Server to server connectivity

The iPlanet™ Trustbase Transaction Manager platform provides a means of abstracting the transport and presentation formats used by the incoming message within the presentation component. The same mechanism (through a different API) is used to allow a business service to make requests of other servers.

The Identrus Certificate Status Check (CSC) service uses this connector functionality to provide a means of determining the validity of a certificate at a particular point in time (See Section on Presentation logic on page 19)

Routing

The router provides a mechanism for imposing structure and ordering on the execution of services in a secure way. It acts as a gatekeeper to ensure that services are only executed by authorised individuals and in an appropriate context. A user of the system will connect to the server and then exchange messages. At the highest level, the user will be trying to accomplish a task. Some tasks will require authorisation (and therefore authentication) prior to being performed; services may also perform tasks in a slightly different fashion depending on the identity of the user making the request.

The Router has been designed with the following in mind:

- Authentication and authorisation is kept separate from business logic.
- Configuration and management of the routing table is easily implement-able and not error prone.
- Complex solutions can be built where required
- Implementing a simple solution is not difficult
- Services can implement atomic business level functions and are independent of one another.

The function of the router is central to the iPlanet™ Trustbase Transaction Manager platform. All messages are passed through a router, and, based on the current context of the message and its contents, the router will accept or reject the message for processing.

In order to define a flexible mechanism for routing, capable of working within a variety of complex environments, iPlanet™ Trustbase Transaction Manager provides rule based routing. This allows a means of modifying, and extending, the behaviour of the iPlanet™ Trustbase Transaction Manager installation over a period of time without the need to modify existing modules or services. See also the section on Routing on page 40 for more details on this.

Authorisation

The basic requirement of being able to gate service access is met by the ability to route a message based not only on the message type, but also on its current level of authorisation. Within iPlanet™ Trustbase Transaction Manager, authorisation is considered an extension of authentication i.e. in understanding who a person is, we can determine what they are allowed to do.

The authentication mechanisms of iPlanet™ Trustbase Transaction Manager are not a separate component. Authentication data is gathered by the default iPlanet™ Trustbase Transaction Manager framework. This can then be added to by domain specific services. The platform provides a default authorisation service that is capable of mapping both a username and password, or a digital certificate onto a user group or *role*. The router then ensures that when a service is accessed, the role has been authorised for access to that service.

Developers are at liberty to replace the default authorisation service with a mechanism that maps user information onto existing repositories such as an enterprise directory service. See also the section on Routing on page 40 for more information on this.

Services

Business services are at the heart of an e-commerce application, and the iPlanet™ Trustbase Transaction Manager provides a means of registering services written by the developer into the platform. These services need not be concerned with processing transport specific information, presentation specific information, authentication of the user, or authorisation of a users request. This allows the developer to concentrate on the function of the application, and integration of existing systems into a web enabled infrastructure. See also the section on Standard Services on page 57 for more information on this.

Chapter 3

Presentation logic

When an iPlanet™ Trustbase Transaction Manager implementation, on receiving a message, routes the message through to its appropriate service, and the service cannot directly act upon this message, iPlanet Transaction Manager utilises the Connection Manager component to process messages appropriately.

Overview

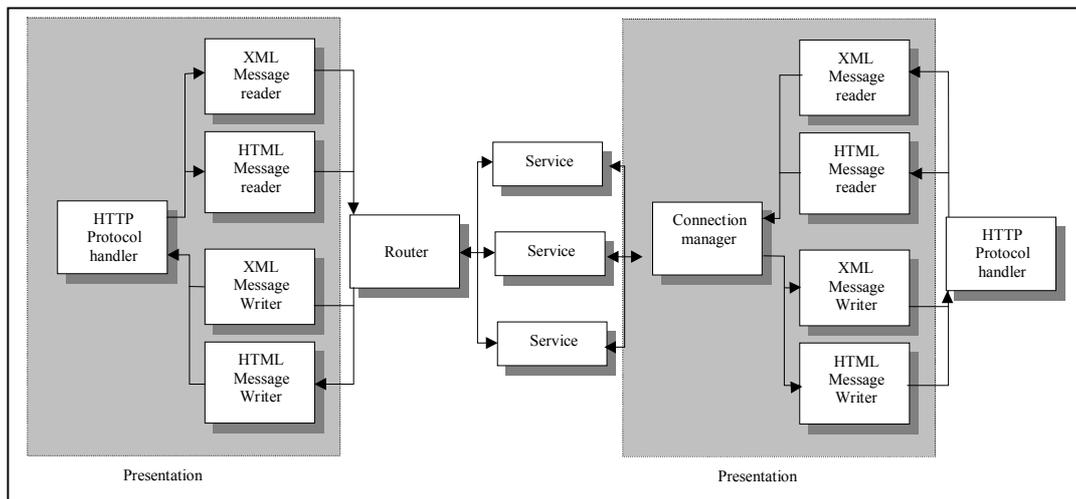
The iPlanet™ Trustbase Transaction Manager presentation components comprise of a set of protocol handlers, message readers and message writers. The purpose of the protocol handlers is to extract the transport specific headers and footers, determine the MIME type and message type, and to forward the message content to an appropriate message handler.

The message handlers are selected on the basis of the MIME and Message type determined by the protocol handlers. The iPlanet™ Trustbase Transaction Manager provides two default message handlers, these are:

- Identrus XML messages
- HTML messages

In the case of the Identrus XML message, the appropriate Identrus network processing is performed as described in the following sections.

Figure 4 Presentation layer components



In the case of HTML requests, the message reader converts all of the fields in the HTML form into an XML structure represented by a set of DOM nodes.

In both the cases of XML and HTML the output to the router is wrapped with the session information into an iPlanet™ Trustbase Transaction Manager message. These iPlanet™ Trustbase Transaction Manager messages are then sent, by the router, to the appropriate service.

The services perform the appropriate message transformation or construction and return an iPlanet™ Trustbase Transaction Manager message back to the router. The router will then send the iPlanet™ Trustbase Transaction Manager message back to an appropriate message writer. In the case of the HTML message writer this will extract the DOM values in the message into an HTML template determined by the message type. The template may contain a set of directives that allow conditional or recursive construction of form data.

Protocol handlers

When the platform receives a message, the protocol handler is selected on the basis of the MIME type of the incoming message and is then invoked. The function of this component is to determine the application protocol information, specifically:

- Message type
- Context ID

Each MIME type, and therefore protocol handler, represents a class of messages e.g. application/OCSP. The client must be capable of generating the appropriate MIME type. If the client is a Web browser it normally generates a MIME type of application/x-www-form-urlencoded. iPlanet™ Trustbase Transaction Manager comes supplied with the default Protocol Handler for use in this situation. Its default actions are to:

- Search the form fields for:
 - Message Type – Identifying the message type
 - Context ID – Identifying the Context Identifier
- To set the response MIME type to `text/html` indicating to the Browser that the reply is a standard HTML page.

Identrus protocol handler

The protocol handler does the initial message processing for all messages arriving at the TC with a mime-type that begins with the following string: “application/identrus-”.

Once at the protocol handler the following actions are carried out in the following order:

- The raw XML stream is read from the client.
- The raw XML is parsed into the internal iPlanet object structure (known as the message tree)
- The message tree is then validated against the DTD for that message that is specified by the DOCTYPE in the XML message.
- The message tree is traversed and any ID attributes are checked to ensure that there are no duplicates. This is a mandatory XML processing step, and is especially important where XML structures will be digitally signed.
- The system identifier from the DOCTYPE tag is placed into a iPlanet™ Trustbase Transaction Manager Message Attribute identified by [IdentrusConstants.SYSTEM_ID_ATTR](#).
- Elements from the NIB are extracted ready for Raw Message Logging.
- The Certificate Bundle is extracted from the message and all of the certificates are logged in the Certificate Log.

- The message, minus its Certificate Bundle, is sent to the Raw Message Log.
- The Raw Log Id is returned from the log and placed into a iPlanet™ Trustbase Transaction Manager Message Attribute identified by [IdentrusConstants.RAW_RECORD_MARKER](#)
- The iPlanet™ Trustbase Transaction Manager Message type is set to be [IdentrusConstants.IDENTRUS_MESSAGE](#)
- The message tree is attached to the iPlanet™ Trustbase Transaction Manager Message as serialised content, ready to be sent to the appropriate message reader.

Note: The variable DOCTYPE is an XML concept. See for instance <http://www.w3.org/TR/REC-xml>. The constants SYSTEM_ID_ATTR, RAW_RECORD_MARKER and IDENTRUS_MESSAGE are Identrus Constants see [com.iplanet.trustbase.identrus.IdentrusConstants](#) for more information on Identrus constants.

Message Readers

The appropriate message reader is invoked based on the information gathered by the protocol handler. Each Message Reader will have an implicit understanding of the class of messages it can convert into an internal iPlanet™ Trustbase Transaction Manager message.

Default Message Reader, HTTP Reader

iPlanet™ Trustbase Transaction Manager comes with a pre-registered HTTP/POST Message Reader that takes an `application/x-www-form-urlencoded` type. This is registered as a default message reader, meaning that if the message reader registry is unable to make a specific match to the message type detected by the host environment adapter then the default HTTP/POST message reader will be returned. If the MIME type is anything other than `application/x-www-form-urlencoded` an error is returned.

Once the appropriate message reader has been found it is then passed an iPlanet™ Trustbase Transaction Manager Message object and the input stream containing the remaining unprocessed message content. The message reader reads and parses this message content from the input stream and creates a data block containing an internal representation of the message content. This is then stored in the iPlanet™ Trustbase Transaction Manager Message object for use by the appropriate service(s).

Once the message reader has parsed the input stream, and completed construction of the iPlanet™ Trustbase Transaction Manager Message, the message analyser passes the iPlanet™ Trustbase Transaction Manager Message on to the Router, and handles any errors or exceptions that result from the subsequent processing of the iPlanet™ Trustbase Transaction Manager Message.

Using the default HTTP Reader

The HTTP reader performs a relatively simple extraction of information from the input stream. Each input field is extracted from the HTTP Post or Get string and placed into a DOM structure. Once all the fields have been extracted this DOM Object is stored within the iPlanet™ Trustbase Transaction Manager Message. If the user does not specify the structure of the DOM object then each field is stored as an attribute of a node `HTTP_VALUES`,

```
HTTP
HTTP_VALUES Attributes [<field_name> = <value>]*
```

However it is possible for the developer to force the fields to be stored in a more complex structure using certain escape sequences.

In order to create your own DOM structure the HTTPReader class has the ability to recognise an escape sequence contained in place of an input field name: `<input name="$|xml:<node1>...<noden>|$">`. For example:

```
<input name="$|xml:message.user.name|$">
<input name="$|xml:message.user.address|$">
<input name="$|xml:message.user?jobtitle|$">
```

This will lead to a set of nodes being created under the root HTTP node containing the following data:

```
HTTP
HTTP_VALUES Attributes [<field_name> = <value>]*
message
user Attribute jobtitle=<value>
name = <value>
address = <value>
```

Default Message Reader, Identrus Reader

This message reader will handle messages with a mime-type that begins with "application/identrus-" and a message type of [IdentrusConstants.IDENTRUS_MESSAGE](#).

Once at the message reader the following actions are carried out:

- Signature check the mandatory DSIG signature
- Verify the certificate chain in the message, and ensure that the root certificate of the chain is in the database. Ensure that the root of the chain has a Certificate Store attribute of [IdentrusConstants.IR_CA](#).
- Set the security context on the iPlanet™ Trustbase Transaction Manager Message ready for the authorisation phase in the router.
- Make a billing log entry for the message

The message is then sent for routing.

Default Message Reader, Identrus Error Reader

This message reader will handle messages with a mime-type of "application/identrus-identrustransporterror" and a message type of [IdentrusConstants.IDENTRUS_TRANSPORT_ERROR](#)

Once at the message reader the following actions are carried out:

- There is no signature to check on an incoming "Identrus Transport Error", so the security context on the iPlanet™ Trustbase Transaction Manager Message is left unset. Messages arriving through this reader should only be coming through the connector, so they should never go to the router.

Message Writers

Once a message has been routed through to one or more services and returned to the Router, it is passed back to the Message Analyser (this assumes no exceptions or errors are thrown). At this point the message analyser is required to determine which Message Writer to be used to process the returned iPlanet™ Trustbase Transaction Manager Message into a suitable form for return to the user. As with message readers the analyser uses the Messages returned type and response MIME type values to select the appropriate message writer.

The message writer is responsible for extracting information contained in the iPlanet™ Trustbase Transaction Manager Message and processing it into a format that can be written back to the client.

To register the message writer with the registry iPlanet™ Trustbase Transaction Manager is supplied with an HTML management interface. This allows the developer to register new Message Writer sub-classes in the internal configuration store associated with iPlanet™ Trustbase Transaction Manager. Alternatively registration may be undertaken by adding the appropriate entries to the [tbase.properties](#) file. To achieve this, new Message Writers should be added to the [MessageWriter] section and take the form:

```
message.writer=<name>:<classpath>
message.writer=
Script:uk.co.jcp.tbaseimpl.parse.message.http.ScriptWriter
```

where the <name> is a name given by the developer and the <classpath> is the fully qualified classpath of the new [MessageWriter](#) instance. It is important to note that if you add new Message Writers through the management interface and at a later date it becomes necessary to restart iPlanet™ Trustbase Transaction Manager using the properties file then all changes will be lost. So it is as well to add entries to the properties file even if you are using the management interface.

Default HTML Message Writer

iPlanet™ Trustbase Transaction Manager comes with a pre-registered HTML message writer that takes a text/html MIME type. This is registered as a default message writer, meaning that if the message writer registry is unable to make a specific match to the message type and MIME type detected by the host environment adapter then the default HTML message reader will be returned. If the MIME type is anything other than text/HTML an error is returned.

This [MessageWriter](#) is capable of translating XML fields contained in the message into an HTML form using pattern matching and templates similar to many other Dynamic HTML writers.

The ScriptWriter, the default HTML Message Writer, uses a scripting language that is intended to provide a simple way to display uncomplicated information quickly and efficiently without having to recourse to a full XSL.

This is intended for use where the response is expected to contain text (usually HTML), and the information returned by the router is simple enough to allow straightforward substitution of values from the returned XML into the Template.

The ScriptWriter class works by taking the DOM object contained in the returned message and extracting all the fields therein. These fields are then substituted into the template that is selected by using the type and format of the returned message type. The substitution is achieved using a series of Script Tags that allow the developer to specify which parts of the returned message are to be placed where in the template. The templates used by the ScriptWriter may be configured in one of two ways, either through the iPlanet™

Trustbase Transaction Manager HTML management interface or by adding the appropriate entries to the [tbase.properties](#) file in the [ScriptWriter] section as illustrated below:

```
writer.typeandformat=<format>:<type>:<template>
writer.typeandformat=text/html:TimeServiceTimeResponse:Config\Templates\
TimeService\TimeServiceTimeResponse.html
```

Where the <format> is text/html, <type> is the message type that will have been set by the service that last processed the message before it was returned, and <template> is the name of the template file. This file may either be an absolute path or a relative path, if it is a relative path then an entry is required to specify the location of the 'root':template.directory=.\
(note the trailing '\ ' which is required)

Script Tags

There are four main tags used in the script templates:

- \$\$xml:<XML Variable Name>\$\$

This represents the single substitution of a variable from the XML Document into the template. <XML Variable Name> is the fully qualified 'location' of a node in the XML Document returned by the router which is expected to contain a string value that may be substituted into the template. For instance, to retrieve the name from an XML Document one might specify the tag,

```
$$xml:message.content.user.name$$
```

This tag is also capable of specifying an attribute from a given XML node using the "?" escape sequence to select the given attribute. For example, in order to retrieve users middle initials that are stored as an attribute of their name, one might specify the following tag,

```
$$xml:message.content.user.name?middle_initials$$
```

- `$$repeat:<XML Array Name>[<UserDefined Tag Name>:<Iterations>]$$...$$/repeat$$`

Represents the start point for a repeating array of values contained in XML Document.

`<XML Array Name>` is the fully qualified 'location' of an array in the XML Document.

`<UserDefined Tag>` Name allows the template author to provide a shorthand name for this array location, e.g. If the array represents the address of a user, `Response.Content.User.Address`, then the author could tag this as `address` for ease of use within the repetition block.

`<Iterations>` tag allows the author to specify which 'parts' of the array are to be displayed.

Figure 5 XML Repeat Iterators

Iterator	Description	Example
*	Iterate over all items contained in the specified array	*
<n>	Select the n th element of the array, if the array has less than n elements then nothing is displayed	2
<m>-<n>	Select from the m th to the n th elements (inclusive), if the array has less than n elements then this selects from m to the end, if the array has less than m elements then nothing is displayed	3-5
<m>-*	Select all elements from the m th up to the last element in the array	2-*
<attribute>=<value>	Select ALL the elements in the array where the specified attribute equals the given value	location= London
,	Allow multiple iterators to be specified together	0,3- 5,9,postcode =EC1

Every repeat block is terminated with a `$$/repeat$$` statement.

- `$$xif:<XML Variable Name>[="<value>"]$$... $$/xif$$`

Represents a conditional tag. In its simplest form this tag says iterate everything contained within the conditional statement if the node (or attribute - using the same "?" escape sequence) is found to exist within the XML.

If the "=" is used however, not only does the node or attribute have to exist but the value has to be an exact (case sensitive) match of the one specified. For example one might specify a conditional block say only print the user details if the job title matches "SysAdmin" as follows:

```
$$xif:message.content.user?jobtitle="SysAdmin"$$
```

Associated with the `$$xif:...$$` tag there is also an `$$xelse:...$$` tag to provide optional switching. This tag has two modes of use:

```
$$xelse$$
```

Simple usage, if the `$$xif:...$$` condition is unfulfilled then this condition is used instead. One might use this sequence if it is necessary to use two different display formats based on the contents of an XML node.

```
$$xelse:<XML Variable Name>[="<value>"]$$
```

Uses the same format as the xif statement to determine if this branch should be executed. This provides extra conditional switching which can be used to make more complex decisions about which sections of the template are to be used.

- `$$xprop:[<section>]<property>[=<default_value>]$$`

This is a special tag that is used to incorporate property's values from the [tbase.properties](#) file into the template. This allows different application server environments to be supported by the same template by changing the property value for each server type. The `$$xprop:...$$` tag is also special in that it is pre-processed before any of the other tags are dealt with, thus it is possible to include other tag information (except `$$xprop:...$$` itself) in the properties.

- `<section>` the name of the section that contains the property, if this is left out, the section used is the ScriptWriter section.
- `<property>` the name of the property that is to be used.
- `<default_value>` a default value to be used if the property is not found, if this is not specified and the property is not found then an empty string is returned.

Using the ScriptWriter tags

- `$$xml:...$$` substitution

As documented above this tag represents a straightforward substitution of a single value contained in the returned XML Document, into the HTML Template. In the following template snippet, we can see that the author expects the XML Document returned from the router to contain data indicating an administrative support mail address.

```
<P>
<FONT SIZE=-2>
Please try our
<A HREF="mailto:$$xml:response.content.supportaddress$$">Support
Desk</A>
if you have any enquiries
</FONT>
</P>
```

- `$$xml:...$$` filtering

It is also possible to introduce an attribute filter into the qualified path, as shown in the code snippet below:

```
<P>
Email the
<A
HREF="mailto:$$xml:response.content.user?title="SysAdmin".email$$">System
Administrator</A>
if you have any enquiries regarding your login details.
</P>
```

Here the XML Document contains an array of one or more user's details and each user contains a node that contains their email address. We use the user title attribute to determine which of the users job description is System Administrator and use the resultant user to extract the email address.

- `$$repeat:...$$` loop

The following example shows the complete usage of a repeat block, including the `$$/repeat$$`. Here the author is indicating the XML Document response will contain an array of transactions, under the node `http_response.content.transaction`. This section of the XML Document tree is then tagged as `txn`.

Inside the repeat block, one then uses the standard `$$xml:...$$` notation to select data contained in each member of the array returned.

