

webMethods Trading Networks Administrator's Guide

Version 10.2

April 2018

This document applies to webMethods Trading Networks Version 10.2 and to all subsequent releases.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © 2007-2018 Software AG, Darmstadt, Germany and/or Software AG USA Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.

The name Software AG and all Software AG product names are either trademarks or registered trademarks of Software AG and/or Software AG USA Inc. and/or its subsidiaries and/or its affiliates and/or their licensors. Other company and product names mentioned herein may be trademarks of their respective owners.

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <http://softwareag.com/licenses>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices, license terms, additional rights or restrictions, please refer to "License Texts, Copyright Notices and Disclaimers of Third Party Products". For certain specific third-party license restrictions, please refer to section E of the Legal Notices available under "License Terms and Conditions for Use of Software AG Products / Copyright and Trademark Notices of Software AG Products". These documents are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

Table of Contents

About this Guide.....	15
Document Conventions.....	15
Online Information.....	16
Understanding webMethods Trading Networks.....	17
webMethods Trading Networks.....	18
Architecture.....	18
Partners in a Trading Network.....	19
Understanding the Trading Networks Terminology.....	21
Asset Definition.....	25
Document Attributes.....	26
Document Types.....	27
XML Document Types.....	28
Flat File Document Types.....	29
Processing Rules.....	30
Profiles.....	32
Overview of Creating Partner Profiles.....	32
Overview of the Partner Onboarding Process.....	33
TPAs.....	34
Business Process Definition.....	35
Document Delivery.....	36
Immediate Delivery.....	37
Scheduled Delivery.....	38
Queue Documents for Polling.....	39
Document Processing.....	40
Error Logging.....	43
Document Recognition Errors.....	43
Document Processing Errors.....	43
Monitoring Transactions.....	45
Security.....	46
Managing, Authenticating, and Authorizing Trading Networks Users.....	46
Protecting Access to User Interfaces.....	47
Protecting Partner Profile Passwords.....	47
Protecting Access to Trading Networks Processing.....	48
Certificates for Verifying, Signing, Encrypting, and Decrypting Documents and Authenticating Connections.....	49
Overlapping of Certificates.....	50
Verifying Digital Signatures.....	59
Digitally Signing Documents.....	60
Encrypting and Decrypting Data.....	60
Communicating Securely Using SSL.....	61

Run-Time Event Notifications.....	61
Caching.....	62
Dashboards and Charts.....	62
Database Partitioning.....	63
Getting Started.....	65
Before Configuring Trading Networks.....	66
Starting the Host Integration Server.....	66
Starting Integration Server on Microsoft Windows.....	66
Starting Integration Server on UNIX.....	66
Starting Integration Server from the Command Line.....	67
Shutting Down Integration Server.....	67
Planning Your Trading Network.....	67
Creating Services and Viewing Specifications.....	68
Additional Documentation.....	68
Sending Documents to Trading Networks.....	69
Sending Documents to Your Trading Networks from Another System.....	69
Sending Documents from One Trading Networks to Another.....	70
Sending a Document Back to Your Trading Networks.....	71
Summary of the Configuration Steps.....	71
Using Command Central to Manage Trading Networks.....	75
Configuring Trading Networks.....	77
Overview.....	78
Configuring Database Settings.....	78
Configuring a High Availability Environment.....	78
Configuring a Cluster to Share a Database.....	79
Configuring Multiple Clusters Instances to Share a Database.....	79
Configure a Non-Clustered Instance to Each Have Its Own Database.....	80
Configuring Trading Networks for MySQL Community Edition 5.7.....	80
Configuring Trading Networks for MySQL Enterprise Edition 5.7.....	82
Configuring Trading Networks for a Clustered Environment.....	82
Configuring E-mail Settings for Document Delivery.....	84
Configuring Alias for ActiveTransfer on Remote Server.....	84
Configuring Task Settings.....	85
Configuring User Status.....	86
Configuring Trading Networks to Work with My webMethods.....	87
Configuring Integration Server to Use Central User Management.....	87
Specifying the SAML Resolver Location.....	87
Connecting Trading Networks Servers to My webMethods Server.....	87
Configuring the Timeout for Trading Networks Web Services.....	88
Configuring Resubmission and Reprocessing Settings.....	89
Granting Access to Trading Networks.....	91
Overview.....	92

Defining Roles.....	92
Adding My webMethods Users to the TN Administrators Role.....	93
Granting Users Authority to Connect Trading Networks Server Instances to My webMethods.....	93
Granting Users Access to Trading Networks Server Instances.....	94
Granting Users Authority to Execute Trading Networks Services.....	94
Granting Users General Functional Permissions.....	94
Granting Users Data Permissions.....	96
Defining a Data Set.....	96
Granting Data Permissions to a Data Set.....	98
Setting Partner Permissions.....	101
Removing a Data Set and Its Data Permissions.....	101
Removing References to Users.....	101
Creating Custom Document Attributes.....	103
Create Custom Attributes.....	104
Creating XML Document Types.....	105
Before You Begin.....	106
Obtaining XML Document Samples.....	106
Creating IS Schemas.....	106
Creating IS Document Types.....	106
Creating Custom Transformations for Extracted Attributes.....	106
Creating a Service for the Check for Duplicate Document Action.....	107
Creating an XML Document Type.....	108
Creating and Naming the Document Type.....	108
Specifying Criteria for Matching Inbound Documents to the Document Type.....	108
Specifying the Root Tag.....	109
Specifying the DOCTYPE Identifier.....	109
Defining Identifying Queries.....	109
Specify Pipeline Variables.....	110
Specifying Document Attributes to Extract.....	110
Specifying the Attributes.....	111
SenderID and ReceiverID System Attributes.....	111
DocumentID.....	112
UserStatus.....	112
GroupID.....	112
ConversationID.....	112
SignedBody and Signature.....	112
Custom Attributes.....	112
Indicating Whether the Attribute is Required.....	112
Transforming Extracted Attributes.....	113
Built-In Transformations for DATETIME and DATETIME LIST Data Type Attributes.....	113
Built-In Transformations for STRING or STRING LIST Data Type Attributes.....	113
Built-In Transformations for NUMBER Data Type Attributes.....	114