Each transaction will contain three nodes `date`, `description` and `amount`. These will be filled into a single table row and the table will contain as many rows as there are transactions returned.

```
<table BORDER="1" BGCOLOR="#fff0a0">
  <th COLSPAN="3" BGCOLOR="ffd080"><b><i><font SIZE="+1"
COLOR="#00a000">Statement Details</font></i></b></th>
  <tr>
    <td WIDTH="50" ALIGN="center"><B>Date</B></td>
    <td WIDTH="400" ><B>Description</B></td>
    <td WIDTH="200" ALIGN="center"><B>Amount</B></td>
  </tr>
  $$repeat:message.content.transaction[txn:*]$$
    <tr>
      <td>$$xml:[txn].date$$</td>
      <td>$$xml:[txn].description$$</td>
      <td ALIGN="right">$$xml:[txn].amount$$</td>
    </tr>
  $$/repeat$$
</table>
```

- Nesting `$$repeat:...$$` blocks

It is also possible to nest repeat blocks, if we consider the users example again it might be that each user node contains an array of address nodes (one per line of their address) and hence we might wish to iterate over the address lines for each user. The following template snippet shows how this might be done:

```
<table BORDER="1" BGCOLOR="#fff0a0">
  <th COLSPAN="3" BGCOLOR="ffd080"><b><i><font SIZE="+1"
COLOR="#00a000">System Users</font></i></b></th>
  <tr>
    <td WIDTH="100" ALIGN="center"><B>Name</B></td>
    <td WIDTH="400" ALIGN="center"><B>Address</B></td>
    <td WIDTH="100" ALIGN="center"><B>Telephone</B></td>
  </tr>
  $$repeat:message.content.user[user:*]$$
    <tr>
      <td>$$xml:[user].name$$</td>
      <td>
        $$repeat:[user].address[addr:*]$$
          $$xml:[addr]$$<br>
        $$/repeat$$
      </td>
      <td>$$xml:[user].telno$$</td>
    </tr>
  $$/repeat$$
</table>
```

- `$$xif...$$`

In this case, we will extend the repeat example to show how we might need to perform a conditional display. In this case the phone number column either displays a home phone number for all teleworkers or an internal extension for those based in the office.

```
<table BORDER="1" BGCOLOR="#fff0a0">
  <th COLSPAN="3" BGCOLOR="ffd080"><b><i><font SIZE="+1"
COLOR="#00a000">System Users</font></i></b></th>
  <tr>
    <td WIDTH="100" ALIGN="center"><B>Name</B></td>
    <td WIDTH="400" ALIGN="center"><B>Address</B></td>
    <td WIDTH="100" ALIGN="center"><B>Telephone</B></td>
  </tr>
  $$repeat:http_response.content.user[user:*]$$
    <tr>
      <td>$$xml:[user].name$$</td>
      <td>
        $$repeat:[user].address[addr:*]$$
          $$xml:[addr]$$<br>
        $$/repeat$$
      </td>
      <td>
        $$xif:[user].workType="Teleworker"$$
          $$xml:[user].homePhone$$
        $$xelse$$
          $$xml:[user].workExtn$$
        $$/xif$$
      </td>
    </tr>
  $$/repeat$$
</table>
```

- `$$else:...$$`

Lastly, it is possible to perform a multiple conditional statement using the extended `else` construct. Taking the previous example we can extend the switch to include an extra case for salesmen.

```
<td>
    $$xif: [user].workType="Teleworker"$$
        $$xml: [user].homePhone$$
    $$xelse: [user].workType="Sales"$$
        $$xml: [user].mobile$$
    $$xelse$$
        $$xml: [user].workExtn$$
    $$/xif$$
</td>
```

- `$$xprop:...$$` property substitution

The code snippet below shows how to get a property from the [tbase.properties](#) file into a template:

```
<hr>
<form METHOD=post
ACTION="$$xprop: [ApplicationServcer] form.string$$/appservlet">
    <h3>
Please enter your card details so that we may confirm your identity:
    </h3>
```

Default Identrus Message Writer

This message writer will handle messages with a mime-type that begins with “application/identrus-“ and a message type of [IdentrusConstants.IDENTRUS_MESSAGE](#).

Once at the message writer the following actions are carried out:

- Sign the message with the [IdentrusConstants.L1_IP_SC certificate](#), unless the iPlanet™ Trustbase Transaction Manager Message attribute identified by [IdentrusConstants.SIGNING_CERT_ATTR](#) is present. If this attribute is present then its value is taken as the purpose Id of the certificate/key used to sign the outgoing message.
- The outgoing message is DTD validated according to its DOCTYPE, this prevents malformed messages from leaving the TC.
- The DOCTYPE tag is set and the mime-type of the message is determined from the DOCTYPE
- The Certificate Bundle is extracted from the message and logged with the Certificate Log
- The XML of the outgoing message, minus the Certificate Bundle, is recorded in the message log
- The message is written back to the client.

Default Identrus Error Writer

This message writer will handle messages with a mime-type of “application/identrus-identrustransporterror” and a message type of [IdentrusConstants.IDENTRUS_TRANSPORT_ERROR](#)

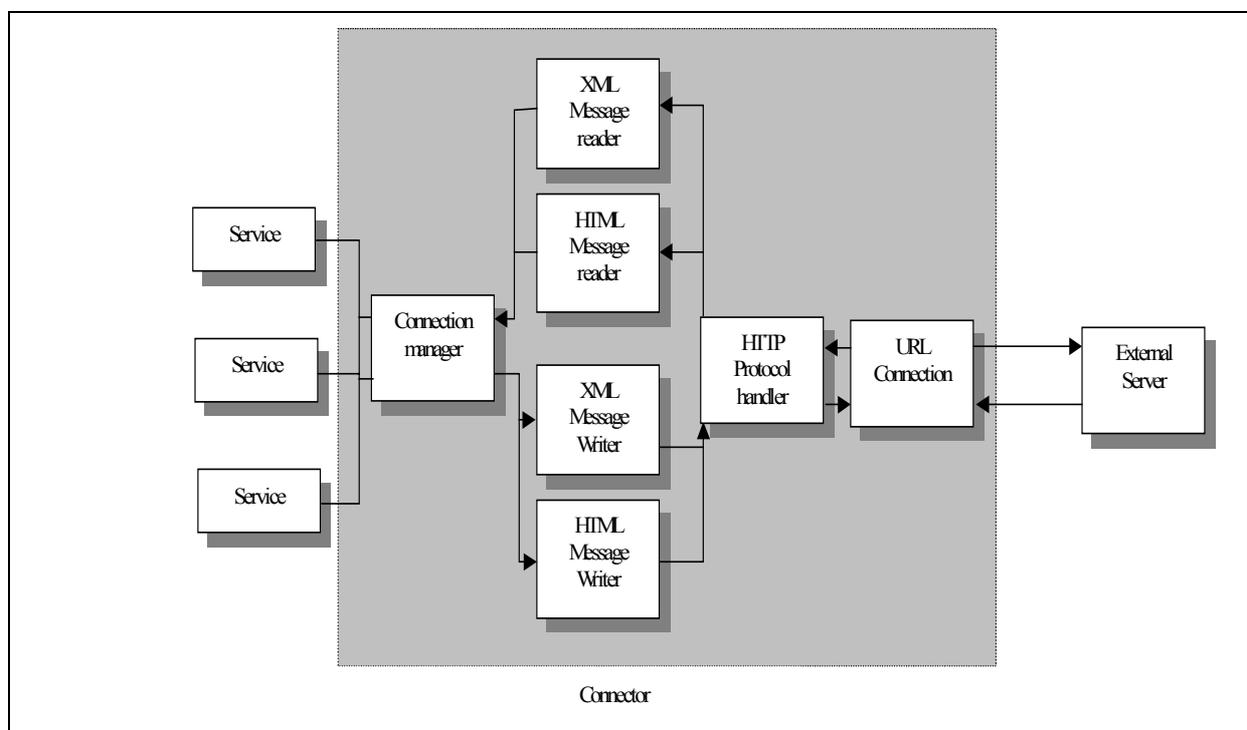
Once at the message writer the following actions are carried out:

- The outgoing message is DTD validated according to its DOCTYPE, this prevents malformed messages from leaving the TC.
- The DOCTYPE tag is set and the mime-type of the message is determined from the DOCTYPE
- The XML of the outgoing message is recorded in the raw log
- The message is written back to the client.

Connection Manager

On receiving a message iPlanet™ Trustbase Transaction Manager routes this message to an appropriate service which can process the message accordingly, and can, if required return a message in response. A problem arises when the particular iPlanet™ Trustbase Transaction Manager implementation, on receiving a message, routes the message through to its appropriate service, and the service cannot directly act upon this message. It needs additional information from another source in order to complete the task presented to it by the received message. Hence, iPlanet™ Trustbase Transaction Manager needs the facility to create new messages using specific message transport protocols and send these messages to specific external destinations. To accomplish this, the service needs to call the “Connection Manager”.

Figure 6 Connection Manager Architecture



Note: please consult your API for more information about this:

uk.co.jcp.tbase.connector

uk.co.jcp.tbase.xurl

uk.co.jcp.tbaseimpl.connector

The Connection Manager is essentially a service that can be used to send and receive messages to and from external entities. The Connection Manager is passed a iPlanet™ Trustbase Transaction Manager message and [Destination](#) Object. It uses this information to send and receive messages related to the specific message/destination object supplied. The Connection Manager Message Process works as follows:

- The Service calls the [Connector](#), passing the iPlanet™ Trustbase Transaction Manager Message containing the request to the external resource, and a [Destination](#) Object describing the external resource. [Destination](#) Objects are application-defined implementations, and are recognised by [ProtocolMap](#) objects that are application defined plug-in components to the [Connector](#). A null Message parameter may be passed, in which case no attempt will be made to write data to the [URLConnection](#) object when a connection has been established.
- The [Connector](#) calls the [ProtocolMapManager](#) to translate the [Destination](#) Object it was supplied with into a [ProtocolDescriptor](#) Object that specifies a URL for the connections, and a mime-type for the request content. The [ProtocolMapManager](#) determines which [ProtocolMap](#) from those registered with it (an administration activity) is responsible for dealing with the supplied Destination implementation, and passes the [Destination](#) Object to the [ProtocolMap](#).
- The [ProtocolMap](#) translates the [Destination](#) Object into a [ProtocolDescriptor](#) . This action may be simple and local, or it may be complex and involve further external requests.
- The [Connector](#) makes a connection to the external resource using the URL in the [ProtocolDescriptor](#) to form a [URLConnection](#).
- If a Message was passed to the [Connector](#), it sets the doOutput variable on the [URLConnection](#) to true, to signify it's intention to write data to the connection. The [Connector](#) then selects a [MessageWriter](#) from its registry of [MessageWriters](#), according to the type of the Message to be written (in the message Type attribute), and the required format, as specified in the [ProtocolDescriptor](#). The [MessageWriter](#) is called to translate the Message to the appropriate format, and write it to the [OutputStream](#) of the [URLConnection](#). Should no Message have been passed to the [Connector](#), the doOutputfield of the [URLConnection](#) is set to false, then nothing is written to the [URLConnection](#)..
- The [Connector](#) hands the [InputStream](#) of the [URLConnection](#) to the [ConnectionProtocolAnalyser](#), along with the mime-type of the response [determined from the [URLConnection](#)]. The [ConnectionProtocolAnalyser](#) determines a [ProtocolHandler](#) to call, using the mime-type of the response.
- The [ConnectionProtocolAnalyser](#) gives the selected [ProtocolHandler](#) a new iPlanet™ Trustbase Transaction Manager Message, and the [InputStream](#) . The [ProtocolHandler](#) reads from the [InputStream](#) to determine the Message type of the response, which it sets on the Message (it may also set a Context identifier, although this will be ignored in this setting), before returning.
- The [Connector](#) selects a [MessageReader](#) based on the response Message type, and the response [InputStream](#) mime-type. The

[MessageReader](#) is called to complete the parsing of the response. After the [MessageReader](#) completes, the [Connector](#) releases any resource used in making the external connection, and returns the parsed Message to the requesting Service.

Note: iPlanet™ Trustbase Transaction Manager provides a default connector that implements the connector interface defined in uk.co.jcp.tbaseimpl.connector.DefaultConnector. Most of these classes are supplied as part iPlanet™ Trustbase Transaction Manager. Some are not. For instance, java.net.URLConnection is part of the core Java API.

Protocol Map Manager

The Connection Manager provides the ability for Services in the iPlanet™ Trustbase Transaction Manager framework to communicate with other external entities. To do this the Service passes a message and [Destination](#) to the [Connector](#) (Connection Manager), which routes the message to the specific external entity. To find the address of the communicating party, the supplied [Destination](#) is passed to the [ProtocolMapManager](#). This [ProtocolMapManager](#) then passes the Destination to its associated [ProtocolMap](#) registered with the [ProtocolMapManager](#). The specific [ProtocolMap](#) then returns a [ProtocolDescriptor](#) object associated with this [Destination](#). Each specific [ProtocolMap](#) decides which [ProtocolDescriptor](#) to return for any supplied [Destination](#).

To create a new application specific protocol mapping system, the following steps have to be followed:

- Creation of a Specific [ProtocolMap](#).
- Creation of a Specific [Destination](#)
- Registering the new [ProtocolMap](#) with the [ProtocolMapManager](#).
- Populating the [ProtocolMap](#) with data.

Note: please consult the framework uk.co.jcp.tbase.connector and the default connection manager class implementation can be found in uk.co.jcp.tbaseimpl.connector

These are now described in turn.

- Creation of a Specific [ProtocolMap](#)

All [ProtocolMap](#) objects must implement the [ProtocolMap](#) interface. They must also provide an empty/default constructor, due to the fact that the [ProtocolMapManager](#) will dynamically instantiate the [ProtocolMap](#) object in order to register the class with itself.

The [ProtocolMap](#) enforces the following methods:

```
public Enumeration getDestinationTypes ()
```

Returns a list of the package qualified class names of the [Destination](#) implementations that this [ProtocolMap](#) recognises.

```
public ProtocolDescriptor getProtocolDescriptor(Destination destination)
    throws InvalidDestinationException
```

Translates an application's specified [Destination](#) object into a [ProtocolDescriptor](#), specifying the URL and mime type for the connection.

Note: The application specific [ProtocolMap](#) must not contain any objects that are not serializable, as the [ProtocolMapManager](#) has persistence. Consult your API uk.co.jcp.tbase.connector.ProtocolMap for more information on this.

- Creation of a Specific Destination

All Destinations must implement the [Destination](#) interface. This interface does not enforce any methods, leaving all methods and data purely application specific, thus relying on its associated [ProtocolMap](#) to have prior knowledge of its format.

Note: The application specific Destination must not contain any objects that are not serializable, as the [ProtocolMapManager](#) has persistence. Consult your API uk.co.jcp.tbase.connector.Destination for more information on this.

- Registering the new [ProtocolMap](#) with the [ProtocolMapManager](#).

Once an application specific [ProtocolMap](#) and [Destination](#) have been constructed, the [ProtocolMap](#) has be registered with the [ProtocolMapManager](#) before any Service in iPlanet™ Trustbase Transaction Manager can utilise its function.

Note: If a new application specific [ProtocolMap](#) contains a reference to a specific Destination used by an existing [ProtocolMap](#) registered with the [ProtocolMapManager](#) then the new [ProtocolMap](#) will be prevented from registering with the [ProtocolMapManager](#).

Protocol Maps can be registered in the [tbase.properties](#) file that contains a property section called "ProtocolMapManager". In this section the application specific [ProtocolMap](#) objects can be supplied. For instance,

```
[ProtocolMapManager]
protocol.map=uk.co.jcp.tbase.connector.SimpleProtocolMap
```

Note: This [tbase.properties](#) file will only be read once unless the configuration object associated with the instance of iPlanet™ Trustbase Transaction Manager is removed. Note also, the full package name must be supplied.

URL Connection Implementation

Internally in the Connection Manager, once the [ProtocolMapManager](#) returns a [ProtocolDescriptor](#) from a supplied [Destination](#), the URL string representation is obtained from the [ProtocolDescriptor](#). From this string representation of a URL, an [URLConnection](#) object is obtained, from which the output and input streams are used to send and receive the message from the Connection Manager ([Connector](#)). This conversion of a string representation of a URL to a [URLConnection](#) object is carried out through the XURLxxxx classes. The classes used are:

- [XURL](#)
- [XURLStreamHandler](#)
- [XURLStreamHandlerFactory](#)

These classes are needed, due to the fact that the existing URL class hierarchy supplied by javasoft is not extensible in an EJB format. This is because in the existing framework, if the URL class does not recognise an URL format then it looks to see if a [XURLStreamHandlerFactory](#) is loaded, which could supply the unknown URL format. This factory can only be set once, so in an EJB environment it is impossible to ascertain whether [URLStreamHandlerFactory](#) has been instantiated, so making this framework unusable. Hence, we have produced an extensible framework that can work in any environment. These objects are identified as [XURL](#), [XURLStreamHandler](#), and [XURLStreamHandlerFactory](#).

The process flow is as follows:

- A [XURL](#) object is constructed using the string representation of the URL supplied by the [ProtocolDescriptor](#).
- To obtain the [URLConnection](#) object the [XURL](#)'s method `openConnection()` is called.
- Inside this method the [XURLStreamHandlerFactory](#) is called.
- The [XURLStreamHandlerFactory](#), if not called before initialises itself from the `tbase.properties` file, registering [XURLStreamHandler](#) classes in the section "[XURLStreamHandlerFactory](#)", each class mapped against the value "url.stream.protocol".
- The [XURLStreamHandlerFactory](#) is searched using the protocol that will be used to connect to the specified URL as the key.
- If the protocol supplied maps to a [XURLStreamHandler](#) registered with the [XURLStreamHandlerFactory](#), then the [XURLStreamHandler](#) object is called in order to provide the [URLConnection](#) object required.
- If the protocol does not match any [XURLStreamHandler](#)s stored within the [XURLStreamHandlerFactory](#), then an URL object is created from the supplied string representation in the [XURL](#), and the [URLConnection](#) object is obtained directly from that.

Providing this extra framework, JCP supplies the ability to override existing protocol formats, in order to introduce any specific enhancements, i.e. HTTPS, where our SSL is directly tied in with our Certificate Storage interfaces.

Note: please consult your API for more information this: uk.co.jcp.tbase.xurl

Routing

The router component has two specific functions:

- Provide a means of authorising requests to particular services
- Provide a means of controlling the flow of transactions

The router has a rule based architecture that allows the developer to modify the behaviour of a transaction in many cases without recourse to writing Java Code. This is achieved by producing a set of XML rules that determine the flow of a message through the system, and comparing these rules to the values held in the iPlanet™ Trustbase Transaction Manager's message.

The following sections define the concepts used in the router architecture, the function of the router, and the default implementation for Identrus message processing.

Messages

An iPlanet™ Trustbase Transaction Manager message is an entity that moves from one part of iPlanet™ Trustbase Transaction Manager to another. The message is an encapsulation of the message sent to iPlanet™ Trustbase Transaction Manager together with internal data about the message. The basic components of an iPlanet™ Trustbase Transaction Manager message are as follows:

- Message Type
- Attributes
- Data Content

Note: Further details of the content of an iPlanet™ Trustbase Transaction Manager message can be found in the API documentation for uk.co.jcp.tbase.environment.Message.

The Message Type and Attributes determine how the framework itself will handle the iPlanet™ Trustbase Transaction Manager Message. Anything in the Data Content is treated as opaque by the iPlanet™ Trustbase Transaction Manager framework; i.e. any changes to the Data Content will not be visible to the iPlanet™ Trustbase Transaction Manager framework.

The Router determines the state of a message by interacting with the message's attributes. The router uses a set of pre-conditions based on attribute values to determine the state of the message and therefore which actions should be applied to that message. These pre-conditions and actions are stored in the form of router rules that are described later.

Message Attributes

Message Attributes are just name-value pairs in the same way as Java properties. Attribute names and values are always string types.

A protocol handler, message reader, message writer, the router or a service may read and write message attributes. This allows each stage in the message processing pipeline to add its own state information.

Attributes may be defined by any user application code, but a number of standard attributes are used by the iPlanet™ Trustbase Transaction Manager framework. These are defined below:

- `messageType`
This attribute is treated specially by the Message API class, the message type is used in conjunction with the message MIME type to determine which Message Readers, Message Writers and Protocol Handlers should be used in the processing of this message. It is the responsibility of the Protocol Handler for the MIME type to set this attribute correctly. The `messageType` attribute may change during the processing of a message.
- `security.role`
This attribute contains the security role to which this message has been attributed. The authenticator service sets this attribute to one of the following values: the role name for successful authentication, default when no role can be assigned

- security.cert
The encoding of the certificate used as authentication evidence, this is normally the SSL client certificate, but in the case of Identrus Messages it is the certificate that is used to verify the mandatory level 1 signature on the message.
- security.cert_encoding
This is always BASE64.
- security.cert_present
Is there a valid certificate present for authentication purposes?
- security.username
A username for authentication purposes.
- security.password
A password for authentication purposes.
- security.user_pass_present
Is there any valid username/password evidence for authentication present on this message.
- security.auth_failed
Authorisation failed when an executeServiceDirective was called, this means that there was no role/serviceName mapping in iPlanet™ Trustbase Transaction Manager. This means that the message is not authorised to be sent to the requested service.

Note: The attribute names for security specific attributes are declared as constants in uk.co.jcp.tbbaseimpl.authenticator.SecurityContext Also see uk.co.jcp.tbbase.environment.attribute

Identrus Message Attributes

In the case of the Identrus message processing, the following additional message attributes are used:

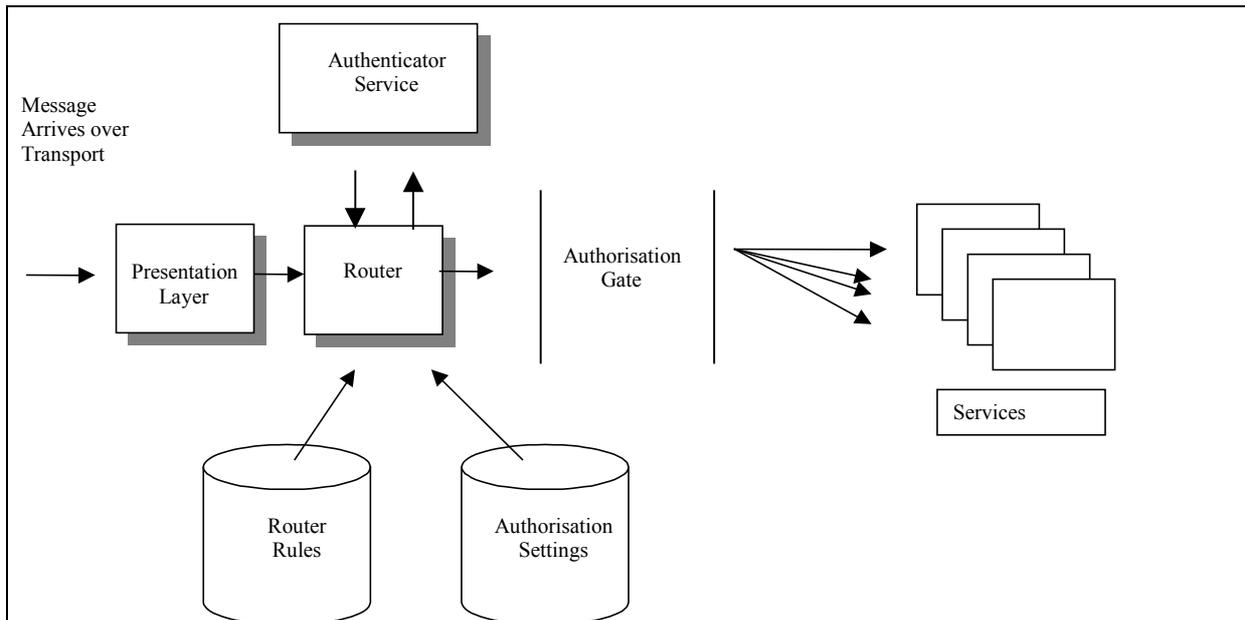
- DocType
The DOCTYPE of the XML message that will identify the type of message. e.g. CSCRequest or PingRequest
- SigningCertPurposeld
This optional attribute is only interpreted by the default Identrus Message Writer, if present it specifies the Certificate Store Attribute to be used to extract the signing certificate and private key. If this is not present then the message writer assumes that the Level 1 Inter-participant Signing Certificate is to be used.
- XMLSystemId
This contains the 'system identifier' attribute from the incoming XML message – it is used by the Identrus Message Writer to fill in the XML System identifier on the corresponding response message.