Custom Transformation Services.....	115
Specifying the Namespace Mapping.....	115
Defining Pre-Processing Actions for the Document Type.....	116
Format as an IS Document Type Option.....	116
Enable Processing Rule Routing Option.....	116
Verify Digital Signature Pre-Processing Action.....	116
Validate Structure Pre-Processing Action.....	116
Check for Duplicate Document Pre-Processing Action.....	117
Save Document to Database Pre-Processing Action.....	117
Editing an XML Document Types.....	118
Testing XML Document Types.....	118
Creating Flat File Document Types.....	121
Before You Begin.....	122
Registering Flat File Content Types.....	122
Creating Flat File Schemas.....	122
Creating Custom Transformations for Extracted Attributes.....	122
Creating a Service for the Verify Signature Action.....	123
Creating a Service for the Check for Duplicate Document Action.....	123
Creating a Document Gateway Service.....	123
Reading the Input Stream.....	124
Specifying and Passing the Outputs.....	124
DoctypeID or DoctypeName.....	125
SenderID and ReceiverID.....	125
DocumentID.....	125
GroupID.....	125
ConversationID.....	125
UserStatus.....	126
processingRuleID or processingRuleName.....	126
\$receiveSvc.....	126
\$contentType and \$contentEncoding.....	126
Custom Attributes.....	127
Pass the Outputs to Trading Networks.....	127
Creating a Flat File Document Type.....	127
Creating and Naming the Document Type.....	127
Specifying Criteria for Matching Inbound Documents to the Document Type.....	128
Specifying Attributes to Extract.....	128
Specifying the Attributes.....	128
SenderID and ReceiverID System Attributes.....	128
DocumentID.....	129
UserStatus.....	129
GroupID.....	129
ConversationID.....	130
Custom Attributes.....	130
Indicating Whether the Attribute is Required.....	130

Transforming Extracted Attributes.....	130
Built-In Transformations for DATETIME and DATETIME LIST Data Type Attributes.....	130
Built-In Transformations for STRING or STRING LIST Data Type Attributes....	130
Built-In Transformations for NUMBER Data Type Attributes.....	131
Custom Transformation Services.....	132
Defining Pre-Processing Actions for the Document Type.....	132
Enable Processing Rule Routing Option.....	133
Verify Digital Signature Pre-Processing Action.....	133
Validate Structure Pre-Processing Action.....	133
Check for Duplicate Document Pre-Processing Action.....	133
Save Document to Database Pre-Processing Action.....	134
Editing a Flat File Document Type.....	135
Preparing for Document Delivery.....	137
Overview.....	138
Creating Custom Immediate Delivery Services.....	138
Creating a Custom Immediate Delivery Service.....	138
Retrieving the Document Content to Deliver.....	139
Registering the New Delivery Service.....	139
Enabling Document Delivery Using a Web Service.....	141
Creating Custom Scheduled Delivery Services.....	143
Creating a Custom Scheduled Delivery Service.....	143
Scheduled Delivery Service Inputs.....	143
Scheduled Delivery Service Outputs.....	143
Retrieving the Document Content to Deliver.....	144
Required Logic to Include in the Scheduled Delivery Service.....	144
Typical Logic to Include in the Scheduled Delivery Service.....	145
Handling Exceptions.....	146
The wm.tn.transport:batchFtp Built-in Service.....	147
Registering the New Delivery Service.....	150
Adding a Public Queue to Trading Networks.....	152
Setting Up the Queue for Polling for webMethods for Partners.....	154
Creating Processing Rules.....	157
Overview.....	158
Before You Begin.....	158
Creating a Service for the Check for Duplicate Document Action.....	158
Creating a Service for the Execute a Service Action.....	158
Using Alert Email Messages.....	159
Performing Tasks for the Deliver Document By Action.....	160
Creating a Processing Rule.....	160
Creating and Naming the Rule.....	160
Specifying Criteria for the Rule.....	160
Specifying Sender and Receiver Criteria.....	160
Specifying Document Type Criteria.....	161

Specifying User Status Criteria.....	161
Specifying Recognition Error Criteria.....	162
Setting Up Extended Criteria for the Rule.....	162
Attribute Field.....	162
Operator Field.....	162
Value Field.....	162
Defining Pre-Processing Actions for the Rule.....	163
Verify Digital Signature Pre-Processing Action.....	163
Validate Structure Pre-Processing Action.....	163
Check for Duplicate Document Pre-Processing Action.....	164
Save Document to Database Pre-Processing Action.....	164
Defining Processing Actions for the Rule.....	165
Execute a Service Action.....	165
Alert e-Mail Message Action.....	166
Email To Field.....	166
Subject Field.....	166
Mail Body Field.....	166
Change User Status Action.....	167
Deliver Document By Action.....	167
Respond Action.....	168
Defining a Processing Rule Based on a Saved Document.....	169
Editing the Default Processing Rule.....	170
Reordering Processing Rules.....	170
Testing the Order of Processing Rules.....	171
Deleting a Processing Rule.....	171
Onboarding New Partners.....	173
Overview.....	174
Summary of the Partner Onboarding Process.....	174
Uploading Partner Information.....	175
Uploading Partner Information from a Spreadsheet.....	175
Adding Partner Information Using the Quick Add Feature.....	176
Managing Partner Onboarding Templates.....	176
Searching for a Template by Name.....	177
Adding a New Template.....	177
Creating a New Group from an Existing Group.....	177
Adding New Fields to an Existing Group.....	178
Managing Fields in a Group.....	178
Creating a Mapping Service for New Templates.....	179
Managing Invitation Emails.....	180
Configuring Invitation Emails.....	180
Sending Invitation Emails.....	181
Managing Partner Invitation Emails.....	181
Granting Partners Access to the Questionnaire Page.....	182
Approving or Rejecting a Partner.....	183

Viewing and Changing the Status of a Partner.....	183
Creating Profiles.....	185
Overview.....	186
Before You Begin.....	186
Creating Contact Types.....	186
Creating External ID Types.....	186
Creating Services for Use with TPAs.....	186
Creating a Profile.....	187
Creating and Naming a Profile.....	187
Adding External IDs to a Profile.....	188
Adding Addresses to a Profile.....	189
Copying or Sharing Addresses.....	189
Adding Contacts to a Profile.....	190
Copying or Sharing Contacts.....	190
Specifying Document Delivery in a Profile.....	190
Adding an Immediate Delivery Method to a Profile.....	191
Copying or Sharing Immediate Delivery Methods.....	195
Adding Scheduled Delivery Method to a Partner Profile.....	196
Specifying Reliable Delivery Settings in a Partner Profile.....	198
Specifying Queue for Polling Settings in a Partner Profile.....	199
Suspending or Resuming Document Delivery for a Partner.....	200
Scheduling a Delivery Suspension.....	201
Editing an Ongoing Delivery Suspension Schedule.....	202
Canceling an Ongoing Delivery Suspension Schedule.....	203
Associate My webMethods User Accounts with a Profile.....	203
Adding Security Information.....	204
Adding Certificate Sets to a Profile.....	205
Adding the Default Certificate Set to a Profile.....	205
Adding a Certificate Set to Use with a Specific Partner.....	206
Updating Certificate Sets.....	207
Updating Certificate Sets for Your Enterprise Profile.....	207
Updating Certificate Sets for a Partner Profile.....	208
Making Standard Fields Required.....	208
Adding Extended Fields to a Profile.....	209
Opening a Profile's Extended Fields.....	209
Adding Extended Fields to a Profile.....	209
Associating Partners with Partner Groups.....	210
Working with TPAs.....	211
Defining a TPA.....	211
Opening a TPA for Editing.....	213
Editing a TPA.....	213
Editing a TPA's Fields.....	213
Changing a TPA's Agreement Status.....	214
Deleting a TPA.....	215

Deleting a Partner Profile.....	215
Finding Partner Certificates that Are Expired or Expiring Soon.....	215
Finding Expired Partner Certificates.....	215
Finding Partner Certificates that Expire Soon.....	216
Managing Partner Access to APIs.....	216
Integrating Trading Networks with API Gateway.....	216
Configuring the API Gateway Property.....	217
Viewing the APIs Accessible to a Partner.....	217
Viewing the List of APIs of a Partner Group.....	217
Providing API Access for a Partner using Partner Profiles.....	218
Revoking API Access for a Partner from Partner Profile.....	218
Granting Permissions to a User to Access APIs.....	218
Providing API Access to a Partner Group.....	218
Revoking API Access for a Partner Group.....	219
Searching for Assets.....	221
Setting Search Options.....	222
Searching Using Keywords (Simple Search).....	222
Searching Using Keywords and Filters (Advanced Search).....	223
Improving Trading Networks Performance During Searches.....	224
Setting a Limit for Search Results.....	224
Deleting Search Results Periodically.....	225
Migrating Assets.....	227
Overview.....	228
Trading Networks Assets You Can Export.....	228
Migration Methods.....	229
Trading Networks Asset Dependencies.....	229
How Trading Networks Detects Assets in the Target System Before Importing Them.....	232
Migrating Assets Using webMethods Deployer.....	233
Selecting Trading Networks Assets For Export.....	233
Building the Exported Assets.....	234
Deploy the Assets to Other Servers.....	235
Migrating Using My webMethods.....	235
Exporting Assets Using My webMethods.....	235
Importing Assets Using My webMethods.....	236
Migrating Assets Using the tnexport and tnimport Utilities.....	238
tnexport Utility.....	238
Specifying Export Options on the Command Line.....	239
Specifying Export Options in a File.....	240
Examples.....	242
tnimport Utility.....	242
Specifying Import Options on the Command Line.....	242
Specifying Import Options in a File.....	243
Examples.....	245