Note: The attribute names for Identrus specific attributes are declared as constants in com.iplanet.trustbase.identrus.IdentrusConstants

Router Architecture

The Router performs the core rule based message routing which is central to the flexibility offered by iPlanet™ Trustbase Transaction Manager, but it is also responsible for gating access to services on the basis of authentication information and authorisation settings. The figure below shows the basic elements involved with getting a message through iPlanet™ Trustbase Transaction Manager and into the service that can process it.

Figure 7 Router Architecture



- The message arrives at iPlanet™ Trustbase Transaction Manager and makes its way through the presentation layers (protocol handlers and message readers), during which time the messageType is identified and set and a number of other standard attributes are also set. These attributes form the basis of the authentication evidence for the message.
- The message then arrives at the router. All messages are initially sent to the service identified by the name “Authenticator”. There is a default ‘Authenticator’ service that comes with iPlanet™ Trustbase Transaction Manager, but this may be replaced by an alternative implementation if required.
- The message is sent to the default Authenticator Service that examines the attributes on the message to assess the authentication level of the message. The function of this service is defined in more detail below.
- The message arrives back at the router with a security context attribute (security.role) set, this context attribute determines the Role into which the message has been authenticated.

- The target service is located by processing the message through the routing rules and then the message is sent to the service. The detailed definition of these rules is given later on in this section.
- Before the message arrives at the service it must pass through an authorisation gate, this authorisation gate checks that there is a mapping between the role attribute assigned to the message and the target service. If such a mapping exists then the message is passed to the service, if no such mapping exists then the message is sent to an authorisation error service.

The stages described above are the basic stages which all iPlanet™ Trustbase Transaction Manager messages go through. The processing pipeline has been streamlined from the developers point of view when working with Identrus XML messages and is defined in the Default Routing section in this chapter.

Authentication and Authorisation

The authentication mechanisms of iPlanet™ Trustbase Transaction Manager are not a separate component. Authentication data is gathered by the default iPlanet™ Trustbase Transaction Manager framework, this can be added by domain specific services. These services must map an identity onto a particular set of attributes, and iPlanet™ Trustbase Transaction Manager then uses these attributes to impose access control on other services.

Authentication

In order to ensure that the authentication and authorisation mechanisms are suitable for the widest variety of organisation, the iPlanet™ Trustbase Transaction Manager framework does not perform the authentication operations itself, but expects a service to be capable of verifying identity and mapping this to a particular role or roles. The system comes supplied with a basic identity that uses a digitally signed data item and an accompanying X509 digital certificate. The service is capable of mapping this data into a role or set of roles. The basic authentication information may be presented in one of two ways:

- Implicitly - within the Transport protocol e.g. SSL
- Explicitly - within the message content e.g. signed message block of an Identrus XML message

The architecture is designed to reduce both of these to explicit authentication by removing the certificate used in the SSL handshake and using it to decorate the message passed to the authentication service. This allows the target system to impose different levels of authentication that would be difficult to achieve with an implicit 'black box' approach to transport level authentication.

The default service also provides a means of mapping a username and password onto a role for use by the router when authorising requests. This mechanism is commonly used when presenting configuration and administration information using the HTML templating mechanism. It is recommended that the closed community mechanism be replaced by an enterprise wide username and password mechanism if the platform is to be used for large scale HTML based applications.

Authorisation

Given that we can identify the originator of a message, either through the digital certificate or an alternative means, we now require a mechanism of:

- Identifying the appropriate level of authorisation for the user
- Enforcing access to each service on the basis of authorisation

The identification of the proper authorisation level (or role) is performed by a service named 'Authenticator' in iPlanet™ Trustbase Transaction Manager. This service has a default implementation that uses the username/password tables and certificate tables in the authorisation database to determine a mapping for the security attributes on the message. The mapping of this information is described in the iPlanet™ Trustbase Transaction Manager Configuration & Installation Guide. Enforcing the access to each service is done automatically by the router. If a routing rule is matched and the body contains an 'executeServiceDirective' then before that directive is executed the router establishes if there is a mapping between the role assigned to the message and the required service. These mappings are held in the

authorisation table and this table's use is described in the Trustbase Configuration & Installation Guide. If the mapping exists then the message is passed to the service. If no mapping exists then a new attribute is added to the message "security.auth_failed" and the message is sent back to the router. The routing rules should always detect an authorisation failure and execute an appropriate service to send an 'unauthorised' response to the client. Services may be executed unconditionally by using the 'unauthorisedExecuteServiceDirective' that does not attempt check for authorisation.

Default routing

In order to provide a simple development path for Identrus Messaging services, iPlanet™ Trustbase Transaction Manager provides a default routing mechanism for Identrus Messages. This default routing is applied to all Identrus Services developed and deployed as described in the section on Building Identrus solutions on page 74 of this document.

It is anticipated that this default routing will be sufficient for most Identrus Services in the near term. Where the default routing is inadequate, the developer must write new routing rules.

Router Rules

The updates made to the Router for version 2.2 of the iPlanet™ Trustbase Transaction Manager allow a very simple approach to basic rules handling to be taken. In previous versions of the system, all services had to define a set of rules that explicitly mapped the transport of a message to and from the service. Now that rules can be defined dynamically by the system, this basic behaviour can be automatically defined. We now discuss how this is achieved.

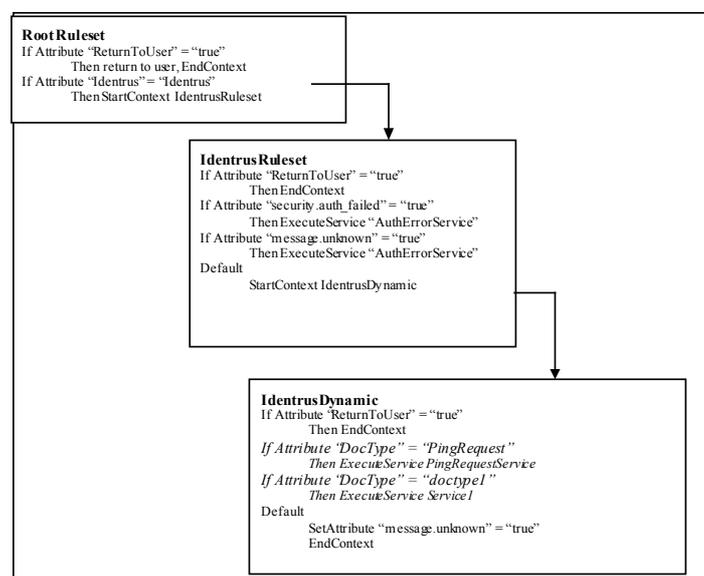
Routing to service

The deployment descriptor contained in the service JAR file will name the service that it contains. At system startup, the rules handler will determine which services are to be loaded, and for each service in the list will dynamically create a routing rule that forwards messages with the specified doctype to the service. If a service has a user-defined ruleset, the dynamically created rule will pass the message to this ruleset instead of directly to the service. This user defined ruleset lives in the JAR file with the service implementation and will normally not be required for services with simple message flows.

Return path

In order that the routing be kept as simple as possible, a message will be returned to a user if a specified attribute has been set on the message. This "ReturnToUser" attribute should be set by the service once processing is complete. This also applies to error messages produced at the service level.

Figure 8 Default routing rules



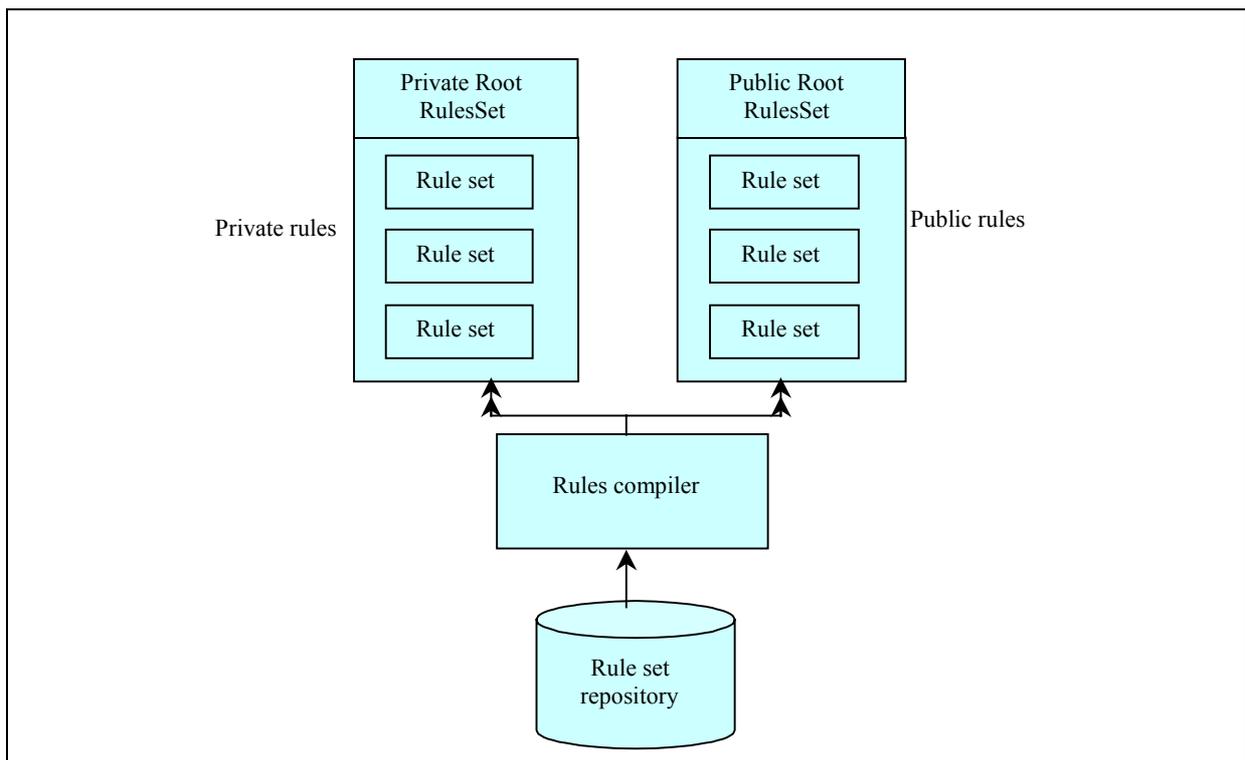
Advanced Routing

This section describes routing rules in detail, for simple Identrus services the developer should not have to write any routing rules in order to get messages to their service. If the default routing is inadequate or non-Identrus message processing is required then developers must resort to writing their own rules.

Routing Rulesets

The routing tables have two logical areas, one private and one public. Public rule sets may be added or configured through the management interface, private rule sets are used by iPlanet™ Trustbase Transaction Manager to provide basic services and may not be configured by the developer or systems administrator. At the top of this structure is the Root Rule Set which governs the way in which the private rules are executed and if none of the private rules match the parameters then the Public Root Rule Set is executed. System developers may then link their own task Rule Sets from this Rule Set (See Rule set definition in the next section)

Figure 9 Rule Sets



Each area contains a number of rule sets. Each public rule set will have a name, usually indicating the business task performed by the set of rules. The public rule sets are loaded from disk each time the platform starts up, and may be loaded explicitly by name as the result of a rule execution in the root rule set.

When a message is passed to the router, it will have already been associated with a context. Each context has an association with a rule set. This implies:

- The message is operating in its root context, and has the root rule set associated with it.

- The message is operating in a sub-context, and has a rule set from the public rules area associated with it.
- The exception to this is if an iPlanet™ Trustbase Transaction Manager message is being processed in which case a private rule set will be associated.

Rule Sets are processed sequentially, and the first rule to match the current conditions is used. The router takes each rule from the rule set in turn, starting with the first, and attempts to match the message type with a rule pre-condition. If a match is found then the action is executed. If no match is found then an error condition is returned. It is important to bear in mind that because rules are processed sequentially it is necessary to keep the more explicit rules at the beginning of the rule set with more general rules occurring later. Otherwise it is entirely possible that a general rule could preclude other rules from ever being executed.

It is considered good practice to keep the root rule set as small as possible. To achieve this, a rule in the root rule set should only match the first message in a business task i.e. an unsolicited message. The action should always be to start a new context, indicating the name of a rule set (and potentially the rule name to start at), that the message should be executed in. In doing this, the router will create a new context and associate the named rule set with it. The message will now be processed in the same manner.

A rule set is a single complete XML document that conforms to the DTD outlined in the next section. When the router starts up it looks in the Router section of the iPlanet™ Trustbase Transaction Manager initialisation file for the directory to find the rules that have been written by hand. It then looks for .xml files, and compiles these into the public rule set area. Each Rule Set contains zero or more rules as shown in the DTD fragment shown below:

```
<!ELEMENT ruleSet (rule*)>
<!ATTLIST ruleSet name CDATA #REQUIRED>
```

Each rule in the ruleset has a structure that is defined in the next section.

Router Rule Syntax

This section works through the Rule set DTD describing the purpose of each component. Each rule has three components:

- An attribute which names the rule. This is optional; the name may be used when starting new contexts if a specific rule should be evaluated.
- A set of pre-conditions that must all be met before the execution of the directives in the body.
- The body of the rule, this is the list of actions to be carried out if the pre-conditions are all met.

A typical rule name is illustrated below.

```
<!ELEMENT rule ( preconditions, body )>
<!ATTLIST rule name CDATA #IMPLIED>

<!ELEMENT preconditions (attributeCondition*)>
```

The preconditions are made up of zero or more Attribute conditions. Each attributeCondition must be satisfied for the body to be executed.

```
<!ELEMENT attributeCondition EMPTY>
<!ATTLIST attributeCondition name CDATA #REQUIRED>
<!ATTLIST attributeCondition operator ( greaterThan | greaterThanEqualTo
| lessThan | lessThanEqualTo | notEqual | startsWith | contains |
endsWith | equals ) "equals">
<!ATTLIST attributeCondition valueType ( any | null | string ) "string">
<!ATTLIST attributeCondition value CDATA #IMPLIED>
```

Each AttributeCondition is evaluated by iPlanet™ Trustbase Transaction Manager as though it were in the form:

```
IF ( name operator value ) then pre-condition met
```

The valueType is normally defaults to 'string' and is not required in rules that have a value field. Where there is a requirement to test for the presence of an attribute, or if it needs to be tested against null, then the valueType field is used and the value field is ignored/not required.

```
<!ELEMENT body ((setAttribute | executeServiceDirective |
unauthorisedExecuteServiceDirective )*, (startContextDirective |
endContextDirective | returnToUserDirective )?)>
```

Once the pre-conditions have been met the body of the rule can be executed. The body of the rule may contain multiple actions that are carried out sequentially in the order that they appear. There are three actions, that are considered to be 'terminal' actions, which may only appear once in a rule and they must appear as the final action. Figure 10 DTD Rule Body shows each of the allowable body directives.

Figure 10 DTD Rule Body

Directive	Description
SetAttribute	Sets a iPlanet™ Trustbase Transaction Manager attribute on a message
executeServiceDirective	Transfers control from the router to the named service, ensuring that the mandatory authorisation checks are passed
unauthorisedExecuteServiceDirective	Transfers control from the router to the named service, without performing any of the mandatory authorisation checks.
StartContextDirective	Transfers router control from one rule set to another
EndContextDirective	Ends the current context, and starts rule search again
ReturnToUserDirective	Returns control from the router to the system user

```

<!ELEMENT setAttribute EMPTY>
<!ATTLIST setAttribute name CDATA #REQUIRED>
<!ATTLIST setAttribute value CDATA #REQUIRED>
<!ATTLIST setAttribute location (MESSAGE | CONTEXT) "CONTEXT">
<!ATTLIST setAttribute type (ASCEND_TRANSITIVE | NONE |
INHERIT_TRANSITIVE | INHERIT) "ASCEND_TRANSITIVE">

```

The `setAttribute` action allows the router to add attributes to either the message or the context. The name, value, location and type fields are all the same as the settings for creating an attribute via the API.

Note: See API documentation for `uk.co.jcp.tbase.environment.attribute.Attribute` for more details

```

<!ELEMENT executeServiceDirective EMPTY>
<!ATTLIST executeServiceDirective name CDATA #REQUIRED>

```

An `executeServiceDirective` must specify the name of the service to invoke. This form of executing a service will always perform an authorisation check before passing the message to the requested service.

The name of the service relates to the name given to the service in the ServiceRegistry section of the [tbase.properties](#) file where each new `service.description` entry consists of a service name and a classpath for that particular service.

```

<!ELEMENT unauthorisedExecuteServiceDirective EMPTY>
<!ATTLIST unauthorisedExecuteServiceDirective name CDATA #REQUIRED>

```

An `unauthorisedExecuteServiceDirective` must specify the name of the service to invoke. This form of executing a service will not perform any authorisation check before passing the message on to the service.

The name of the service relates to the name given to the service in the ServiceRegistry section of the [tbase.properties](#) file where each new `service.description` entry consists of a service name and a classpath for that particular service.

```

<!ELEMENT startContextDirective EMPTY>
<!ATTLIST startContextDirective ruleset CDATA #REQUIRED>
<!ATTLIST startContextDirective rule CDATA #IMPLIED>

```

In executing a `startContextDirective` the router discards the current rule set it has been searching and loads the rule set named in the directive. If an individual rule is also named in the directive then the router attempts to find that rule in the new rule set, otherwise the router begins at the start of the rule set and searches for the first rule with fulfilled preconditions.

A start context directive must contain a name for the rule set to be used. It may optionally contain a rule name to be executed in the specified rule set.

```

<!ELEMENT endContextDirective EMPTY>

```

An end context directive has no attributes or children.

```
<!ELEMENT returnToUserDirective EMPTY>  
<!ATTLIST returnToUserDirective endContext ( true | false ) #REQUIRED>
```

A return to user directive tells the router whether to end the current context or not.

In some cases the context serves no further purpose once a service has been executed on behalf of the user and may therefore be discarded. However it is possible to envisage scenarios e.g. whenever authentication is required for a user to run services, where one would wish to maintain the 'Authorised' context between user interactions, thus allowing the user to repeatedly use services without being required to be authorised every time.

Note Example rules and rulesets for core Trustbase services that comply with these data structures can be found in <install_dir>/Trustbase/TTM/current/Config/Rules/Public

Complete Router Rule DTD

The complete DTD that specifies all of the allowable routing rules within iPlanet™ Trustbase Transaction Manager is shown below.

```

<!-- A rule set is zero or more rules -->
<!ELEMENT ruleSet (rule*)>
<!ATTLIST ruleSet name CDATA #REQUIRED>
<!-- A rule optionally has a name attribute, preconditions & body elements -->
<!ELEMENT rule ( preconditions, body )>
<!ATTLIST rule name CDATA #IMPLIED>
<!-- preconditions element is a set of zero or more attribute condition elements -->
<!ELEMENT preconditions (attributeCondition*)>
<!--An attribute condition is an empty tag with name, valueType & value attributes -->
<!-- the name is mandatory. valueType may be any, null or string. if string is -->
<!-- selected then value must be present. If any is selected then the value -->
<!-- matches any value and if null is selected then the value is disregarded -->
<!-- and assumed to be null -->
<!ELEMENT attributeCondition EMPTY>
<!ATTLIST attributeCondition name CDATA #REQUIRED>
<!ATTLIST attributeCondition operator ( greaterThan | greaterThanEqualTo | lessThan |
lessThanEqualTo | notEqual | startsWith | contains | endsWith | equals ) "equals">
<!ATTLIST attributeCondition valueType ( any | null | string ) "string">
<!ATTLIST attributeCondition value CDATA #IMPLIED>
<!-- The body contains one of four directives -->
<!ELEMENT body ((setAttribute | executeServiceDirective |
unauthorisedExecuteServiceDirective )*, (startContextDirective | endContextDirective
| returnToUserDirective )?)>
  <!-- A setAttribute directive -->
  <!ELEMENT setAttribute EMPTY>
  <!ATTLIST setAttribute name CDATA #REQUIRED>
  <!ATTLIST setAttribute value CDATA #REQUIRED>
  <!ATTLIST setAttribute location (MESSAGE | CONTEXT) "CONTEXT">
  <!ATTLIST setAttribute type (ASCEND_TRANSITIVE | NONE | INHERIT_TRANSITIVE |
INHERIT) "ASCEND_TRANSITIVE">
  <!-- A start context directive must contain a name for the rule set to be used. -->
  <!-- optionally contains a rule name to be executed in the specified rule set -->
  <!ELEMENT startContextDirective EMPTY>
  <!ATTLIST startContextDirective ruleset CDATA #REQUIRED>
  <!ATTLIST startContextDirective rule CDATA #IMPLIED>
  <!-- An end context directive has no attributes or children -->
  <!ELEMENT endContextDirective EMPTY>
  <!-- execute directive must specify a name for the service to be executed -->
  <!ELEMENT executeServiceDirective EMPTY>
  <!ATTLIST executeServiceDirective name CDATA #REQUIRED>
  <!-- unauthorisedExecute directive must specify service name to be executed -->
  <!ELEMENT unauthorisedExecuteServiceDirective EMPTY>
  <!ATTLIST unauthorisedExecuteServiceDirective name CDATA #REQUIRED>
  <!-- A return to user directive is empty and has no attributes -->
  <!ELEMENT returnToUserDirective EMPTY>
  <!ATTLIST returnToUserDirective endContext ( true | false ) #REQUIRED>

```

Configuration Management

There are two elements to the Configuration Manager sub-system, the Configuration Manager itself and its Configuration Store. The Configuration Manager is the central access point for all configuration requests and is responsible for maintaining Configuration Objects in a store. The Configuration Manager co-ordinates access to these objects for both read and update purposes.