Archiving and Deleting Documents.....	247
About Archiving and Deleting Documents.....	248
Setting Up Basic Archive and Deletion.....	249
Setting the Number of Days to Retain Documents.....	249
Scheduling Archive and Deletion.....	250
Setting Up Extended Archive and Deletion.....	251
Setting Up Extended Archive and Deletion Using Integration Server.....	251
Setting Archive and Deletion Criteria.....	251
Scheduling Archive and Deletion.....	251
Setting Up Extended Archive and Deletion Using My webMethods.....	252
Setting the Archive and Deletion Criteria.....	252
Scheduling Archive and Deletion.....	254
Setting Up Stored Procedure Archive and Deletion.....	255
Scheduling Archive and Deletion.....	255
Setting Up Stored Procedure Archive and Deletion Using Integration Server.....	257
Setting the Archive and Deletion Criteria.....	257
Scheduling Archive and Deletion.....	257
Using a Partitioned Database for Archival and Deletion.....	257
Database Partitioning.....	259
Overview.....	260
Partitioning Trading Networks Database.....	260
Deleting Run-time Data from a Partitioned Database.....	261
Dropping a table in a Partitioned Database.....	261
Using a Partitioned Database for Archiving Run-time Data.....	261
Deleting Archived Data from a Partitioned Database.....	262
Caching Assets and Query Results.....	263
Overview.....	264
Asset Caching.....	264
Query Results Caching.....	264
Viewing or Modifying System Cache Settings.....	265
Profile Cache Settings.....	266
Profile Summary Cache Settings.....	267
Profile ID Cache Settings.....	269
Document Type Cache Settings.....	270
Document Attribute Cache Settings.....	272
TPA Cache Settings.....	274
TPA ID Cache Settings.....	275
Query Results Cache Settings.....	277
Managing Events.....	279
Overview.....	280
Run-Time Events.....	280
Event Groups.....	286

Configuring Run-time Events.....	288
Working with Dashboards and Charts.....	289
Overview.....	290
Before You Begin.....	290
Dashboards.....	291
Viewing Dashboards.....	291
Charts.....	292
Viewing Charts.....	294
Configuration Properties.....	297
Introduction.....	298
Activity Log Properties.....	298
tn.docType.EnableLogDocumentType.....	298
tn.partnerOnboarding.EnableLog.....	298
tn.procRule.EnableLogProcessingRule.....	298
tn.profile.EnableLogProfileChanges.....	299
tn.queue.EnableLogQueueStartStop.....	299
tn.tpa.EnableLogTPA.....	299
Cluster Properties.....	299
tn.cluster.notifyCacheChange.....	299
tn.cluster.notifyProfileAddUser.....	299
tn.cluster.notifyPropertyChange.....	299
tn.cluster.sync.remoteAliases.....	300
Database Properties.....	300
tn.db.log.level.....	300
tn.db.readonly.....	300
tn.store.encoding.....	300
tn.db.TNLocalDateFormat.....	301
Database Query Properties.....	301
tn.query.maxrows.....	301
tn.query.threshold.....	301
Data Level Security Properties.....	302
tn.dls.EnableLogDataSetChanges.....	302
Debugging Properties.....	302
tn.store.logsql.....	302
Document Archiving Properties.....	302
tn.archive.archiveAfterDays.....	302
tn.archive.batchSize.....	302
tn.archive.batchBackoffTime.....	303
tn.archive.deleteAfterDays.....	303
tn.archive.maxRows.....	303
Document Validation Property.....	303
tn.doc.validate.max_errs.....	303
Document Delivery Properties.....	303
tn.delivery.ftpTransferType.....	303

tn.mail.from.....	303
tn.mail.subject.....	304
tn.delivery.active.transfer.server.aliases.....	304
Event Properties.....	304
tn.eda.events.transactions.....	304
tn.eda.events.completedTransactions.....	304
tn.eda.events.deliveryTasks.....	305
tn.eda.events.serviceExecutionTasks.....	305
tn.eda.events.userInitiatedTasks.....	305
Flat File Property.....	306
tn.ff.contenttypes.....	306
Large Document Handling Properties.....	306
tn.BigDocThreshold.....	306
tn.xml.xqlThreshold.....	307
Localization Properties.....	307
tn.locale.language.....	307
tn.locale.country.....	307
tn.locale.variant.....	307
Performance Properties.....	308
tn.clean.routePipeline.....	308
tn.db.fetchMaxRows.....	308
tn.db.sortDocTimestamp.....	308
tn.doc.saveAsDone.....	308
tn.receive.clearTNOjects.....	309
tn.receive.clearKeys.....	309
tn.store.pool.stmt.....	309
tn.resubmit.return.bizdocs.....	309
tn.tpa.cacheSize.....	310
tn.xml.cacheXQLQueryResults.....	310
Profile Properties.....	310
tn.default.idType.....	310
tn.export.deleted.profiles.....	310
tn.api.gateway.server.alias.....	310
Task Properties.....	311
tn.task.attempt.notPersisted.....	311
tn.task.dbupdate.enableNotifications.....	311
tn.task.dbupdate.retryEnabled.....	311
tn.task.dbupdate.retrySweepTime.....	311
tn.task.init.groupSize.....	312
tn.task.init.pauseLength.....	312
tn.task.maxRetries.....	312
tn.task.ttw.....	312
tn.task.retryFactor.....	312
tn.task.notifyFailure.....	313
tn.task.pauseOnMaxThreads.....	313

tn.task.sweepTime.....	313
tn.task.threadpool.pct.....	314
tn.task.orderingMethod.....	314
Miscellaneous Properties.....	314
tn.chart.fetchMaxRows.....	314
tn.check.ambiguous.docs.....	314
tn.checkUser.ignoreCase.....	315
tn.decimalFormatPattern.....	315
tn.eda.events.emittableDocumentAttributeNameFilter.....	315
tn.ism.stats.saveFrequency.....	316
tn.mail.onuserfail.....	316
tn.portal.rule.preventValidation.....	316
tn.portal.savetodisk.chunksize.....	316
tn.recursion.max.....	316
tn.server.seq.....	317
tn.transport.user.....	317
Server Configuration Parameters.....	319
Introduction.....	320
Server Configuration Parameters.....	320
XQL Reference.....	321
XQL Reference Information.....	322
Sample XML File.....	322
Example XQL Queries.....	323
Extracting Custom Attributes From XML Documents.....	324
Extracting STRINGLIST or STRING ATTRIBUTES.....	325
Extracting NUMBERLIST or NUMBER ATTRIBUTES.....	325
Extracting DATETIMELIST or DATETIME ATTRIBUTES.....	326
Large Document Handling.....	329
Overview.....	330
Sending Large XML Documents for Processing.....	330
Minimum Hardware Requirements.....	331
Configuring Large Document Handling.....	331
Defining Document Types for Large Documents.....	333
Creating Services that Recognize Large Document Handling.....	333
Determining Where the Document Content is Stored.....	333
Retrieving Document Content from Hard Disk Drive Space.....	335
Retrieving the Content for Document Delivery.....	335
Increasing the Size of the Largest Document that Can Be Saved When Using DB2.....	336
Example of User Status in Document Processing.....	337
An Example of User Status while Processing a Document.....	338

About this Guide

This guide explains how to configure webMethods Trading Networks. It explains how to plan, build, and manage a trading network, including how to define and update the information to collect about trading partners and how to add and maintain partners to your trading network. It also explains how to define the manner in which you want Trading Networks to process business documents that are sent to your trading network.

This guide assumes you are familiar with webMethods Integration Server.

Note: The webMethods for Partners product provides the same functionality as Trading Networks. For simplicity, this guide refers only to Trading Networks, but the information applies to both products. The difference between the products is that Trading Networks allows you to have an unlimited number of partners, while webMethods for Partners allows you to have only one partner.

Document Conventions

Convention	Description
Bold	Identifies elements on a screen.
Narrowfont	Identifies storage locations for services on webMethods Integration Server, using the convention <i>folder.subfolder:service</i> .
UPPERCASE	Identifies keyboard keys. Keys you must press simultaneously are joined with a plus sign (+).
<i>Italic</i>	Identifies variables for which you must supply values specific to your own situation or environment. Identifies new terms the first time they occur in the text.
Monospace font	Identifies text you must type or messages displayed by the system.
{ }	Indicates a set of choices from which you must choose one. Type only the information inside the curly braces. Do not type the { } symbols.

Convention	Description
	Separates two mutually exclusive choices in a syntax line. Type one of these choices. Do not type the symbol.
[]	Indicates one or more options. Type only the information inside the square brackets. Do not type the [] symbols.
...	Indicates that you can type multiple options of the same type. Type only the information. Do not type the ellipsis (...).

Online Information

Software AG Documentation Website

You can find documentation on the Software AG Documentation website at <http://documentation.softwareag.com>. The site requires Empower credentials. If you do not have Empower credentials, you must use the TECHcommunity website.

Software AG Empower Product Support Website

You can find product information on the Software AG Empower Product Support website at <https://empower.softwareag.com>.

To submit feature/enhancement requests, get information about product availability, and download products, go to [Products](#).

To get information about fixes and to read early warnings, technical papers, and knowledge base articles, go to the [Knowledge Center](#).

Software AG TECHcommunity

You can find documentation and other technical information on the Software AG TECHcommunity website at <http://techcommunity.softwareag.com>. You can:

- Access product documentation, if you have TECHcommunity credentials. If you do not, you will need to register and specify "Documentation" as an area of interest.
- Access articles, code samples, demos, and tutorials.
- Use the online discussion forums, moderated by Software AG professionals, to ask questions, discuss best practices, and learn how other customers are using Software AG technology.
- Link to external websites that discuss open standards and web technology.

1 Understanding webMethods Trading Networks

■ webMethods Trading Networks	18
■ Architecture	18
■ Partners in a Trading Network	19
■ Understanding the Trading Networks Terminology	21
■ Asset Definition	25
■ Business Process Definition	35
■ Document Delivery	36
■ Document Processing	40
■ Monitoring Transactions	45
■ Security	46
■ Run-Time Event Notifications	61
■ Caching	62
■ Dashboards and Charts	62
■ Database Partitioning	63

webMethods Trading Networks

A *trading network* is a group of organizations that have agreed to exchange business documents. Participants might include strategic partners, buyers, suppliers, and marketplaces (for example, Ariba Network), and are referred to as *trading partners*. Business documents typically include purchase orders, order statuses, purchase order acknowledgements, invoices, and other domain-specific business documents.

webMethods Trading Networks enables your corporation to connect to other organizations to form a business-to-business (B2B) trading network. Trading Networks is a format-neutral, business-document gateway that can recognize and process documents that flow between distributed trading partners. Through Trading Networks, you can exchange business documents with the partners in your network to relay production information. The business documents can be in any format that is recognized by two partners, such as XML and flat file.

Trading Networks is also the base through which webMethods products support numerous eBusiness Standards (eStandards) such as RosettaNet, EDI, ebXML Messaging Service, SWIFT, FIX, and CIDX. webMethods eStandards Modules use features of Trading Networks to perform the processing behavior required for the eStandard supported by the module.