Configuration Objects

Configuration Objects hold persistent data for Services and their Configuration Services. A given Service may reference any number of Configuration Objects, and any number of Services or Configuration Services may reference a given Configuration Object. Note that iPlanet™ Trustbase Transaction Manager components themselves use Configuration Objects. A Configurable Entity is any Service or component that uses Configuration Objects and the Configuration Manager. A Configuration Service is a Service that implements a read-write interface to the Configuration Object, which usually involves some graphical or HTML based interface allowing a user to change the values stored in the configuration object.

All configuration objects must implement the [ConfigurationObject](#) interface to ensure that the [ConfigurationManager](#) can manipulate them correctly. Configuration Objects are indexed in the store by ConfigUID objects.

Note: Further details of the interface that a [ConfigurationObject](#) must present can be found in the API documentation for uk.co.jcp.tbase.config.ConfigurationObject

Configuration Manager

A Service uses the [SingletonConfigManager](#) class to gain a reference to the Configuration Manager.

The Configuration Manager allows callers to get readable or writeable versions of Configuration Objects identified by ConfigUID objects.

There are two methods for getting a Configuration Object:

- The first returns a readable deep copy of the Configuration Object that can be used to set up a Service or for populating Configuration Service views. Configuration Objects gained by this method cannot have their changes applied back to the Configuration Store – this is because the returned Configuration Object does not contain a valid [ConfigurationLock](#) object. This form of Configuration Object will always be provided, regardless of whether there is an outstanding lock on the requested object, i.e. requests for read-only copies of [ConfigurationObjects](#) are never refused by the [ConfigurationManager](#).
- The second method provides a deep copy of the Configuration Object that contains a configuration lock valid for a fixed period of time. This lock period will, itself, be configurable. After this period the lock will become invalid and any attempt to apply a change to the Configuration Store will be denied. The lock expires to provide a simple solution to deadlocking issues caused by developers not releasing allocated locks.

Updates and deletions of Configuration Objects can only be performed with a Configuration Object that contains a valid lock object.

Note that a Service can only register a Configuration Object with a given ConfigUID if one does not already exist with that ConfigUID.

Configurable Entities may register for notification when a Configuration Object changes. When a change occurs, registered entities receive an event object that contains a read-only copy of the changed Configuration Object.

Note: Further details of the methods that the [ConfigurationManager](#) provides to manipulate [ConfigurationObjects](#) can be found in the API documentation for [uk.co.jcp.tbase.config.SingletonConfigManager](#) and [uk.co.jcp.tbase.config.ConfigManager](#).

Configuration Store

The Configuration Store is responsible for maintaining persistent storage of Configuration Objects. The version supplied with iPlanet™ Trustbase Transaction Manager uses an underlying JDBC implementation to store data. Alternative Configuration Stores may be implemented to store Configuration Objects in any given format.

The Configuration Store implements only primitive read, write and delete methods to access Configuration Objects via their ConfigUID key – the Configuration Manager handles higher level semantics such as locking.

The default JDBC store implementation cannot be replaced by developers and its interface is not public.

Configuration Services

A Configuration Service is the element that knows of the configurable attributes of a [ConfigurationObject](#). It provides the functionality to present these attributes to an administrator via user interface, iPlanet™ Trustbase Transaction Manager itself uses HTML forms to present configuration data.

The Configuration Services supplied with iPlanet™ Trustbase Transaction Manager use the HTTP Reader and [ScriptWriter](#) classes to generate HTML pages/forms that allow information to be changed by the administrator. The Configuration Service is capable of configuring all the individual iPlanet™ Trustbase Transaction Manager components, including protocol analysers, message analysers, readers and writers, router and services. Upon receipt of updated attributes the Configuration Service is responsible for registering the changes with the Configuration Manager that will in turn notify other interested parties.

At least one Configuration Service will exist for each Configurable Entity registered in the platform. Note that there will only be a single runtime instance of any given Configuration Service per deployment, no matter how many instances of the actual entities that it configures are in existence.

Note: There are no API classes associated with a Configuration Service, it is a normal iPlanet™ Trustbase Transaction Manager service and therefore must implement the [uk.co.jcp.tbase.service.Service](#) interface.

Chapter 6

Standard Services

Services provide actions in decision making process. Each service is a plug in component to the iPlanet™ Trustbase Transaction Manager's framework providing an implementation of the iPlanet™ Trustbase Transaction Manager's Service interface (See Configuration Guide for more details).

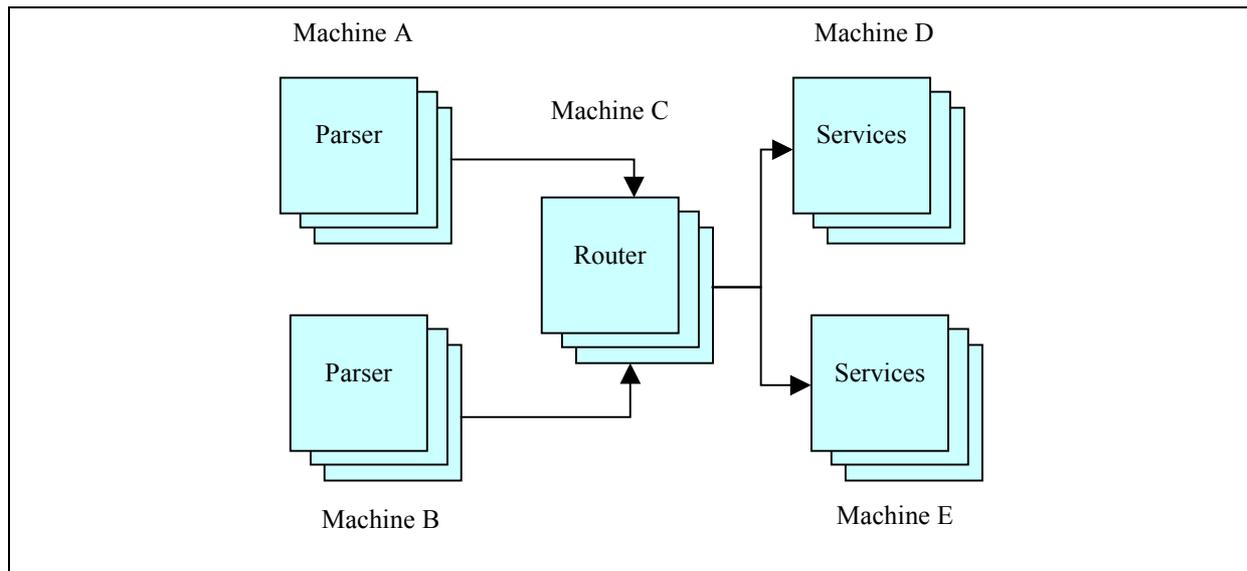
iPlanet™ Trustbase Transaction Manager is designed to run on a variety of platforms, in particular application servers. These provide a solution with the ability to service a large number of users at any single point in time. This is achieved by replicating the components i.e. running classes over a number of Java Virtual or physical machines. The iPlanet™ Trustbase Transaction Manager's Framework has been designed to fulfil all of these requirements. Each of the major sub-systems is capable of being replicated a number of times over a number of physical machines, but to achieve high volume processing the iPlanet™ Trustbase Transaction Manager's Framework imposes a number of constraints on a service.

Overview

A service should be:

- Stateless – Each component should be capable of processing a message irrespective of the context of the message
- Granular – Each component must perform a limited number of tasks, reducing the risk of a single point of failure

Figure 11 Component replication



A service is directly related to the processing of a message and must be stateless to allow replication.

These services usually fall into one of three major categories:

- Authentication and Authorisation
- Business processing including integration into existing systems
- User interaction

In dividing the function into discreet services that service one of these areas a variety of applications may be built over the same common infrastructure.

In particular, a common set of authentication and authorisation services will be used across the enterprise so devising services that only perform one or both of these tasks reduces the work required for each application.

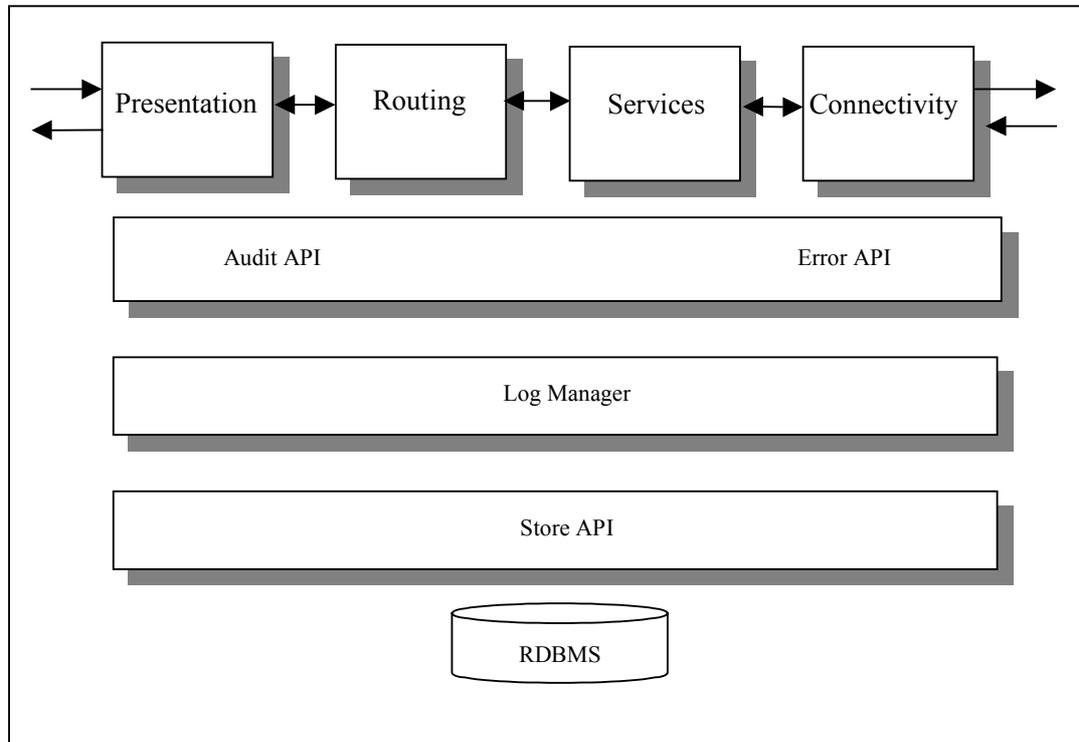
Error and Audit Logging

iPlanet™ Trustbase Transaction Manager provides a set of libraries for use with both error and audit logging. The default implementation of these libraries allow the framework and business components to store information within a relational database for use by external browsers and error management consoles.

Both error and audit logs are controlled by a log manager. The log manager stores the data in an appropriate location depending upon its current configuration. This design provides a location independent mechanism for recording error and audit information, and allows components to be re-used between different iPlanet™ Trustbase Transaction Manager implementations. This log manager is activated each time iPlanet™ Trustbase Transaction Manager boots. Initially the configuration of the log manager is read from an initialisation file, but in subsequent boot sequences this is read from the persistent configuration object generated by the manager.

Overview

Figure 12 iPlanet Trustbase Transaction Manager log manager



The data to be logged by the log manager is supplied by the requesting component in the form of an error or audit log object. The log manager will check the type but not the content of the object being logged. The type checking involved ensures that the object passed is a sub-class of the iPlanet™ Trustbase Transaction Manager defined error log and audit log classes. This allows the developer to define new error and audit log types for use specifically by the solution e.g. Billing logs.

The Log manager is abstracted from the storage of the data by a Store API. This allows different implementations of a Log e.g. Billing information logs to be stored in different physical repositories, or particular log types e.g. Errors to generate events that are passed to Third Party monitoring systems.

Audit logs

Audit logs are generated at particular points in the lifetime of the iPlanet™ Trustbase Transaction Manager platform, specifically:

- Service start up events
- Service shutdown events, both normal and abnormal
- All configuration changes
- All security domain changes

These standard iPlanet™ Trustbase Transaction Manager audit log types are defined in the uk.co.jcp.tbaseimpl.log.audit.type package.

Audit Logging an Event

To cause an entry to be made in the iPlanet™ Trustbase Transaction Manager Audit Log, create a new instance of the `AuditObject` type that you wish to log (which must be a sub-type of uk.co.jcp.tbaseimpl.log.audit.AuditObject) and pass the instance to the static method `AuditLog.log(...)`.

An example of an Audit being logged is shown below

```
package com.iplanet.trustbase.app;
public class DummyService
{
    .....
    private static final String SERVICE_STARTED = "START";

    // Audit the occurrence of this event in the Audit Log
    AuditLog.log(new OperationBeginAudit( this.getClass(), SERVICE_STARTED,
    new String[] { serviceName.toString() } );
    .....
}
```

The three parameters supplied to construct all `AuditObjects` are as follows:

- `this.getClass()`
The first parameter is a class object, whose name will be used as the first part of the internationalisation bundle key for the message. This is normally the class that is logging the audit, to make it easy to relate audit messages to packages. But this may be any class object.
- `SERVICE_STARTED`
This is any string which when concatenated with the class name provides a unique bundle key to locate the text of the message for audit purposes. See below for a description of the `AuditBundle` keys and values.
- `new String[] { serviceName.toString() }`
This is an array of strings which are parameters that are combined with the `AuditBundle` message using the standard `java.text.MessageFormat` class to provide the final audit message.

As mentioned above, there needs to be an `AuditBundle` file on the class path and registered with the bundle manager (see [Defining New Audit Types](#))

section) so that the actual audit message can be looked up. For the example above there would be a file `AuditBundle.properties` created in `.../com/iplanet/trustbase/app` which contained the following line:

```
com.iplanet.trustbase.app.DummyService_START=The service {0} has begun.
```

This is the fully qualified class name provided by the `'this.getClass()'` parameter followed by a `'_'` character, followed by the `SERVICE_STARTED` string.

The message that would be associated with this audit would be `'The service {0} has begun.'` Where the `java.text.MessageFormat` class will replace `'{0}'` with entry 0 in the String array passed to the audit constructor.

Note: The API classes associated with Audit Logging are [uk.co.jcp.tbaseimpl.log.audit.AuditLog](#), [uk.co.jcp.tbaseimpl.log.audit.AuditObject](#) and a number of standard concrete Audit Classes in [uk.co.jcp.tbaseimpl.log.audit.type](#).

Defining New Audit Types

The developer may define new audit log types to allow application or domain specific messages. These must be a sub-class of the [uk.co.jcp.tbaseimpl.log.audit.AuditObject](#) class.

The most common use of these extended audit types is to log information about messages processed by the iPlanet™ Trustbase Transaction Manager Service.

Having implemented the new Audit type, the following steps must be completed in order to get the audit type to appear in the standard iPlanet™ Trustbase Transaction Manager Audit log.

- Create the new Audit class and put it into the JAR file for the service that will utilise the Audit. More details of deploying an iPlanet™ Trustbase Transaction Manager service can be found in later chapters.
- Create the relevant resource bundles for the audit messages. In order to allow simple localisation of messages into various languages, all audit type strings that appear in the iPlanet™ Trustbase Transaction Manager audit log are defined in resource bundles. These bundles are located in the same package (i.e. same directory) as the new AuditType class. The format of the name of this resource bundle must be the same as that defined by the standard Java [java.util.ResourceBundle](#) class. The format of the bundle is a series of name/value pairs. Where the name is constructed as fully-qualified-audit-type-classname_bundleKey, and the value is the String which should be placed in the AuditType field of the Audit Log.
- Create an entry in [tbase.properties](#) to register the new Audit Resource Bundle. To do this, add an entry in the section `[BundleProviderManager/Audit]` that references the fully qualified name of the new Audit Resource Bundle.
- Finally enable the new Audit Type in [tbase.properties](#) by adding an entry in the following section `[uk.co.jcp.tbaseimpl.log.present.audit.AuditLogPresentationConfigService]` for each of the new Audit Type classes defined by the service.

- Restart iPlanet™ Trustbase Transaction Manager to activate the new Audit Types, this is required to enable a new service and is therefore not an action normally associated with new Audit Types.

An example of an [AuditObject](#) implementation is shown below:

```
package uk.co.jcp.tbaseimpl.log.audit.type;
import uk.co.jcp.tbaseimpl.log.audit.*;

/** This class represents an audit object for recording that an
operation has begun
*/
public class OperationBeginAudit extends AuditObject
{
    public OperationBeginAudit ( Class auditClass , String bundleKey ,
String [] params )
    {
        super ( auditClass , bundleKey , params );
    }
    public static String getAuditTypeString()
    {
        // The bundleKey is added to this class name to determine the
// actual string to be written into the AuditType field of the log
String bundleKey = "OPERATION_BEGIN";
return
uk.co.jcp.tbaseimpl.log.audit.AuditObject.getAuditTypeString (
uk.co.jcp.tbaseimpl.log.audit.type.OperationBeginAudit.class, bundleKey
);
    }
}
```

The following file defines the Resource bundle and is located in the package hierarchy in the same place as the [OperationBeginAudit](#) class and will be called [TheNewAuditBundle_en.properties](#).

```
uk.co.jcp.tbaseimpl.log.audit.type.OperationBeginAudit_OPERATION_BEGIN =
OPERATION_BEGIN
```

The new bundle is registered in [tbase.properties](#) with the following entry, note how the locale extension ‘_en’ is not used in the bundle registration and the .properties extension is not required – this is in line with the standard Java [java.util.ResourceBundle](#) class :

```
[BundleProviderManager/Audit]
bundle=uk.co.jcp.tbaseimpl.log.audit.type.TheNewAuditBundle
```

The new Audit type is enabled by adding the following entry into [tbase.properties](#).

```
[uk.co.jcp.tbaseimpl.log.present.audit.AuditLogPresentationConfigService
]
auditlog.type.enabled=uk.co.jcp.tbaseimpl.log.audit.type.OperationBeginA
udit
```

Error handling and logging

Errors that occur in the iPlanet™ Trustbase Transaction Manager platform fall into two categories:

- Logic errors – Something about the data or processing is incorrect and the processing needs to reject a request gracefully.
- Exceptions – An unexpected condition has occurred, that breaks the processing logic

In both cases, iPlanet™ Trustbase Transaction Manager will use the Log manager to record that an error has occurred so that operational staff may perform appropriate analysis and corrective action.

Error Logging

iPlanet™ Trustbase Transaction Manager uses a single error log class that takes a severity, the class of object defining the error, and a programmer defined message, plus a set or zero or more string arguments which may be substituted into the message. The default error logging implementation, uk.co.jcp.tbaseimpl.log.error.ErrorLog, defines four constants that indicate the various severity levels:

Figure 13 Error Severity Types

Constant	Description
INFORMATION	This constant is to be used to log informational events, such as unlikely sections of code being executed that are not necessarily errors - this should be used sparingly. This has a value of 0.
WARNING	This constant is to be used for error conditions that are expected and handled, but require logging for behaviour analysis. This has a value of 1.
ERROR	This constant is to be used for serious errors which indicate that something is inherently incorrect with the system or the information it contains, but that allow processing to continue, or be retried. This has a value of 2.
FATAL	This constant is to be used for fatal errors from which processing cannot recover, these errors would result in the abandoning of processing. This has a value of 3.

Note: The API classes associated with Error Logging are uk.co.jcp.tbaseimpl.log.error.ErrorLog, uk.co.jcp.tbaseimpl.log.error.ErrorObject.

Logging an error from within iPlanet™ Trustbase Transaction Manager is as simple as calling `ErrorLog.log(...)` with an instance of an [ErrorObject](http://uk.co.jcp.tbaseimpl.log.error.ErrorObject). There are many constructors for [ErrorObject](http://uk.co.jcp.tbaseimpl.log.error.ErrorObject), but all of them have at least a string parameter that identifies the unique error code for this error.

The iPlanet™ Trustbase Transaction Manager error logging mechanism requires that every different occurrence of an error be given a code which unique throughout iPlanet™ Trustbase Transaction Manager. All of the error information for the Error Logging sub-system in iPlanet™ Trustbase Transaction Manager is contained in the following database tables:

All unique error codes have their details stored in the error_codes table described below:

Figure 14 Error_codes Table

error_codes table	
errorcode	This is the unique errorcode string that identifies the error; this must be 7 characters exactly. The normal for an error code is XXXnnnn. Where XXX is a three-letter code for the service or subsystem and nnnn is a unique number. e.g. IPH0009 is an error in the Identrus Protocol Handler.
classname	The class from which this error is logged. This places a constraint that each error code may only be used from one place.
severity	This is the severity level of the error, described previously. Constants for each of the error severities can be located in the uk.co.jcp.tbaseimpl.log.error.ErrorLog class.
message	This is the localised version of the error message that will appear in the error log. Parameters may be used in this message as described by the standard Java class java.text.MessageFormat. The values to be placed in these parameters are passed in an array of strings that one of the ErrorObject constructors allows.