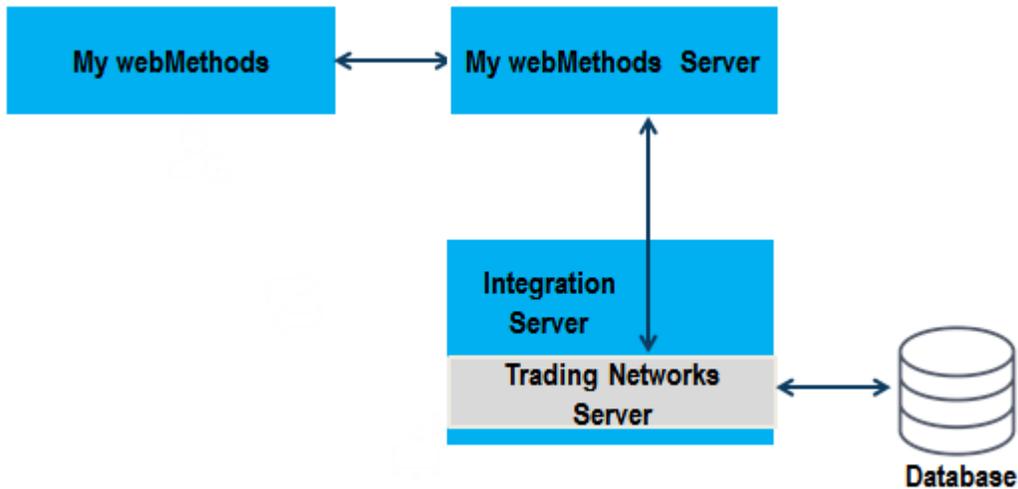
Trading Networks and ActiveTransfer form the primary components of the B2B solution in the webMethods suite of products. Trading Networks provides B2B capabilities such as, partner management, document management, processing rules, TPAs, transaction monitoring, and so on. ActiveTransfer has the ability to send documents to Trading Networks. By using the managed file transfer capabilities of ActiveTransfer, you can leverage the B2B capabilities of Trading Networks to manage and deliver documents to partners securely and efficiently.

Architecture

webMethods Integration Server hosts packages that contain services and related files, and provides an environment for the orderly, efficient, and secure execution of services. By way of the WmTN package, Trading Networks Server manages the partners on your network and the exchange of documents.

My webMethods is a Web-based user interface framework that supports administration and monitoring user interfaces for webMethods products. The Trading Networks user interface in My webMethods lets you perform all Trading Networks tasks. My webMethods Server is the run-time container for functions that webMethods products make available. For example, when a user searches for a Trading Networks asset, My webMethods Server interacts with Trading Networks to perform the search and return the asset to the user.

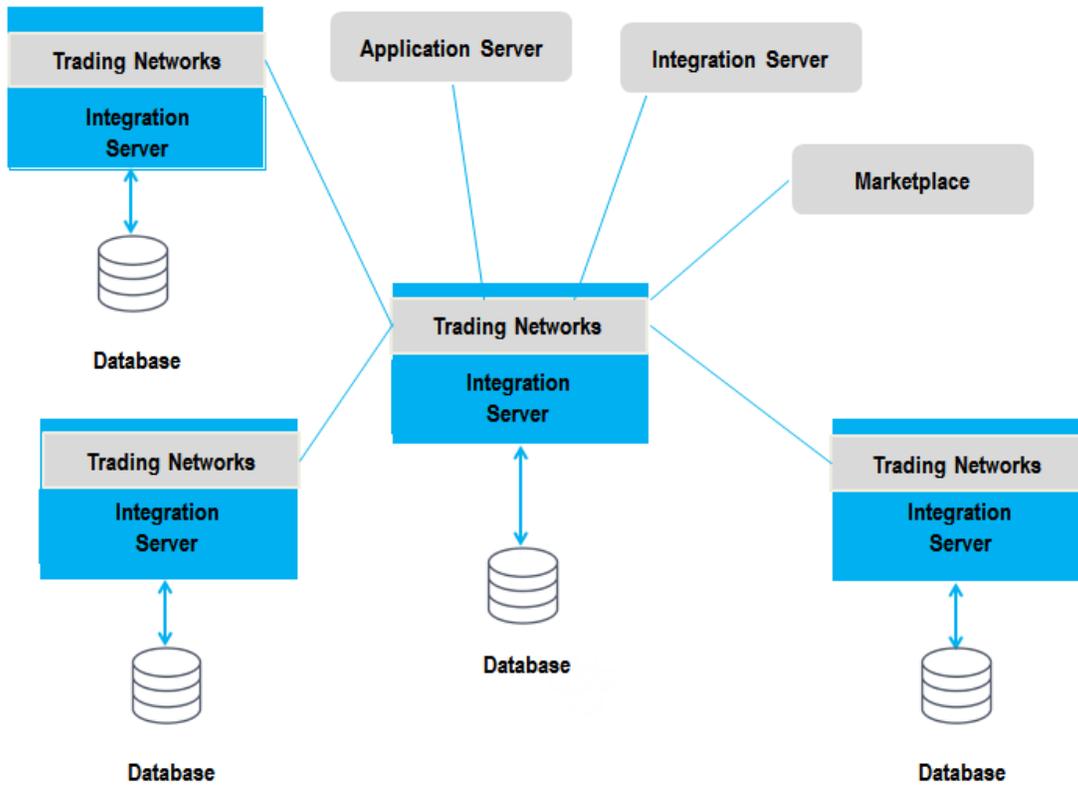
The Trading Networks database stores all information about the trading network, such as partner information, types of documents to process, processing actions, and log activity.



Partners in a Trading Network

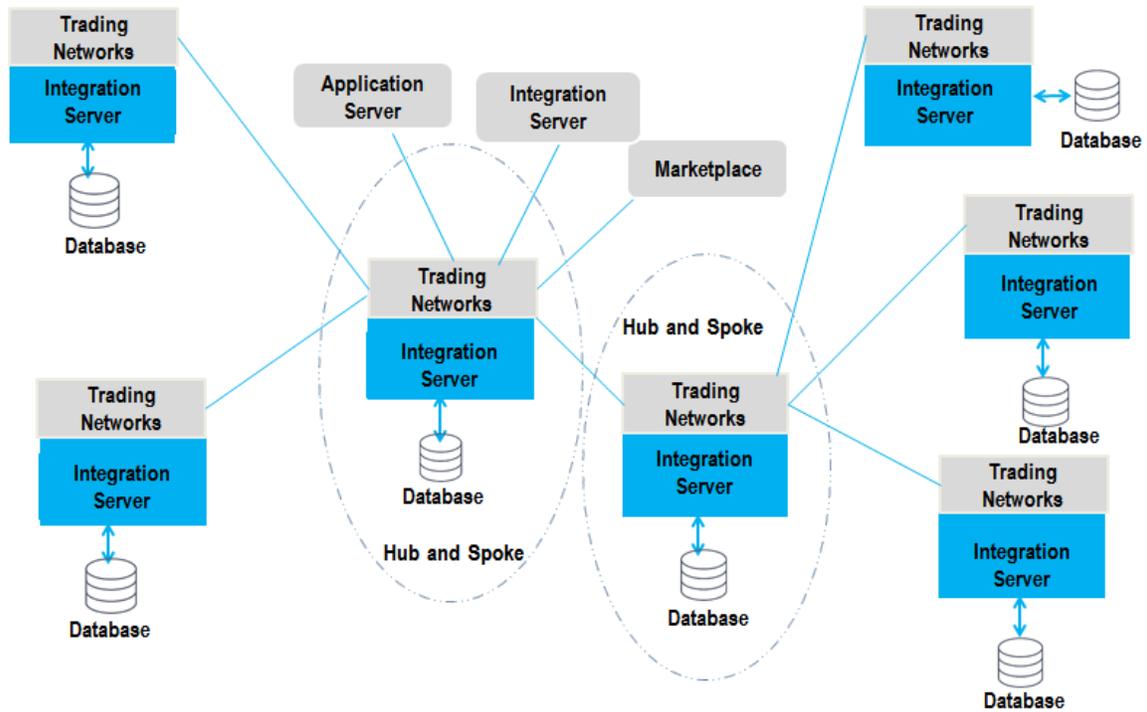
Each of your partners has its own system that communicates with your Trading Networks system. The partner system does not have to use Trading Networks or other Software AG software. When you identify a partner to your Trading Networks system, you provide information about how to connect to and exchange documents with the partner.

The following diagram illustrates various configurations of partners in a trading network.



In this network, one partner system is an Integration Server that is not using Trading Networks. The application server and marketplace partner systems are not using any Software AG software. The partner in the center is referred to as the *hub* of the network. The other partners are referred to as *spokes*. The hub hosts the network and the spokes participate by interacting with the hub.

A Trading Networks partner does not have to be exclusively a hub or a spoke; it can be both, as shown in the following diagram. It can be a hub of its own network and a spoke in another partner's network.



Understanding the Trading Networks Terminology

webMethods Trading Networks is a component that runs on webMethods Integration Server. Trading Networks enables your enterprise to link with other companies (buyers, suppliers, strategic partners) and marketplaces to form a business-to-business network.

The components of Trading Networks are a server and the My webMethods interface. The My webMethods interface is a web-based administration and monitoring user interface for managing your My webMethods components.

Term	Description
Activity Log	A log that Trading Networks maintains to record the activity that occurs within the Trading Networks system. Trading Networks records entries, for example, when you manage trading partner information, when it processes documents, and when you perform administrative tasks.
Business Process	A multi-step interaction among participating systems, people, and trading partners. A business process can be fully automated (involve only interaction among computer systems) or include varying degrees of

Term	Description
	human interaction (for example, review and approval steps). It can be either brief or long-running.
Custom Attribute	A document attribute that you define to identify information within a document that is of interest to you.
Deliver	Sending an outbound document from Trading Networks to the trading partner that is the receiver of the document.
Delivery Method	A method for delivering a document to a trading partner. For example HTTP, HTTPS, FTP, FTPS, e-mail (SMTP), SFTP, Web Service. Trading Networks supports, immediate delivery methods, scheduled delivery methods, receiver's preferred delivery method, and queue for polling.
Delivery Task	A task that Trading Networks establishes to keep track of the attempts to re-deliver a document when it is using reliable delivery.
Document	A business document (For example, purchase order, acknowledgement, confirmation) sent to Trading Networks. The document can be in any format (XML, EDI, etc.) Trading Networks provides out-of-the-box support for XML and flat file documents. The webMethods EDI Module is necessary for EDI documents.
Document Attribute	A Trading Networks object that defines a piece of information within a document that is of interest. For example, document attributes in a purchase order might be the purchase order number, the account number of the purchase order and the total purchase amount. Document attributes can be either a system attributes (those that are provided with Trading Networks) or custom attributes (those that you define for your enterprise).
Document ID	A system attribute for an identifier in a document that is typically a unique value that distinguishes a document from other versions of the same document.