The actual error log table is described below, this table is not normally viewed by the administrator directly, instead there is an Oracle view called errorview that provides a resolved view of the errors that have been logged.

Figure 15 Error Table

Error table	
errorid	This is a unique id for the error log entry; it is generated from a monotonic sequence that means that this field may be used to accurately order error messages in the order that they were logged.
errorcode	This is the errorcode of the error being logged, see the error_codes table description.
message	The final message string that is generated from the message string in the error_codes table combined with the variable parameters from the runtime system substituted in.
timestamp	This is an ORACLE DateTime field that identifies when the error was logged.
severity	This is the severity integer that is taken from the error_codes entry for this error.
classname	This is the classname that logged the error.
machineid	This is a string representing the IP address of the iPlanet™ Trustbase Transaction Manager that logged the error – this may be different in a multi-node IAS installation.
contextid	This context id field is for future expansion.

When an error is logged it is often accompanied by some free form string data which helps to store the context in which the error occurred to aid diagnosis. The most common example of such data is exception stack traces.

Figure 16 Error Support Table

error_support table	
errorid	This links this entry to an entry in the error table
datatype	This datatype is an arbitrary string identifier that categorises the data in the data field. The only value for this field defined by iPlanet™ Trustbase Transaction Manager is “STACKTRACE” which identifies the contents of the data field to be a Java Exception Stack Trace.
data	Free form string data

The tables described above encapsulate the whole data driven error logging mechanism that iPlanet™ Trustbase Transaction Manager supports.

Defining a New Error

Defining a new error is very simple; it involves making a new entry in the error_codes table. There are currently no tools to support this operation, but it may be accomplished through the use of basic SQL. The only consideration is that the error_code field must be unique – this is enforced by an ORACLE level table constraint, so if the new code is inserted without error then the code is unique.

A sample error code is inserted using the SQL defined below:

```
INSERT INTO error_codes values
(
'TST0001',
'com.iplanet.trustbase.sample.service',
'1',
'This is the only exception in the test service'
);
```

This error may then be logged using the following piece of code:

```
.....
}
catch (Exception exc)
{
    ErrorLog.log( new ErrorObject( "TST0001", exc );
}
.....
```

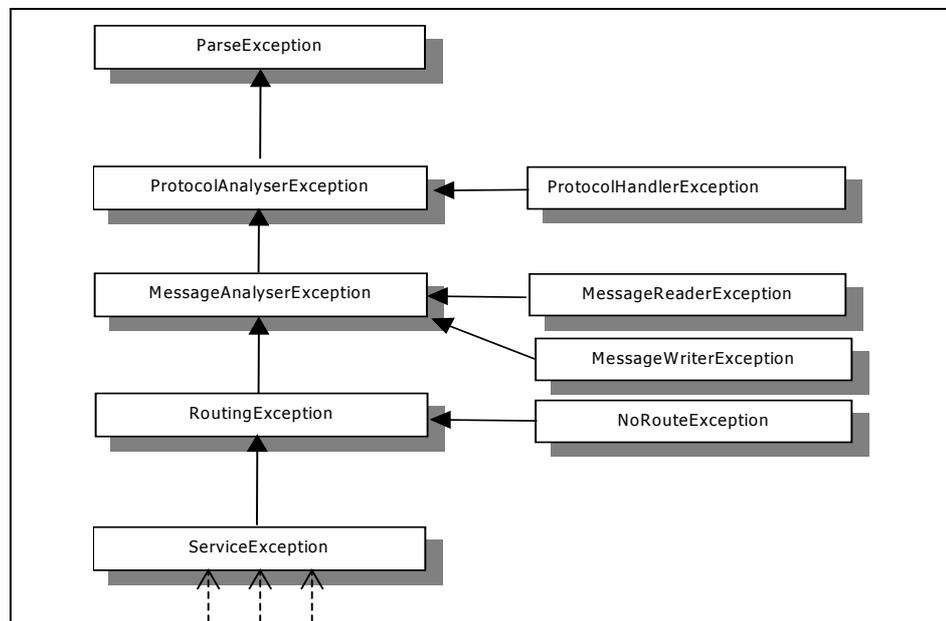
This will result in an entry being made in the error table as well as the related stack trace being logged in the error_support table.

Exception Handling

Exception handling, as with the logging of logical errors, should also log information before the appropriate exception is thrown, in order to enable users and developers of the system to analyse the situation that caused them to arise.

The iPlanet™ Trustbase Transaction Manager exception hierarchy is designed to allow the various components to throw exceptions back to their calling component without the calling component having to explicitly include handling code. This is achieved by having each successive 'level' of exceptions (as related to each successive 'level' of the code) derive from the previous level. For instance consider the example of a service throwing an exception derived from [ServiceException](#). If the service does not explicitly handle the exception it will be passed back to the Router (as the calling component), and because [ServiceException](#) is derived from [RoutingException](#) we can expect the router to be able to handle the exception gracefully.

Figure 17 iPlanet™ Trustbase Transaction Manager Exception Hierarchy



It is imperative that developers of extension to the system maintain this hierarchical approach to exception inheritance in order that the handling code built into the current iPlanet™ Trustbase Transaction Manager system will function correctly.

The hierarchical nature of the iPlanet™ Trustbase Transaction Manager exceptions allows the exception handling to be done differently depending upon where the exception occurs. The components and the associated actions performed are shown in the table below:

Figure 18 Exceptions

Component	Base Exception	Actions
Servlet	TbaseRuntimeException	Catches TbaseRuntimeExceptions and generates the appropriate HTTP error
Protocol Analyser	ProtocolAnalyserException	may throw ParseException which will result in an HTTP error in the servlet
Message Analyser	MessageAnalyserException	Generated if a RemoteException has occurred, will result in an HTTP error in the servlet.
Message Readers	MessageReaderException	Is caught in the MessageAnalyser if it contains an embedded Message, then that Message is processed, otherwise the exception is propagated through to the servlet.
Message Writers	MessageWriterException	Is propagated back through the MessageAnalyser to the servlet that returns an HTTP error.
Router	RoutingException	Is propagated back through the MessageAnalyser to the servlet that returns an HTTP error
Services	ServiceException	May or may not contain a message, and is propagated back through the router to the MessageAnalyser where it is handled in the same fashion as the MessageReaderException

Identrus logging

The Identrus Transaction Coordinator specifications identify two specific logging actions, these being:

- Logging of all messages sent and received by the Transaction Coordinator (Raw logging)
- Generation of data for billing purposes

Overview

The iPlanet™ Trustbase Transaction Manager fulfils both of these requirements as a default action of processing an Identrus message. The data is stored within the RDBMS specified at installation time, and the tables are available for developers via standard JDBC to provide services that use this information. The following sections define the tables stored in the RDBMS and identify the relationships between each table. The iPlanet™ Trustbase Transaction Manager will utilise all of the tables described below for all Identrus messages; there should be no requirement for a developer to write to any of these tables.

Data definitions

Connection information

The SSL proxy and the SMTP mail listener both log data about the connections made through them. The table below provides the column definitions for the SSL Proxy:

Figure 19 SSL Connection Table

ssl_connection table	
ConnectionId	Unique connection identifier
ClientCertIssuerDN	The connecting clients certificate Issuer DN
ClientCertSerialNumber	The connecting clients certificate serial number
CipherSuite	The cipher suite used for the SSL session
ConnectTime	The time at which the connection was made, this is an ORACLE DateTime field
TimeStampType	The type of timestamp
ConnectIPAddr	The connecting client's IP address
ConnectionFailed	Integer value indicating if the connection failed – a value of 1 indicates a failure.
ConnectionFailedReason	If a failure occurred, what was the SSL error code

The tables below provide the column definitions for the SMTP/SMIME connection logs: Data in the smime_transport table is extracted from the SMIME v2 signature body part on the message.

Figure 20 SMIME Transport Table

smime_transport table	
ConnectionId	Provides a link back to the smtp_message table
peer_issuer_dn	The issuer_dn of the certificate that was used to verify the message
peer_cert_serial_number	The serial number of the certificate used to verify the message.
message_protection	The type of protection used to secure the message
time_stamp_type	The type of timestamp LOCAL or NETWORK
time_stamp	The time at which the entry was made

Figure 21 SMTP Connection Table

smtp_connection table	
stream_id	Provides a link back to the smtp_message table
peer_ip_addr	The ip address of the submitting SMTP agent
timestamptype	The type of timestamp LOCAL or NETWORK
timestamp	The time at which the entry was made

Figure 22 SMTP Message Table

smtp_message table	
stream_id	A unique id for the smime_transport
connection_id	A unique id for the smtp connection
recipients	The recipients of this message
sender	The sender of this message
timestamptype	The type of timestamp LOCAL or NETWORK
message_valid	Is the message valid? 1 indicates it is valid
message_invalid_reason	The reason for the invalidity of the message
timestamp	The date and time at which the entry was made

The ssl_connection and smtp_message tables both have connection_id fields that are passed to the iPlanet™ Trustbase Transaction Manager running in the application server. This connection_id is stored within the Raw Log table allowing queries that link the originator information with the actual requests made.

Figure 23 OCSP table

ocsp_data table	
ocspid	A unique identifier for the record
type	OCSPREQUEST or OCSPRESRESPONSE
message	A text summary of the contents of the request or response
machine	The URL to which the request was submitted to or the response was received from
timestamp	The date and time that the entry was made
data	Base64 encoding of the request or response

Raw log tables

The default presentation handlers for Identrus messages record the following data for each message that is sent or received:

Figure 24 Raw log Table

raw_data table	
Sessionid	The id of the raw log session that wrote this record
Logconnectionid	The id of the connection within the session
Recordid	The id of the record within the connection
msggrpId	The Identrus MsgGrpId from the NIB of the message
msgId	The Identrus MsgId from the NIB of the message
doctype	The DOCTYPE of the message. e.g. CSCRequest, PingRequest etc.
recordmarker	A unique monotonically increasing identifier
connectionid	The connection id to link this record to the SSL or SMIME connection logs.
protocoltype	The protocol over which the message arrived. e.g. HTTP or SMTP
input	Was this message inbound to the iPlanet™ Trustbase Transaction Manager or outbound? A value of 1 indicates it was incoming.
timestamp	An integer which represents the UNIX time at which the record was logged.
rawdata	The Identrus Message XML, without the CertBundle fields. The certificates from the bundle are logged separately in the cert_data table.
digestofrecord	A SHA-1 digest of this record.
signeddigestofcalculation	An RSA signature of this record and data from the previous record.
servercertissuerdn	The issuer DN of the certificate used to verify the signature
servercertserialnumber	The serial number of the certificate used to verify the signature.

In order to reduce the volume of data logged with each Identrus message the certificates contained with the message header are stripped out and stored in a certificate table. If the iPlanet™ Trustbase Transaction Manager has already logged a particular certificate in the table it will not be logged again. The information stored within the table is:

Figure 25 Certdata Table

cert_data table	
IssuerDN	The issuer distinguished name of the certificate, RFC 2253 format string.
SerialNumber	The serial number of the certificate
CertData	The Base64 certificate data.

This data is designed to be tamper evident, and services should under no circumstances modify data within the Raw Log or Tamper tables. The tamper checking is achieved by producing a continuous hash that is stored with each record, and the current hash is stored within a signed record within a separate tamper table. The Tamper table fields are not described here, see the Installation and Configuration Guide for information on how to check the tamper status of records in the raw log.

Billing records

Billing records are a sub-set of the information within the raw message log that provides sufficient information to determine who made each transaction. These tables are designed for used by third party tools that generate the actual Bill for the customer. The definitions for the bill table columns are as follows:

Figure 26 Bill data Table

bill_data table	
RawRecordId	This will be the RawRecordId of the associated raw log table record.
SubjectDN	This will be the originator distinguished name extracted from the mandatory Identrus level 1 message signature. This will determine who should be billed.
IssuerDN	This will be the issuer distinguished name extracted from the mandatory Identrus level 1 message signature. This is to enable the identification of the exact key used to sign this message – in conjunction with the serial number field below.
SerialNumber	This will be the originator certificate serial number that may be used to identify the exact key used to sign the message – in conjunction with the issuer distinguished name.

Building Identrus solutions

The iPlanet™ Trustbase Transaction Manager platform and associated development tools provide a means of developing Identrus applications starting from the base DTD through to installing the service for use in a run-time environment.

The following sections walk through the generation of an Identrus application called Ping. This application is designed to highlight the following concepts:

- Overall development process
- Use of iPlanet™ Trustbase Transaction Manager developer tools
- Using configuration objects
- Using the log manager

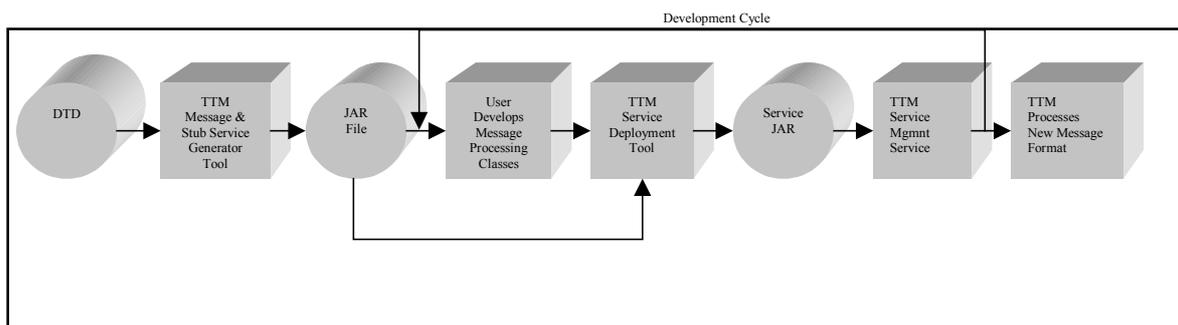
Methodology

Development process

The development of an application requires a base input of an Identrus compliant DTD. This is passed through a class generator that produces Java Classes for all of the elements in the DTD along with a set of stub rules and a service definition within a JAR file. The developer is then at liberty to extend this base data by implementing the appropriate business logic in the form of a set of Java classes called by the service stub that perform the appropriate actions.

The set of generated files and associated developer files are then packaged into a single JAR and deployed into an iPlanet™ Trustbase Transaction Manager directory structure. The iPlanet™ Trustbase Transaction Manager (when re-started) will locate a package descriptor within the JAR file and offer this new service for activation within the running environment. The overall process is outlined in the figure below.

Figure 27 Development process



Class generation

The iPlanet™ Trustbase Transaction Manager class generation tool is designed to produce a set of classes that may be used at two points in the message processing pipeline. These are:

Presentation - The iPlanet™ Trustbase Transaction Manager XML parser uses the classes to validate incoming XML against the DTD.

Services - Business logic that accesses elements of the message

In order to access the class generation tool, a JDK 1.2.x or greater must be installed; a JRE does not contain the correct compilers to build an iPlanet™ Trustbase Transaction Manager service. Also the classpath must be correctly set, executing the following on your iPlanet™ Trustbase Transaction Manager installation can do this:

```
cd ...../Trustbase/TTM/Scripts
. ./setcp
```

This sets the CLASSPATH environment variable to the correct value for the iPlanet™ Trustbase Transaction Manager tools.

The class generation tool is available as a command line option and is invoked using the following:

```
java com.iplanet.trustbase.app.classgen.ClassGen [-options] <PublicId>
<DTD file name>
```

Where options include:

```
-help           displays this help screen
-d <dir>       specify output directory
-o <jar name>   specify output jar file
-v            verbose output
-r            recursively generate classes for all included DTDs
-root <elementName> specify for each root element of the DTD
-stub <className> fully qualified class name of the stub service to
be created
```

And PublicID is the name of the service stub that will be generated. The Public ID is a mandatory requirement for all Identrus DTD's, and is in the standard form shown below:

```
-//IDENTRUS//service_name//EN
```

Where service name represents the Identrus service declared by the DTD. The Public ID is case sensitive and should be carefully entered as it forms part of the package name for the generated classes.

The tool will only generate classes for DTD elements that do not already have classes available on the classpath. Therefore, the core Identrus classes are not regenerated with each new Identrus service because they are already available in Trustbase jars.

A '-root' entry *must* be made for each of the elements in the DTD that can be the root XML element in a message. This is necessary to allow the automatic processing of the Identrus Message by the default presentation layer in the iPlanet™ Trustbase Transaction Manager.

A '-stub' option must be used if the tool is required to generate a stub iPlanet™ Trustbase Transaction Manager service for the new message types. The Java source file for this stub class is placed in the current directory when the tool completes. Do NOT re-package the stub class once generation has completed – the fully qualified name of the stub class is recorded in the [tbasesvc.properties](#) file in the generated jar file.

The Class generator will attempt to recursively load the referenced DTD(s). The references may be:

- Unqualified - The DTD must be located in the directory in which the class generator is run in
- Local - The file:/// qualifier is used
- Fully qualified - The HTTP://www qualifier is used

An invocation of the Classgen tool will produce the following output:

- A .zip containing

- A set of generated classes, one for each entity in the DTD
- A set of DTD's that represent the input actually used by the XML parser for generating the classes
- A [tbasesvc.properties](#) file used by iPlanet™ Trustbase Transaction Manager to register the service
- A service stub in the current directory – only if the '-stub' option is used

Each of the generate classes implements the [TbaseElement](#) interface. This interface provides two specific roles:

- Parser - Allows the iPlanet™ Trustbase Transaction Manager parsing mechanisms to validate incoming XML messages
- Service - Provides a standard means of getting data from incoming messages and constructing valid responses without explicitly producing the XML.

There is a sub-class of [TbaseElement](#) called [TbaseIdentifiedElement](#) that represents elements that have an attribute named 'id' with a type of 'ID'. This allows a type-safe method of checking ID compliance when generating or validating XML DSIG signatures.

The service role of the [TbaseElement](#) provides the following function:

- Constructing the XML document hierarchy
- Reading and writing the element attributes by name
- Validating the object representation of the XML document
- Generating the String representation of the XML document described by the object hierarchy

All elements that are identified by a '-root' option at class generation time, will end up extending the base class [com.iplanet.trustbase.identrus.message.IdentrusMessage](#). This enforces that the message has a valid structure (as defined by Identrus Core Network Messaging Definition) and the framework can carry out that mandatory processing of the message. An example of using this interface is shown in the example section at the end of this document.

Note: The API classes for [com.iplanet.trustbase.identrus.message.IdentrusMessage](#), [com.iplanet.trustbase.xml.message.TbaseElement](#) and [com.iplanet.trustbase.xml.message.TbaseIdentifiedElement](#) provide more information on their function.

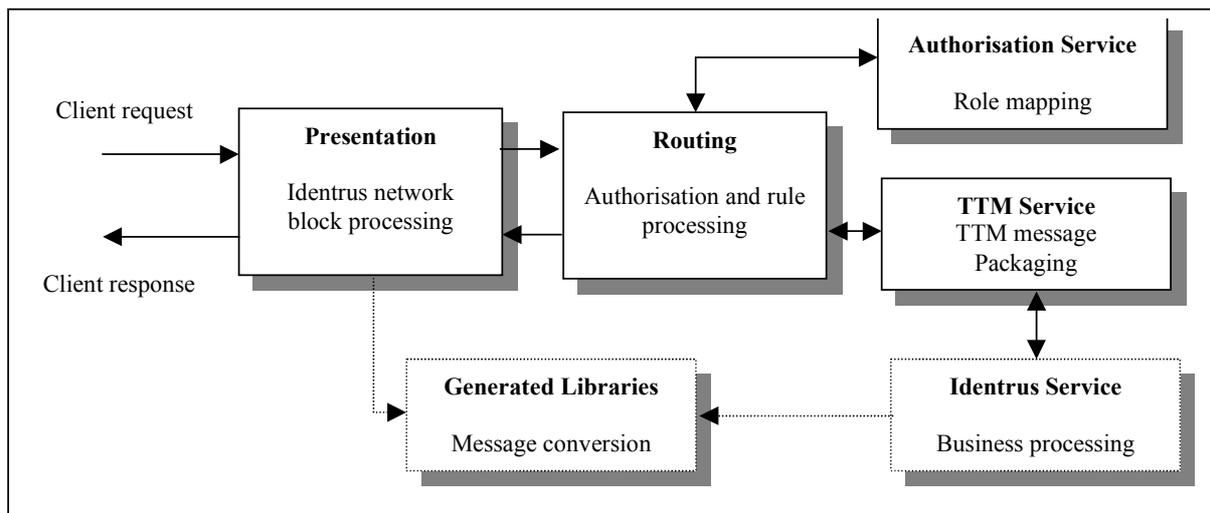
Service development

Development of a service is the act of completing the service stub into a functional business component.

Every service must be stateless so that it may be replicated by the platform. This means that a service must be capable of processing messages from different transactions and different users without holding information about a previous Session State Transaction. If a service is required to hold information about a session or transaction state, this must be persisted in a database, as the subsequent messages in a transaction are not guaranteed to be passed to the same host machine within a cluster.

All Identrus network message processing is performed by the iPlanet™ Trustbase Transaction Manager platform prior to and after the invocation of the [ProcessIdentrusMessage](#) method in the service stub class generated by the Class Generation tool. The figure below shows the message processing path with a Developer written Identrus Service using the generated message libraries and proxied by the iPlanet™ Trustbase Transaction Manager message process stub.

Figure 28 Message path processing



This means that for most Identrus services, the developer only needs to modify the generated stub class to implement the business processing of the incoming message and generate the core content of the outbound message.

The business processing does not need to do any of the mandatory message processing such as mandatory signature verification or generation on the message. These core processing tasks are completed by the iPlanet™ Trustbase Transaction Manager presentation layers.

To develop the service, include the generated jar file on the classpath during development, along with the core iPlanet™ Trustbase Transaction Manager libraries.

Service Building

Having successfully generated the messaging classes and developed the message processing logic extensions to the stub service, the next task is to build a iPlanet™ Trustbase Transaction Manager service jar.

Building the iPlanet™ Trustbase Transaction Manager service jar is a simple matter of the following steps:

- Using the standard JDK 'jar' tool, unpack the generated jar file into a directory.

```
mkdir DumpDirectory
cd DumpDirectory
jar -xvf ../jarfile.jar
cd ..
```

- Beneath the 'DumpDirectory' copy in all of the classes which are required for the processing of the message. Ensure that the correct package structure is maintained for all these classes. n.b. Remember that the service stub class *must not* have been re-packaged since the class generation created it.

- Using the standard JDK 'jar' tool, repack the archive

```
jar -cvf jarfile.jar -C DumpDirectory DumpDirectory/*
```

- The jar file is now ready for deployment.

Service Deployment

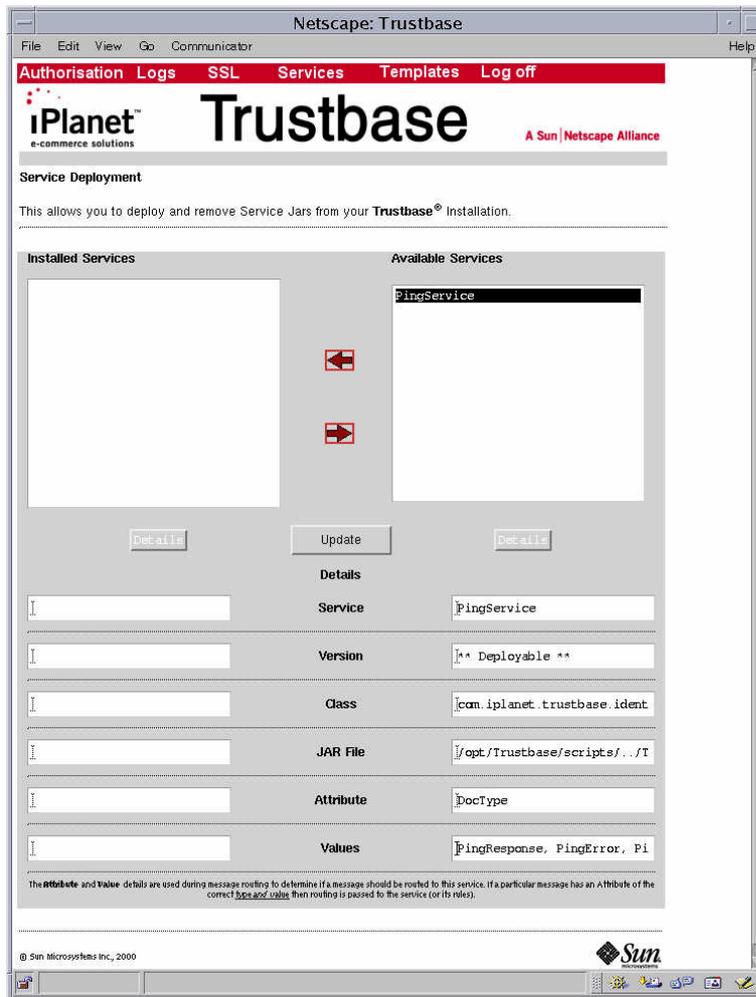
Once the service jar file has been built it can be deployed into iPlanet™ Trustbase Transaction Manager. This is done in three stages:

- Copy the jar file into the directory:
...../Trustbase/TTM/current/deploy
- Use the administrator console to deploy the service into the running iPlanet™ Trustbase Transaction Manager
- Re-start iPlanet™ Trustbase Transaction Manager to activate the service

To deploy the service, logon to iPlanet™ Trustbase Transaction Manager as 'administrator' and select the 'Deployment' option from the services menu. If you have copied the service jar into the 'deploy' directory then the new service should appear on the left hand side of the deployment screen.

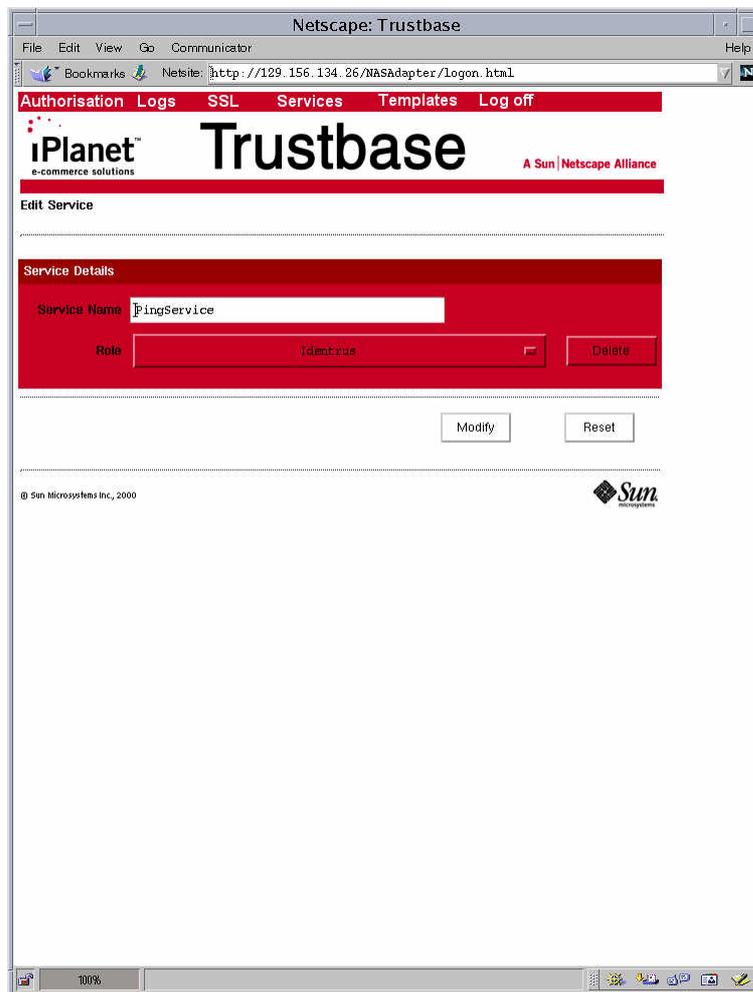
- To deploy the service select  followed by 

Figure 29 Deploying PingService within iPlanet Trustbase Transaction Manager



Services may also require authorisation. Select <Authorisation><Add Service>. The service itself will need to be assigned a role, as illustrated below:

Figure 30 Assigning a role to a service



Note: if no role is available you may have to define a role. Consult the Configuration and Installation Guide for more details about this.

Ping Example

The following sections use the Ping DTD as defined in the Identrus messaging specification to walk through the steps in the development cycle as described in the previous chapter. Developing a service prior to deploying it involves a number of steps.

- 1. Create your DTD definitions that specify the syntactic structure of the messages you wish to send round the system.
- 2. Use Classgen `com.iplanet.trustbase.app.classgen.ClassGen` to generate your java classes from your DTD definitions.
- 3. Write the Java code for the service using the Identrus API that assists the Identrus processing and validating of messages, certificates, keys and digital signatures.
- 4. Use the JDK JAR tool to build the final iPlanet™ Trustbase Transaction Manager service jar file.
- 5. Deploy the service within iPlanet™ Trustbase Transaction Manager by selecting the relevant configuration options as described below in this guide.
- 6. Finally, once it has been deployed within iPlanet™ Trustbase Transaction Manager itself you can run your service.

Create DTD Definitions

- This example uses the 'Ping' DTD defined in the Identrus Messaging specification and included below for completeness. Thus, create your DTD definitions which can be found in ping.dtd. (See IT-TCMPD, the Identrus TC Messaging Specification for the structure and definition of ping.dtd)

```
<!--PUBLIC ID for this DTD is: "-//IDENTRUS//PING DTD//EN"-->

<!ENTITY % CoreNetwork.dtd PUBLIC "-//IDENTRUS//CORE NETWORK
INFRASTRUCTURE DTD//EN" "corenetworkinfrastructure.dtd">
%CoreNetwork.dtd;

<!ELEMENT PingRequest (NIB, Signature, CertBundle, PingData)>
<!ATTLIST PingRequest
  id ID #REQUIRED
>
<!ELEMENT PingResponse (NIB, Signature, CertBundle, PingData)>
<!ATTLIST PingResponse
  id ID #REQUIRED
>
<!ELEMENT PingError (NIB, Signature, CertBundle, ErrorInfo)>
<!ATTLIST PingError
  id ID #REQUIRED
>
<!ELEMENT PingData (#PCDATA)>
<!ATTLIST PingData
  id ID #REQUIRED
>
<!ELEMENT ErrorInfo (VendorData*)>
<!ATTLIST ErrorInfo
  id ID #REQUIRED
  errorCode NMTOKEN #REQUIRED
>
<!ELEMENT VendorData (#PCDATA)>
<!ATTLIST VendorData
  id ID #REQUIRED
  dataType CDATA #REQUIRED
>
```

- These link to other dtd files that are supplied with iPlanet™ Trustbase Transaction Manager and are listed below:

```
XMLDSIG.dtd
corenetworkinfrastructure.dtd
Foundation.dtd
```

Note All dtd's can be found in <install-dir>/Trustbase/TTM/current/apidocs

Prepare to execute the ClassGen tool by completing the following steps:

- Change directory to ../Trustbase/TTM/Scripts
- Set up the classpath ready for the ClassGen tool to be executed. This is done by sourcing the setcp script. E.g. “. ./setcp” in the standard shell.

- Make a new sub-directory called ping and make this the current directory
- Copy ping.dtd and each of the 3 core Identrus DTD's to the new directory. It is essential that the names of the DTD files are correct, they are case sensitive and their exact names can be found by looking inside the DTD files. The last attribute of an <!ENTITY> tag defines the local file name for the included DTD.
- Establish which options must be passed to the ClassGen tool, the table below shows each of the options for this example.

Figure 31 Classgen Options

ClassGen Options for Ping DTD	
-d .	Create the files in the current directory
-o ping.jar	Create the final jar called ping.jar
-r	Recurse through all sub-dtd's and generate any classes that do not already exist on the classpath.
-root PingRequest	The PingRequest element in the DTD is an eligible root element for an XML document.
-root PingResponse	The PingResponse element in the DTD is an eligible root element for an XML document.
-root PingError	The PingError element in the DTD is an eligible root element for an XML document.
-stub com.iplanet.trustbase.sample.PingService	Create the stub Identrus Service source code in the output directory; package the file in com.iplanet.trustbase.sample.
"-//IDENTRUS//PING DTD//EN"	This is the public identifier for the Ping DTD
Ping.dtd	The DTD from which class generation should begin.

This results in the following command being issued:

```
java com.iplanet.trustbase.app.classgen.ClassGen -d . -o ping.jar -r
-root PingRequest -root PingResponse -root PingError -stub
com.iplanet.trustbase.sample.PingService "-//IDENTRUS//PING DTD//EN"
ping.dtd
```

The result of executing this command is the following:

- Many warnings on the screen. These warnings state that classes have not been generated for certain elements in the DTD's referenced from the ping.dtd. This is because these classes already exist on the classpath of iPlanet™ Trustbase Transaction Manager and therefore do not need to be regenerated.
- A file called ping.jar in the current directory. This contains the compiled message classes for the ping DTD.

- A sub-directory called src that contains the source code of all of the generated class files.
- A file called PingService.java in the current directory. This is the stub Identrus Service that needs to be updated with message processing code in the next stage.

With the JAR file and the stub service generated, you can now move to the next stage of the development of a service.

API

Before writing your Java code you'll need to know some of the core API that is supplied with iPlanet™ Trustbase Transaction Manager for handling Identrus Messages:

```
com.iplanet.trustbase.identrus
```

This package allows access to Identrus message Specification

```
com.iplanet.trustbase.identrus.message
```

This package provides the necessary routines to allow Identrus message processing

```
com.iplanet.trustbase.identrus.security
```

This package processes certificates and keys

```
com.iplanet.trustbase.util.tree
```

Searches tree

```
com.iplanet.trustbase.xml.dsig
```

Generates and validates XML Digital signatures, this package is not required if the only signatures on the messages are the mandatory level 1 signatures. These mandatory signatures are checked/generated automatically by iPlanet™ Trustbase Transaction Manager.

```
uk.co.jcp.tbase.config
```

This package interfaces with configuration objects

PingService Source Code

[PingService.java](#) - the generated service stub is an implementation of the abstract [IdentrusService](#) in which the [ProcessIdentrusMessage](#) method parameter is the message received by the TTM platform.

By the time the message arrives at this [ProcessIdentrusMessage](#) method, it has already passed mandatory signature validation checks and been raw logged.

The method must create a response message to the incoming request. Using the supplied Identrus API's this is simple.

```

package com.iplanet.trustbase.sample;
import com.iplanet.trustbase.identrus.message.IdentrusMessage;
import uk.co.jcp.tbase.service.ServiceException;
import com.iplanet.trustbase.generated.IDENTRUS.PING_DTD.*;
import
com.iplanet.trustbase.generated.IDENTRUS.CORE_NETWORK_INFRASTRUCTURE_DTD
.*;
import com.iplanet.trustbase.identrus.IdentrusService;
import com.iplanet.trustbase.identrus.message.*;
import uk.co.jcp.tbaseimpl.log.error.*;
import uk.co.jcp.tbase.service.*;
/** Stub Identrus service implementation. */
public class PingService extends IdentrusService
{
    public IdentrusMessage processIdentrusMessage( IdentrusMessage
message )
    {
        if (message instanceof PingRequest)
        {
            // Handle PingRequest
            PingResponse pr = new PingResponse();
            PingData pd = new PingData();
            pd.setPCDATA( "The ping was successful" );
            pr.setPingData( pd );
            try
            {
                NIBAccessor niba = NIBAccessor.getInstance(
message.getNetworkInfoBlk(), "2" );
                pr.setNetworkInfoBlk( niba );
            }
            catch (NIBAccessorException nie)
            {
                ErrorLog.log( new ErrorObject( "IDT0052", nie ) );
            }
            return pr;
        }
        if (message instanceof PingResponse)
        {
            // Handle PingResponse
        }
        if (message instanceof PingError)
        {
            // Handle PingError
        }
        return message;
    }
}

```

Having written this simple service it may be compiled ready to be built into a Service JAR.

Creating the Idenrus Service JAR

The final step before deploying the service into iPlanet™ Trustbase Transaction Manager is to create the consolidated JAR file that contains all of the generated classes as well as the [PingService](#) class that has been developed by hand.

Building the iPlanet™ Trustbase Transaction Manager service jar is a simple matter of the following steps:

- Using the standard JDK 'jar' tool, unpack the generated jar file into a directory.

```
mkdir DumpDirectory
cd DumpDirectory
jar -xvf ../ping.jar
cd ..
```

- Beneath the 'DumpDirectory' copy in all of the classes which are required for the processing of the message. Ensure that the correct package structure is maintained for all these classes. n.b. Remember that the service stub class must not have been re-packaged since the class generation created it.

- Using the standard JDK 'jar' tool, repack the archive

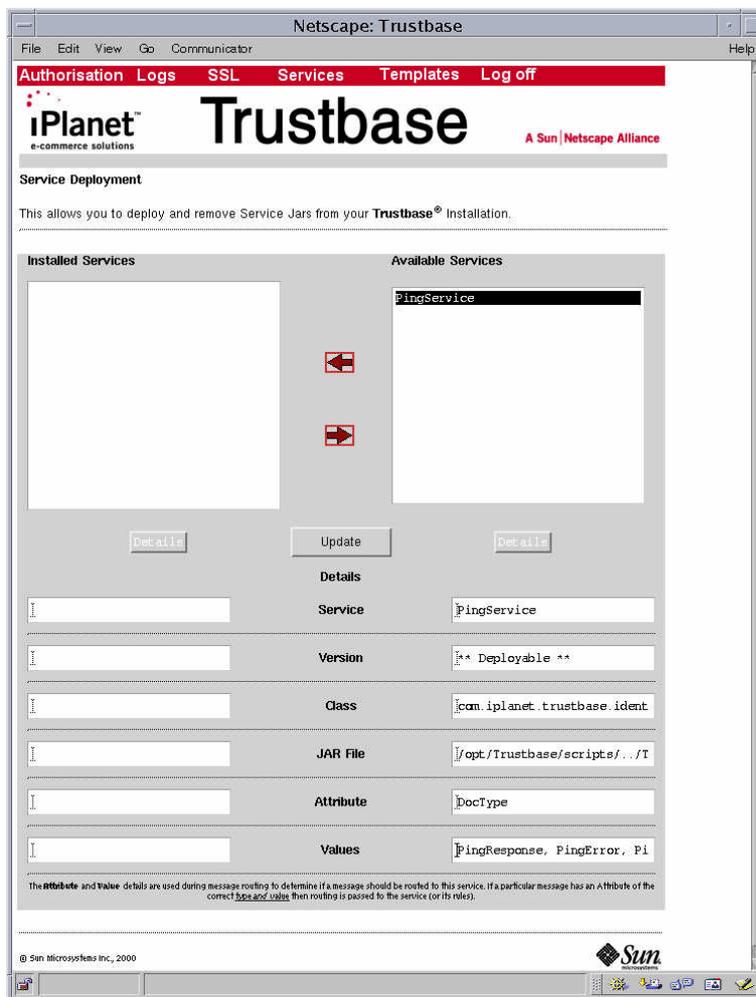
```
jar -cvf ping.jar -C DumpDirectory DumpDirectory/*
```

- The jar file is now ready for deployment.

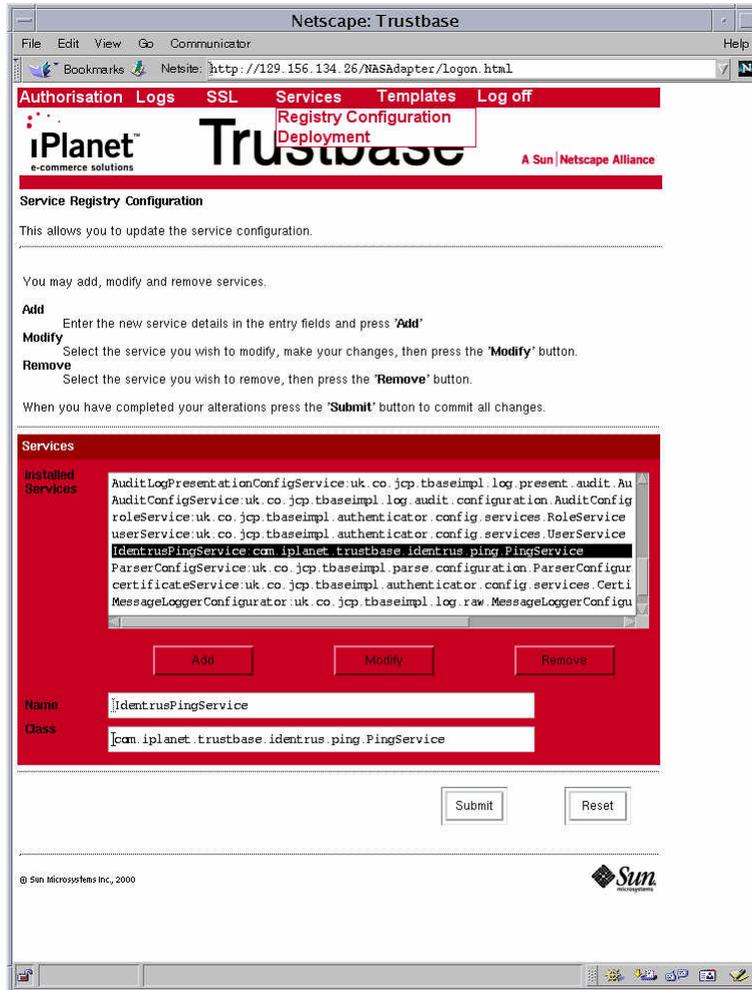
Deploying ping.jar within iPlanet Trustbase Transaction Manager

- Place ping.jar in the deploy directory.
...../Trustbase/TTM/current/deploy
- Logon to iPlanet™ Trustbase Transaction Manager as Administrator
- Deploy PingService within iPlanet™ Trustbase Transaction Manager. Select <Services> followed by <Deployment>.
- Since Ping.jar has been placed in the deploy directory automatically by the builder, iPlanet™ Trustbase Transaction Manager picks up all the relevant information and places this in the screen headed "Available services" as illustrated below.
- To deploy the service select  followed by 

Figure 32 Deploying PingService within iPlanet Trustbase Transaction Manager

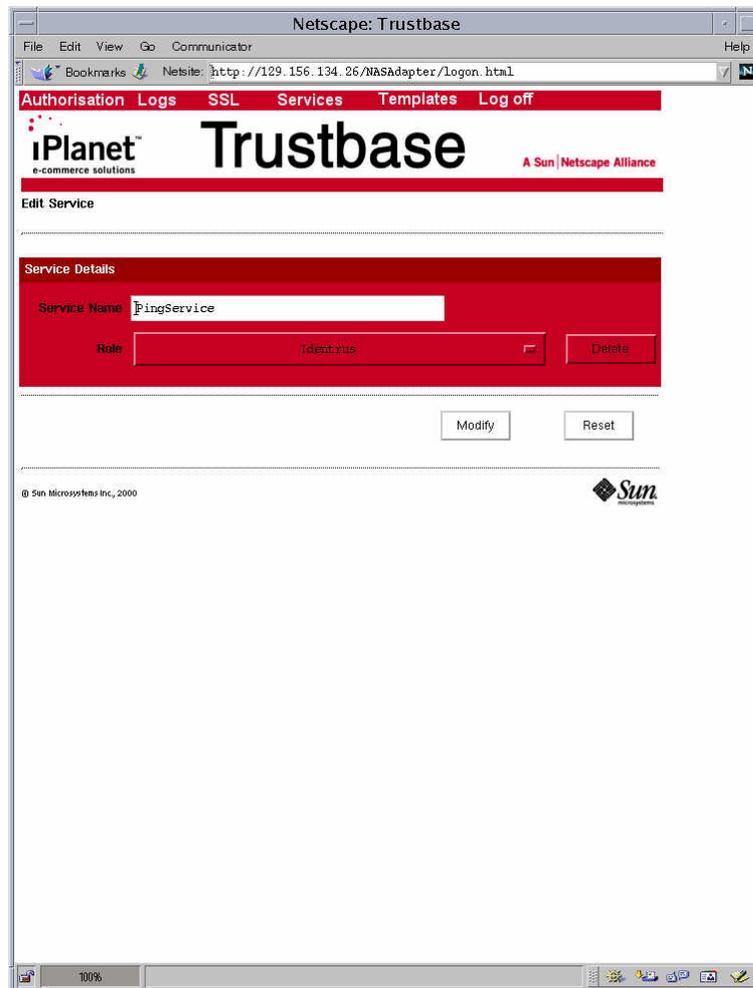


- Services require registration. This is done automatically when the jar file is built from the builder tool. [PingService](#) can in fact be seen from the Configuration Console by selecting <Services> <Registry Configuration> as illustrated below:



- Services may also require authorisation. Select <Authorisation><Add Service>. The service itself will need to be assigned a role, as illustrated below:

Figure 33 Assigning a role to a service



Note if no role is available you may have to define a role. Consult the Configuration and Installation Guide for more details about this.

References and Glossary

Software Platform

Solaris 8 and JDK

<http://www.sun.com/software/solaris/cover/sol8.html>

Java

<http://www.javasoft.com>

iPlanet Application Server 4.1

http://www.iplanet.com/products/infrastructure/app_servers/index.html

iPlanet Web Server 6.0

http://www.iplanet.com/products/infrastructure/web_servers/index.html

Oracle 8i

<http://www.oracle.com>

Hardware Security nCipher KeySafe 1.0 and CAFast

<http://www.ncipher.com>

Transport Protocols

HTTP

HTTP/1.0 or 1.1 protocol:

<http://www.w3.org/Protocols/rfc1945/rfc1945.txt>

<http://www.ietf.org/rfc/rfc1945.txt>

SMTP RFC821

<ftp://ftp.isi.edu/in-notes/rfc821.txt> <http://www.imc.org/ietf-smtp/>

Security Related Protocols

S/MIME Version 2 Message Specification

<ftp://ftp.isi.edu/in-notes/rfc2311.txt>

<http://www.imc.org/ietf-smime>

<http://www.ietf.org/rfc/rfc2311.txt>

DOMHASH

<http://www.ietf.org/rfc/rfc2803.txt>

OCSP

<http://www.ietf.org/rfc/rfc2560.txt>

Certificate requests and responses

PKCS10 requests RFC2314 can be found in <http://www.ietf.org/rfc.html>

PKCS7 responses RFC2315 can be found in <http://www.ietf.org/rfc.html>

Trading protocols

Identrus

<http://www.identrus.com>

Transaction Coordinator requirements (IT-TCFUNC)

Core messaging specification (IT-TCMPD)

Certificate Status Check Messaging specification (IT-TCCSC)

Message Protocols

DOM

<http://www.w3.org/TR/REC-DOM-Level-1/>

DTD

<http://www.w3.org/XML/1998/06/xmlspec-v20.dtd>

XML

<http://www.w3.org/TR/REC-xml>

XML Syntax Processing specification

<http://www.w3.org/TR/xmlsig-core>

HTML

HTML 3.2 as specified in <http://www.w3.org/TR/REC-html32.html>

Security related terms

Application protocol. An application protocol is a protocol that normally layers directly on top of the transport layer (e.g., TCP/IP). Examples include HTTP, TELNET, FTP, and SMTP.

Asymmetric cipher. See public key cryptography.

ASN.1. Abstract Syntax Notation One.

Authentication. Authentication is the ability of one entity to determine the identity of another entity.

base64. A representation of characters in digital format using a 65 character subset of U.S. ASCII.

BBS. A random number generating algorithm.

BER. Basic encoding Rules used with X509.

Block cipher. A block cipher is an algorithm that operates on plaintext in groups of bits, called blocks. 64 bits is a typical block size.

Bulk cipher. A symmetric encryption algorithm used to encrypt large quantities of data.

Cipher Block Chaining Mode (CBC). CBC is a mode in which every plaintext block encrypted with the block cipher is first eXclusive-OR-ed with the previous ciphertext block (or, in the case of the first block, with the initialisation vector).

Certificate. As part of the X.509 protocol (a.k.a. ISO Authentication framework), certificates are assigned by a trusted Certificate Authority and provide verification of a party's identity and may also supply its public key.

Client. The application entity that initiates a connection to a server.

Client write key. The key used to encrypt data written by the client.

Client write MAC secret The secret data used to authenticate data written by the client.

Connection. A connection is a transport (in the OSI layering model definition) that provides a suitable type of service. For SSL, such connections are peer to peer relationships. The connections are transient. Every connection is associated with one session.

CRL Certificate Revocation List. A certificate that is not valid but still within its expiry date.

Data Encryption Standard (DES). DES is a very widely used symmetric encryption algorithm. DES is a block cipher.

3DES. Similar to DES.

DER. Distinguished Encoding rules used in X509.

DH. A public-key cryptographic algorithm for encrypting and decrypting data.

Digital Signature Standard (DSS). A standard for digital signing, including the Digital Signing Algorithm, approved by the National Institute of Standards and Technology, defined in NIST FIPS PUB 186, "Digital Signature Standard," published May, 1994 by the U.S. Dept. of Commerce.

Digital signatures. Digital signatures utilise public key cryptography and one-way hash functions to produce a signature of the data that can be authenticated, and is difficult to forge or repudiate.

DSA. Digital Signature Algorithm.

Handshake. An initial negotiation between client and server that establishes the parameters of their transactions.

Initialization Vector (IV). When a block cipher is used in CBC mode, the initialisation vector is eXclusive-OR-ed with the first plaintext block prior to encryption.

IDEA. A 64-bit block cipher designed by Xuejia Lai and James Massey.

Message Authentication Code (MAC). A Message Authentication Code is a one-way hash computed from a message and some secret data. Its purpose is to detect if the message has been altered.

Master secret. Secure secret data used for generating encryption keys, MAC secrets, and IVs.

MD5. MD5 is a secure hashing function that converts an arbitrarily long data stream into a digest of fixed size.

Message digest. A digest algorithm converts data of any size, via a one-way hashing function, into a small fixed size unique representation. Message digests are used extensively in the generation of digital signatures and integrity checking of data.

MIME. MultiPURPOSE Internet Mail Extension

PBE. Password based encryption

PEM. Privacy enhanced mail

Public Key Infrastructure (PKI). Defines protocols to support online interaction.

Public key cryptography. A class of cryptographic techniques employing two-key ciphers. Messages encrypted with the public key can only be decrypted with the associated private key. Conversely, messages signed with the private key can be verified with the public key.

One-way hash function . A one-way transformation that converts an arbitrary amount of data into a fixed-length hash. It is computationally hard to reverse the transformation or to find collisions. MD5 and SHA are examples of one-way hash functions.

OSI. Open Systems Inter-Connection.

RC2, RC4. Proprietary bulk ciphers from RSA Data Security, Inc. RC2 is block cipher and RC4 is a stream cipher.

RFC. A series of authoritative discussion documents. Requests for Comments.

RSA. A very widely used public-key algorithm that can be used for either encryption or digital signing.

Salt. Non-secret random data used to make export encryption keys resist pre-computation attacks.

Server. The server is the application entity that responds to requests for connections from clients. The server is passive, waiting for requests from clients.

Session. A SSL session is an association between a client and a server. Sessions are created by the handshake protocol. Sessions define a set of cryptographic security parameters that can be shared among multiple connections. Sessions are used to avoid the expensive negotiation of new security parameters for each connection.

Session identifier. A session identifier is a value generated by a server that identifies a particular session.

Server write key. The key used to encrypt data written by the server.

Server write MAC secret. The secret data used to authenticate data written by the server.

SHA. The Secure Hash Algorithm is defined in FIPS PUB 180-1. It produces a 20-byte output.

SSL. Secure sockets layer

Stream cipher. An encryption algorithm that converts a key into a cryptographically-strong keystream that is then eXclusive-OR-ed with the plaintext.

Symmetric cipher. See Bulk Cypher.

TSL. Transport security layer

X690. The ASN.1 specification

X509. An authentication framework based on ASN.1 BER and DER and base64.

Java Related terms

Abstract class. A class that contains one or more abstract methods, and therefore can never be instantiated. Abstract classes are defined so that other classes can extend them and make them concrete by implementing the abstract methods.

Abstract method. A method that has no implementation.

API Application Programming Interface. The specification of how a programmer writing an application accesses the behaviour and state of classes and objects.

Applet A program written in Java to run within a Java-compatible web browser, such as HotJava(TM) or Netscape Navigator(TM).

Atomic. Refers to an operation that is never interrupted or left in an incomplete state under any circumstance.

Bean. A reusable software component. Beans can be combined to create an application.

Class In Java, a type that defines the implementation of a particular kind of object. A class definition defines instance and class variables and methods, as well as specifying the interfaces the class implements and the immediate superclass of the class. If the superclass is not explicitly specified, the superclass will implicitly be Object.

Classpath A classpath is an environmental variable that tells the Java Virtual Machine and other Java applications (for example, the Java tools located in the JDK1.1.X\bin directory) where to find the class libraries, including user-defined class libraries.

Codebase Works together with the code attribute in the <APPLET> tag to give a complete specification of where to find the main applet class file: code specifies the name of the file, and codebase specifies the URL of the directory containing the file.

Core class A public class (or interface) that is a standard member of the Java Platform. The intent is that the Java core classes, at minimum, are available on all operating systems where the Java Platform runs. A 100%-pure Java program relies only on core classes, meaning it can run anywhere.

Critical section A segment of code in which a thread uses resources (such as certain instance variables) that can be used by other threads, but that must not be used by them at the same time.

Deprecation Refers to a class, interface, constructor, method or field that is no longer recommended, and may cease to exist in a future version.

Derived from Class X is "derived from" class Y if class X extends class Y. See also subclass, superclass.

Exception An event, during program execution, that prevents the program from continuing normally; generally, an error. Java supports exceptions with the try, catch, and throw keywords. See also exception handler.

Exception handler A block of code that reacts to a specific type of exception. If the exception is for an error that the program can recover from, the program can resume executing after the exception handler has executed.

Extends Class X extends class Y to add functionality, either by adding fields or methods to class Y, or by overriding methods of class Y. An interface extends another interface by adding methods. Class X is said to be a subclass of class Y. See also derived from.

GUI Graphical User Interface. Refers to the techniques involved in using graphics, along with a keyboard and a mouse, to provide an easy-to-use interface to some program.

HotJava(TM) Browser An easily customisable Web browser developed by Sun Microsystems that is written in Java.

HTML HyperText Markup Language. This is a file format, based on SGML, for hypertext documents on the Internet. It is very simple and allows for the embedding of images, sounds, video streams, form fields and simple text formatting. References to other objects are embedded using URLs.

HTTP Hypertext Transfer Protocol. The Internet protocol, based on TCP/IP, used to fetch hypertext objects from remote hosts. See also TCP/IP.

IDL Java Interface Definition Language. Java API's that provide standards-based interoperability and connectivity with CORBA (Common Object Request Broker Architecture).

Instance An object of a particular class. In Java programs, an instance of a class is created using the new operator followed by the class name.

Interface In Java, a group of methods that can be implemented by several classes, regardless of where the classes are in the class hierarchy.

IP Internet Protocol. The basic protocol of the Internet. It enables the unreliable delivery of individual packets from one host to another. It makes no guarantees about whether or not the packet will be delivered, how long it will take, or if multiple packets will arrive in the order they were sent. Protocols built on top of this add the notions of connection and reliability. See also TCP/IP.

JAR file format JAR (Java Archive) is a platform-independent file format that aggregates many files into one. Multiple Java applets and their requisite components (.class files, images, sounds and other resource files) can be bundled in a JAR file and subsequently downloaded to a browser in a single HTTP transaction. It also supports file compression and digital signatures.

JavaBeans(TM) A portable, platform-independent reusable component model.

Java Database Connectivity (JDBC(TM)) An industry standard for database-independent connectivity between Java and a wide range of databases. The JDBC(TM) provides a call-level API for SQL-based database access.

Java(TM) Development Kit (JDK(TM)) A software development environment for writing applets and application in Java.

Java(TM) Foundation Class (JFC) An extension that adds graphical user interface class libraries to the Abstract Windowing Toolkit (AWT).

Java Platform The Java(TM) Virtual Machine and the Java core classes make up the Java Platform. The Java Platform provides a uniform programming

interface to a 100% Pure Java program regardless of the underlying operating system.

Java Remote Method Invocation (RMI) A distributed object model for Java-to-Java applications, in which the methods of remote Java objects can be invoked from other Java virtual machines, possibly on different hosts.

Java Runtime Environment (JRE) A subset of the Java(TM) Development Kit for end-users and developers who want to redistribute the JRE. The JRE consists of the Java Virtual Machine, the Java Core Classes, and supporting files.

JavaScript(TM) A Web scripting language that is used in both browsers and Web servers. It's only loosely related to Java and the name causes unnecessary confusion. Like any scripting language, it's used mostly to tie other components together or to accept user input.

Java(TM) Virtual Machine (JVM) The part of the Java Runtime Environment responsible for interpreting Java bytecodes.

JDK(TM) Java(TM) Development Kit. A software development environment for writing applets and application in Java.

JFC Java(TM) Foundation Class. An extension that adds graphical user interface class libraries to the Abstract Windowing Toolkit (AWT).

JRE Java Runtime Environment. A subset of the Java Developer Kit for end-users and developers who want to redistribute the JRE. The JRE consists of the Java Virtual Machine, the Java Core Classes, and supporting files.

Just-in-time (JIT) Compiler A compiler that converts all of the bytecode into native machine code just as a Java program is run. This results in run-time speed improvements over code that is interpreted by a Java Virtual Machine.

JVM Java Virtual Machine. The part of the Java Runtime Environment responsible for interpreting Java bytecodes.

Multithreaded Describes a program that is designed to have parts of its code execute concurrently. See also thread.

NCSA National Center for Supercomputer Applications.

Package A group of types. Packages are declared with the package keyword.

Process A virtual address space containing one or more threads.

RPC Remote Procedure Call. Executing what looks like a normal procedure call (or method invocation) by sending network packets to some remote host.

Sandbox Comprises a number of co-operating system components, ranging from security managers that execute as part of the application, to security measures designed into the Java Virtual Machine and the language itself. The sandbox ensures that a distrusted, and possibly malicious, application can't gain access to system resources.

Secure Socket Layer (SSL) A protocol that allows communication between a Web browser and a server to be encrypted for privacy. It can also provide communication between other entities.

Synchronized A Java keyword that, when applied to a method or code block, guarantees that at most one thread at a time executes that code.

TCP/IP Transmission Control Protocol based on IP. This is an Internet protocol that provides for the reliable delivery of streams of data from one host to another. See also IP.

Thin Client A system that runs a very light operating system with no local system administration and executes Java applications delivered over the network.

Thread The basic unit of program execution. A process can have several threads running concurrently, each performing a different job, such as waiting for events or performing a time-consuming job that the program doesn't need to complete before going on. When a thread has finished its job, the thread is suspended or destroyed. See also process.

Unicode A 16-bit character set defined by ISO 10646. See also ASCII. All Java source is written in Unicode.

URL Uniform Resource Locator. A standard for writing a text reference to an arbitrary piece of data in the WWW. A URL looks like "protocol://host/localinfo" where protocol specifies a protocol to use to fetch the object (like HTTP or FTP), host specifies the Internet name of the host on which to find it, and localinfo is a string (often a file name) passed to the protocol handler on the remote host.

Virtual machine An abstract specification for a computing device that can be implemented in different ways, in software or hardware. You compile to the instruction set of a virtual machine much like you'd compile to the instruction set of a microprocessor. The Java Virtual Machine consists of a bytecode instruction set, a set of registers, a stack, a garbage-collected heap, and an area for storing methods.

Wrapper An object that encapsulates and delegates to another object to alter its interface or behaviour in some way.

Server Definitions

API Application Programming Interface,

ASP Active Server Pages

Attribute An attribute is a string value that may be used in conjunction with a set of rules by the router to determine the next action to perform. Attributes are used to populate contexts with information about a message.

Business Logic Business logic is the 'user' code in the system. Business logic executes tasks such as 'debit account', 'retrieve balance' etc.

Configurable Entity is any Service or component that uses Configuration Objects and the Configuration Manager.

Configuration Object Configuration Objects hold persistent configuration data for services.

Configuration Service is a Service that implements a read-write interface to the Configuration Object.

Connection Manager Describes the process with which iPlanet™ Trustbase Transaction Manager communicates with external entities. It utilises the following objects to accomplish this task... Protocol Maps, Protocol Analysers, Handlers, Message Readers and Writers.

Connector The Connector is the main Connection Manager interface. It makes requests external to iPlanet™ Trustbase Transaction Manager. It takes a iPlanet™ Trustbase Transaction Manager Message containing the request, and a Destination Object describing the endpoint for the request.

Context Keeps a record of the current state of a given transaction.

Context Directive The action components that make up a ruleset.

CORBA Common Object Request Broker Architecture.

CSS Cascading Stylesheet.

Destination. Represents the destination of an external request, made by the Connector. An application specifies an implementation of Destination, and a

ProtocolMap that can transform the destination into a ProtocolDescriptor for the Connector that can then make and manage the actual connection.

Directive The 'action' part of a rule that is executed when the preconditions are true.

DOM Document Object Model.

DMZ De-militarised Zone.

DTD Data Type Definition or Document Type Definition.

Environment A set of contexts that are associated with a particular message.

EJB Enterprise Java Bean.

Host Environment Adaptor. The environment adapter forms the interface between a host such as a web server or application server and iPlanet™ Trustbase Transaction Manager itself.

HSM Hardware Security Module.

HTML HyperText Markup Language.

IDL Interface Definition Language.

JWS Java Web Server.

JDBC Java Database Connectivity.

lastService An attribute containing the name of the most recently executed service.

messageType An attribute contained within a message which holds the type of a given message. Message types are externally defined by the user.

Message An internal representation of a request from the user or a response from the server. Messages are routed within the system.

Message Analyser Provides the logic to identify which message reader or writer to use for a particular message based on the transport and the external format of the content.

Message Reader a Message Reader parses the remaining content of a Message from the InputStream, into the Message's content field. Message Readers may be a part of an application, and have specific knowledge of Message types, or they may be general purpose and have general knowledge of Message formats

Message Writer A Message Writer translates processed Message objects into the clients required presentation protocol, and write the results onto an Output Stream that is provided by the Protocol Analyser.

Message Registry A section of the tbase.properties file which provides a mapping between a message type and the message readers and writers used to process it.

Message Log Manager instantiates and allows access to Message Loggers. The message loggers are accessed according to which mime type they have registered interest in.

Message Logger Logs incoming and outgoing messages in their raw unprocessed form. The log can then be later queried and manipulated through the logManager or directly through the back end database engine.

NAS Netscape Application Server

NSK Non Stop Kernel

OAS Oracle Application Server.

PKI Public Key Infrastructure.

Precondition A precondition is a boolean expression which must be true for its corresponding directive to be executed. Preconditions are expressed in terms of attributes and their values. There are two types of precondition; an assertion that an attribute with a particular name exists and an assertion that the attribute with a given name not only exists but has a specific value.

Protocol Analyser Provides the logic to identify which protocol handler to use for a particular message type.

Protocol Handler The protocol Handler Component extracts the message type and context ID from the header of a message. There is usually one protocol handler for a particular message class e.g. iPlanet™ Trustbase Transaction Manager messaging, OFX etc. The protocol handler then routes appropriate protocol to the Message Analyser.

RMI Remote Method Invocation.

Private Rule Set Repository A collection of rule sets that ships with the iPlanet™ Trustbase Transaction Manager and allows built in services (such as the configuration services) to function.

Protocol Descriptor. Holds a description of the endpoint, transport and presentation protocols for a connection, in the form of a URL, and the format of the message to be sent, and as a mime type Implements Destination. It can be used with the SimpleProtocolMap for direct Destination addressing.

Protocol Map. An application specifies ProtocolMap implementations to map it's Destination implementations to URLs and mime types that the Connector can use to make an actual connection.

Protocol Map Manager. Manages a set of ProtocolMap implementations, selecting an appropriate ProtocolMap to translate a particular Destination implementation into a ProtocolDescriptor.

Public Rule Set Repository A user configurable collection of user-defined Rule Sets Rules in the private Rule Set Repository take precedence over rules in the public Rule Set Repository.

Role Role is not set of attributes, it is the name of a particular attribute which the system recognizes. There are several such attributes including:

lastService - the name of the most recently executed service

messageType - the type of a given message

role - a string representing the capacity in which the user is

using the system, eg role - operator, role =administrator etc.

Router The router provides a mechanism for imposing structure and ordering on the execution of services in a secure way which doesn't necessitate code changes.

Rule A Rule contains three components: a Rule Name, Precondition and Directive. If the precondition is true then the directive is executed.

Ruleset A collection of rules that route messages to one or more services to achieve a given task.

Rule Name. Every rule has a name. The rule is referred to by its name in the context of a ruleset.

Service An object implementing business logic. Services are written by the user.

Service Registry A registry of services! Used to provide a lookup between service names and the classes that implement them.

Session A session is the container for all of the tasks a user is performing over a period of time.

SP Service Provider

State A collection of attributes associated with a task at a given instant in time.

Stub The client portion of a distributed object using mechanisms such as CORBA or RMI. The Stub is designed to hide the fact that the implementation of it's related object is not locally located. See also Skeleton.

Skeleton The server/remote portion of a distributed object under CORBA and RMI. The skeleton is invoked by the Stub. See also stub.

Task A unit of work at the business level. A ruleset defines how a task will be executed.

TISS Transport Independent Stub Service

URL Uniform Resource Locator

X500 Set of Open Standards for directory services. See, for instance, Country code that is defined as a ISO3611 standard <http://www.iso.ch> and X500 standard <http://www.itu.int/itudoc/itu-t/rec/x/x500up/x500.html>

XML Extensible Markup Language

XSL Extensible Stylesheet Language

Index

3

3DES, 93

A

Abstract class, 95
Abstract method, 95
Advanced Routing, 48
API, 8, 16, 35, 37, 38, 39, 41, 51, 55, 56, 60, 62, 64, 77, 82, 85, 86, 95, 96, 98
API packages
 com.iplanet.trustbase.identrus, 22, 42, 77, 85, 86
 com.iplanet.trustbase.identrus.message, 77, 85, 86
 com.iplanet.trustbase.identrus.security, 85
 com.iplanet.trustbase.util.tree, 85
 com.iplanet.trustbase.xml.dsig, 85
 uk.co.jcp.tbase.config, 55, 56, 85
 uk.co.jcp.tbase.connector, 35, 37, 38
 uk.co.jcp.tbase.xurl, 35, 39
 uk.co.jcp.tbaseimpl.connector, 35, 37
 uk.co.jcp.tbaseimpl.log.audit.type, 61, 62, 63
 uk.co.jcp.tbaseimpl.parse.message.http, 25
Applet, 95
Application protocol, 93
ASN.1, 93, 95
ASP, 98
Assigning a role to a service, 81, 90
Asymmetric cipher, 93
Atomic, 95
Attribute, 21, 22, 24, 41, 42, 50, 51, 98
Audit, 61, 62, 63
Audit Logging an Event, 61
Audit logs, 61
AuditObject, 61, 62, 63
Authentication, 11, 17, 45, 58, 93, 94
Authentication and Authorisation, 11, 45, 58
Authorisation, 17, 42, 45, 81, 90

B

base64, 93, 95
BBS, 93
Bean, 95, 99
BER, 93, 95
bill_data table, 73
Billing records, 73
Block cipher, 93
Building Identrus solutions, 47, 74
Bulk cipher, 93
Business Logic, 98

C

cert_data table, 72
CertData, 72
Certificate, 8, 9, 10, 16, 21, 22, 24, 34, 39, 42, 92, 93
Certificate requests and responses, 92
Cipher Block Chaining Mode, 93

CipherSuite, 70
Class, 63, 75, 76, 78, 95, 96, 97
Class generation, 75
Classes
 Attribute, 21, 22, 24, 41, 42, 50, 51, 98
 Audit, 61, 62, 63
 AuditObject, 61, 62, 63
 Certificate, 8, 9, 10, 16, 21, 22, 24, 34, 39, 42, 92, 93
 ConfigManager, 56
 ConfigUID, 55, 56
 ConfigurationLock, 55
 ConfigurationObject, 55, 56
 ConnectionProtocolAnalyser, 36
 Connector, 36, 37, 39, 98, 100
 DefaultConnector, 37
 Destination, 36, 37, 38, 39, 98, 100
 ErrorLog, 64, 65, 66, 86
 ErrorObject, 64, 65, 66, 86
 HTTPReader, 23
 IdentrusConstants, 21, 22, 24, 34, 42
 IdentrusMessage, 77, 86
 IdentrusService, 86
 InvalidDestinationException, 38
 Message, 16, 20, 21, 22, 23, 24, 25, 34, 36, 41, 42, 68, 71, 72, 76, 78, 92, 94, 98, 99, 100
 MessageAnalyserException, 68
 MessageReader, 36
 MessageReaderException, 68
 MessageWriter, 25, 36
 MessageWriterException, 68
 NIBAccessor, 86
 NIBAccessorException, 86
 OperationBeginAudit, 61, 63
 ParseException, 68
 Password, 94
 ProtocolAnalyserException, 68
 ProtocolDescriptor, 36, 37, 38, 39, 99, 100
 ProtocolHandler, 36
 ProtocolMap, 36, 37, 38, 99, 100
 ProtocolMapManager, 36, 37, 38, 39
 Router, 17, 23, 25, 41, 43, 47, 49, 53, 67, 68, 100
 RoutingException, 67, 68
 ScriptWriter, 25, 28, 29, 56
 SecurityContext, 42
 Service, 36, 37, 38, 43, 55, 56, 57, 61, 62, 77, 78, 79, 81, 84, 85, 87, 90, 98, 100, 101
 ServiceException, 67, 68, 86
 Signature, 24, 83, 93
 SimpleProtocolMap, 38, 100
 SingletonConfigManager, 55, 56
 TbaseElement, 77
 TbaseIdentifiedElement, 77
 TbaseRuntimeException, 68
 XURL, 39
 XURLStreamHandler, 39
 XURLStreamHandlerFactory, 39
classname, 65
Classpath, 95
Client, 93
Client write key, 93
ClientCertIssuerDN, 70

ClientCertSerialNumber, 70
 com.ipplanet.trustbase.identrus, 22, 42, 77, 85, 86
 com.ipplanet.trustbase.identrus.message, 77, 85, 86
 com.ipplanet.trustbase.identrus.security, 85
 com.ipplanet.trustbase.util.tree, 85
 com.ipplanet.trustbase.xml.dsig, 85
 Complete Router Rule DTD, 53
 Component replication, 58
 ConfigManager, 56
 ConfigUID, 55, 56
 Configurable Entity, 55, 56, 98
 Configuration Management, 54
 Configuration Manager, 54, 55, 56, 98
 Configuration Object, 54, 55, 56, 98
 Configuration Objects, 54, 55, 56, 98
 Configuration Service, 55, 56, 98
 Configuration Services, 55, 56
 Configuration Store, 54, 55, 56
 ConfigurationLock, 55
 ConfigurationObject, 55, 56
 Connection, 19, 35, 36, 37, 39, 70, 71, 93, 98
 Connection information, 70
 Connection Manager, 19, 35, 36, 37, 39, 98
 Connection Manager Architecture, 35
 connection_id, 71
 ConnectionFailed, 70
 ConnectionFailedReason, 70
 connectionid, 72
 ConnectionId, 70
 ConnectionProtocolAnalyser, 36
 ConnectIPAddr, 70
 Connector, 36, 37, 39, 98, 100
 ConnectTime, 70
 Context, 21, 36, 98
 Context Directive, 98
 contextid, 65
 CORBA, 15, 16, 96, 98, 101
 Core class, 95
 Create DTD Definitions, 83
 Creating the Identrus Service JAR, 87
 Critical section, 95
 CRL, 93
 CSS, 98

D

data, 12, 15, 16, 17, 20, 23, 24, 29, 30, 36, 37, 38, 41, 45,
 55, 56, 59, 60, 64, 66, 69, 70, 72, 73, 75, 77, 93, 94,
 95, 97, 98
 Data definitions, 70
 Data Encryption Standard, 93
 datatype, 66
 Default HTML Message Writer, 25
 Default Identrus Error Writer, 34
 Default Identrus Message Writer, 34
 Default Message Reader, HTTP Reader, 23
 Default Message Reader, Identrus Error Reader, 24
 Default Message Reader, Identrus Reader, 24
 Default routing, 47
 Default routing rules, 47
 DefaultConnector, 37
 Defining a New Error, 66
 Defining New Audit Types, 62
 Deploying ping.jar within iPlanet Trustbase Transaction
 Manager, 88
 Deploying PingService within iPlanet Trustbase
 Transaction Manager, 80, 88
 Deprecation, 96

DER, 93, 95
 DES, 93
 Destination, 36, 37, 38, 39, 98, 100
 Development process, 75
 DH, 93
 digestofrecord, 72
 Digital Signature Standard, 93
 Digital signatures, 85, 93
 Directive, 50, 99, 100
 DMZ, 99
 doctype, 47, 71, 72
 DOM, 20, 23, 25, 92, 99
 DOMHASH, 91
 DSA, 93
 DTD, 21, 34, 49, 50, 53, 74, 75, 76, 77, 82, 83, 84, 86,
 92, 99
 DTD Rule Body, 50

E

EJB, 2, 14, 16, 39, 99
 Enterprise connectivity, 16
 Environment, 99
 Error and Audit Logging, 59
 Error handling and logging, 64
 Error Logging, 64
 Error Severity Types, 64
 Error table, 65
 error_codes table, 65, 66
 error_support, 66
 errorcode, 65
 errorid, 65, 66
 ErrorLog, 64, 65, 66, 86
 ErrorObject, 64, 65, 66, 86
 Exception, 66, 67, 68, 96
 Exception handler, 96
 Exception Handling, 67
 Exceptions, 64, 68
 External interfaces, 15

G

GUI, 96

H

Handshake, 94
 Hardware Security nCipher KeySafe 1.0 and CAFast, 9,
 91
 Host Environment Adaptor, 99
 HSM, 99
 HTML, 10, 15, 16, 20, 21, 25, 26, 29, 55, 56, 92, 96, 99
 HTTP, 10, 15, 16, 23, 24, 56, 68, 72, 76, 91, 93, 96, 98
 HTTPReader, 23

I

IDEA, 94
 Identrus, 9, 10, 11, 13, 15, 16, 20, 21, 22, 24, 34, 40, 42,
 44, 45, 47, 48, 65, 69, 70, 72, 73, 74, 75, 76, 77, 78,
 82, 83, 84, 85, 86, 92
 Identrus logging, 69
 Identrus Message Attributes, 42
 Identrus Message Specifications, 9
 Identrus protocol handler, 21
 Identrus Transaction Coordinator, 10, 11, 69
 IdentrusConstants, 21, 22, 24, 34, 42
 IdentrusMessage, 77, 86

IdentrusService, 86
 IDL, 96, 99
 Initialization Vector, 94
 input, 23, 39, 72, 75, 77, 97
 Interface, 2, 95, 96, 98, 99
 Introduction, 7, 10
 InvalidDestinationException, 38
 IP, 34, 65, 70, 96, 97
 iPlanet Application Server 4.1, 9, 91
 iPlanet Certificate Management System, 9
 iPlanet Trustbase Transaction Manager Architecture, 13
 iPlanet Trustbase Transaction Manager Interfaces, 15
 iPlanet Trustbase Transaction Manager log manager, 60
 iPlanet Trustbase Transaction Manager Platform, 10
 iPlanet Web Server 6.0, 9, 91
 IssuerDN, 72, 73

J

JAR, 47, 62, 75, 82, 85, 87, 96
 Java, 2, 8, 37, 40, 41, 57, 62, 63, 65, 66, 75, 76, 82, 85, 91, 95, 96, 97, 98, 99
 Java Database Connectivity, 96, 99
 Java Platform, 95, 96
 Java Remote Method Invocation, 97
 Java Runtime Environment, 97
 JavaBeans, 2, 96
 JavaScript, 2, 97
 JDBC, 2, 15, 16, 56, 70, 96, 99
 JDK, 2, 75, 79, 82, 87, 96, 97
 JFC, 96, 97
 JRE, 75, 97
 Just-in-time, 97
 JVM, 2, 97
 JWS, 99

K

KeySafe, 9

M

machineid, 65
 Master secret, 94
 MD5, 94
 message, 10, 11, 13, 14, 15, 16, 17, 20, 21, 22, 23, 24, 25, 26, 28, 30, 31, 34, 35, 36, 37, 39, 40, 41, 42, 43, 44, 45, 47, 48, 49, 50, 51, 56, 58, 64, 65, 68, 70, 71, 72, 73, 75, 76, 77, 78, 79, 84, 85, 86, 87, 94, 98, 99, 100
 Message, 16, 20, 21, 22, 23, 24, 25, 34, 36, 41, 42, 68, 71, 72, 76, 78, 92, 94, 98, 99, 100
 Message Analyser, 25, 68, 99, 100
 Message Attributes, 41
 Message Authentication Code, 94
 Message Log Manager, 99
 Message Logger, 99
 Message path processing, 78
 Message Protocols, 92
 Message Reader, 23, 41, 68, 98, 99
 Message Readers, 23, 41, 68, 98, 99
 Message Registry, 99
 Message Writer, 25, 34, 41, 42, 68, 99
 Message Writers, 25, 41, 68
 message_invalid_reason, 71
 message_protection, 70
 message_valid, 71
 MessageAnalyserException, 68

MessageReader, 36
 MessageReaderException, 68
 Messages, 24, 25, 41, 42, 47, 85, 94, 99
 messageType, 41, 43, 99, 100
 MessageWriter, 25, 36
 MessageWriterException, 68
 Methodology, 75
 MIME, 20, 21, 23, 25, 41, 94
 msggrpid, 71, 72
 msgid, 71, 72
 Multithreaded, 97

N

nCipher, 9
 NCSA, 97
 NIBAccessor, 86
 NIBAccessorException, 86
 NSK, 99

O

OAS, 99
 OCSP, 21, 92
 One-way hash function, 94
 OperationBeginAudit, 61, 63
 Oracle 8i, 9, 91
 Oracle 8i Installation and Configuration Guides, 9
 OSI, 93, 94
 Overall Layout, 8
 Overview, 14, 20, 58, 60, 70

P

ParseException, 68
 Password, 94
 PBE, 94
 peer_cert_serial_number, 70
 peer_ip_addr, 71
 peer_issuer_dn, 70
 PEM, 94
 Ping Example, 82
 PingService Source Code, 86
 PKI, 8, 94, 99
 Precondition, 100
 Presentation layer components, 20
 Presentation logic, 16, 19
 Private Rule Set Repository, 100
 Product features, 12
 Product Features, 12
 Protocol Analyser, 68, 98, 99, 100
 Protocol Descriptor, 100
 Protocol Handler, 21, 41, 65, 100
 Protocol handlers, 21
 Protocol Map, 37, 38, 98, 100
 Protocol Map Manager, 37, 100
 ProtocolAnalyserException, 68
 ProtocolDescriptor, 36, 37, 38, 39, 99, 100
 ProtocolHandler, 36
 ProtocolMap, 36, 37, 38, 99, 100
 ProtocolMapManager, 36, 37, 38, 39
 protocoltype, 72
 Public key cryptography, 94
 Public Key Infrastructure, 94, 99
 Public Rule Set Repository, 100

R

Raw log tables, 72
 raw_data table, 71, 72
 rawdata, 71, 72
 RawRecordId, 73
 RC2, 94
 RC4, 94
 recipients, 71
 recordmarker, 71, 72
 References and Glossary, 91
 Related Documents, 9
 Return path, 47
 RFC, 72, 94
 RMI, 15, 16, 97, 100, 101
 RMI Remote Method Invocation., 100
 Role, 43, 100
 Router, 17, 23, 25, 41, 43, 47, 49, 53, 67, 68, 100
 Router Architecture, 43
 Router Rule Syntax, 49
 Router Rules, 47
 Routing, 14, 17, 40, 44, 47, 48
 Routing Rule Sets, 48
 Routing to service, 47
 RoutingException, 67, 68
 RPC, 97
 RSA, 72, 94
 Rule, 48, 49, 53, 100
 Rule Name, 100
 Rule Sets, 48, 49, 100
 Ruleset, 100

S

S/MIME Version 2 Message Specification, 91
 Salt, 94
 Sandbox, 97
 Script Tags, 25, 26
 ScriptWriter, 25, 28, 29, 56
 Secure Socket Layer, 97
 Security Related Protocols, 91
 Security related terms, 93
 SecurityContext, 42
 sender, 71
 SerialNumber, 72, 73
 Server, 14, 16, 94, 95, 98, 99
 Server Definitions, 98
 Server to server connectivity, 16
 Server write key, 94
 Server write MAC secret, 95
 servercertissuerdn, 72
 servercertserialnumber, 72
 Service, 36, 37, 38, 43, 55, 56, 57, 61, 62, 77, 78, 79, 81, 84, 85, 87, 90, 98, 100, 101
 Service Building, 79
 Service Deployment, 79
 Service development, 78
 Service Registry, 100
 ServiceException, 67, 68, 86
 Services, 13, 17, 18, 37, 46, 47, 55, 57, 68, 75, 81, 88, 89, 90, 100
 Session, 78, 94, 100
 severity, 64, 65
 SHA, 94, 95
 Signature, 24, 83, 93
 signedigestofcalculation, 72
 SimpleProtocolMap, 38, 100
 SingletonConfigManager, 55, 56

Skeleton, 101
 smime_transport table, 70
 SMTP RFC821, 91
 smtp_connection table, 71
 smtp_message table, 70, 71
 Software Platform, 91
 Solaris 8 and Java Development Kit 1.2.1, 9
 Solaris 8 and JDK, 91
 SP, 101
 SSL, 15, 16, 39, 42, 45, 70, 72, 93, 94, 95, 97
 ssl_connection, 70, 71
 Standard Services, 18, 57
 State, 78, 101
 Stream cipher, 95
 stream_id, 71
 Stub, 86, 101
 SubjectDN, 73
 SUN Microsystems Java Related terms, 95
 Synchronized, 97

T

Tables

bill_data table, 73
 cert_data table, 72
 CertData, 72
 CipherSuite, 70
 classname, 65
 ClientCertIssuerDN, 70
 ClientCertSerialNumber, 70
 connection_id, 71
 ConnectionFailed, 70
 ConnectionFailedReason, 70
 connectionid, 72
 ConnectionId, 70
 ConnectIPAddr, 70
 ConnectTime, 70
 contextid, 65
 data, 12, 15, 16, 17, 20, 23, 24, 29, 30, 36, 37, 38, 41, 45, 55, 56, 59, 60, 64, 66, 69, 70, 72, 73, 75, 77, 93, 94, 95, 97, 98
 datatype, 66
 digestofrecord, 72
 doctype, 47, 71, 72
 Error table, 65
 error_codes table, 65, 66
 error_support, 66
 errorcode, 65
 errorid, 65, 66
 input, 23, 39, 72, 75, 77, 97
 IssuerDN, 72, 73
 machineid, 65
 message, 10, 11, 13, 14, 15, 16, 17, 20, 21, 22, 23, 24, 25, 26, 28, 30, 31, 34, 35, 36, 37, 39, 40, 41, 42, 43, 44, 45, 47, 48, 49, 50, 51, 56, 58, 64, 65, 68, 70, 71, 72, 73, 75, 76, 77, 78, 79, 84, 85, 86, 87, 94, 98, 99, 100
 message_invalid_reason, 71
 message_protection, 70
 message_valid, 71
 msggrpid, 71, 72
 msgid, 71, 72
 peer_cert_serial_number, 70
 peer_ip_addr, 71
 peer_issuer_dn, 70
 protocoltype, 72
 raw_data table, 71, 72
 rawdata, 71, 72

RawRecordId, 73
 recipients, 71
 recordmarker, 71, 72
 sender, 71
 SerialNumber, 72, 73
 servercertissuerdn, 72
 servercertserialnumber, 72
 severity, 64, 65
 signeddigestofcalculation, 72
 smime_transport table, 70
 smtp_connection table, 71
 smtp_message table, 70, 71
 ssl_connection, 70, 71
 stream_id, 71
 SubjectDN, 73
 time_stamp, 70
 time_stamp_type, 70
 timestamp, 65, 70, 71, 72
 timestamptype, 71
 TimeStampType, 70
 Task, 101
 TbaseElement, 77
 TbaseIdentifiedElement, 77
 TbaseRuntimeException, 68
 TCP/IP, 93, 96, 97
 Thin Client, 97
 Three Tier Architecture, 14
 time_stamp, 70
 time_stamp_type, 70
 timestamp, 65, 70, 71, 72
 timestamptype, 71
 TimeStampType, 70
 TISS, 101
 Trading protocols, 92
 Transport protocols, 15
 Transport Protocols, 91
 TSL, 95

U

uk.co.jcp.tbase.config, 55, 56, 85
 uk.co.jcp.tbase.connector, 35, 37, 38
 uk.co.jcp.tbase.xurl, 35, 39
 uk.co.jcp.tbaseimpl.connector, 35, 37
 uk.co.jcp.tbaseimpl.log.audit.type, 61, 62, 63
 uk.co.jcp.tbaseimpl.parse.message.http, 25
 URL, 36, 38, 39, 95, 98, 100, 101
 URL Connection Implementation, 39
 Using the default HTTP Reader, 23
 Using the ScriptWriter tags, 29

V

Virtual machine, 98

W

Wrapper, 98

X

X500, 101
 X509, 45, 93, 95
 X690, 95
 XML, 8, 10, 16, 20, 21, 22, 25, 26, 27, 28, 29, 30, 34, 40,
 42, 44, 45, 49, 72, 75, 76, 77, 84, 85, 92, 101
 XML Repeat Iterators, 27
 XML Syntax Processing specification, 92
 XSL, 25, 101
 XURL, 39
 XURLStreamHandler, 39
 XURLStreamHandlerFactory, 39