Term	Description
Enterprise Partner	The partner that hosts the trading network. On your Trading Networks system, this would typically be your corporation. (Also known as the hub, local partner, or sponsor.)
External ID	The value of the external ID type within a document. For example, if the external ID type is a D-U-N-S number, the external ID is the actual value of the D-U-N-S number.
Flat File	Any file or document that has a format that is non-describing, that is, a document that does not contain metadata. A flat file document presents hierarchical data in a record-based storage format, which unlike XML, does not embed structural information within the data.
Immediate Delivery Method	A delivery method where Trading Networks attempts to immediately deliver a document directly to the receiving partner. You can create immediate delivery methods using all the supported delivery methods.
Local Partner	The enterprise partner that hosts Trading Networks. (Also known as the enterprise partner, hub or sponsor.)
Private Queue	A scheduled delivery queue that you define to schedule the delivery of documents that are aimed at one specific trading partner. You define a private queue in the profile of the partner to receive the documents.
Processing Rule	A Trading Networks object that contains a set of actions that determine how Trading Networks is to process an inbound document and criteria that indicates when to select a processing rule for an incoming document.
Profile	A Trading Networks object that contains a summary of information about a corporation that is part of a trading network. A profile contains standard fields that Trading Networks provides and extended fields that are site-defined.

Term	Description
Public Queue	A scheduled delivery queue that you define to schedule the delivery of documents that are aimed at multiple trading partners.
Reliable Delivery	A feature of Trading Networks where Trading Networks attempts to re-deliver a document to a trading partner one or more times if previous attempts to deliver the document fails. For an immediate delivery method, Trading Networks automatically uses reliable delivery when the pre-processing action Save Document to Database indicates that Trading Networks is to save the document content to its database. For a scheduled delivery method, Trading Networks always uses reliable delivery.
Scheduled Delivery Method	A delivery method where Trading Networks batches multiple documents in a scheduled delivery queue. The documents in the queue are acted on at scheduled times to deliver them.
Trading Networks Document Type	A Trading Networks object that defines how Trading Networks is to recognize a document and initial actions to take on a recognized document. Trading Networks recognizes the document by using identification information in the TN document type. The actions specified in aTN document type indicate the document attributes that Trading Networks is to extract from the document (including information about XML namespaces the documents might use) and specify options for pre-processing the document (which include verification, validation, and whether to save the document attributes, document content, and log entries for the document to the database).
Trading Partner Agreement	A Trading Networks object that you can use to tailor how documents are exchanged between two trading partners.
Trading Partner	A trading partner may be an organization in your trading network, for example, a strategic partner, marketplaces, buyer, or supplier. Each trading partner requires a profile. You can exchange business documents with the trading partners in your network to relay mission critical production information.

Term	Description
Transaction	The documents that have passed through Trading Networks.
Unknown Document	A document that does not match any Trading Networks document type.
Unknown Partner	A trading partner (sender or receiver) of a document is considered unknown if Trading Networks is unable to determine the sender or receiver; that is match the sender or receiver to a profile in the Trading Networks system.
User Status	A system attribute that contains a status that a user can associate with a document. For example, "Needs Approval".

Asset Definition

When Trading Networks receives a document, it performs run-time processing. You define this processing by designing Trading Networks assets, as follows:

Assets	Description
Document attributes	Identify the pieces of document content you need to process documents. For example, you might be interested in document senders and receivers, or the total amounts of purchase orders.
Document types	<p>Define document types for documents you and your trading partners will exchange. Document types are definitions that represent particular categories of documents (for example, XML or flat file). The document type can represent an industry standard, such as a cXML Purchase Order, FIXML Quote Request, or Biztalk Envelope, or a custom standard, such as a purchase order format that you and a partner have agreed on.</p> <p>Document types can also specify actions to perform for documents, such as saving documents to the Trading Networks database.</p>
Processing rules	Specify actions to perform for documents, such as delivering documents to partners.

Assets	Description
Profiles	Identify your corporation and the corporations of the partners in your trading network, and specify how to connect to each other and exchange documents.
Trading Partner Agreements (TPAs)	Specify transaction-dependent information that is specific to a group of transactions between two trading partners. A <i>transaction</i> is the passage of a document through Trading Networks.
Note:	You define assets differently for large documents. For details, see " Large Document Handling " on page 329.

Document Attributes

Trading Networks supports two types of document attributes: *system attributes* and *custom attributes*. Trading Networks defines the system attributes as follows:

System Attribute	Description
SenderID	Identifier for the sender of a document.
ReceiverID	Identifier for the receiver of a document.
DocumentID	Unique identifier for a document.
UserStatus	Status that you or a partner assign to a document (for example, Needs Approval).
GroupID	Identifier that associates a document with other documents in a group. Grouping documents is helpful for end users doing document searches.
ConversationID	Identifier that associates a document with other documents that are processed by a business process (also called a <i>conversation</i> of documents). For information about business processes, see " Business Process Definition " on page 35. This identifier is only present if you are using BPM. You need to extract it from the document and add it to the document type. The transaction with the ConversationID then gets attached to the business processes.

System Attribute	Description
SignedBody	For XML documents, data that was digitally signed to create the digital signature for the document.
Signature	For XML documents, digital signature of the document.

You define *custom document attributes*. For example, if you are interested in PO numbers and total order amounts, you might define PO_Number and Total_Order_Amount custom attributes.

You tell Trading Networks to extract document attributes from documents you receive for these reasons:

- To use extracted attributes as a criterion for using a particular processing rule. For example, if you want to use one processing rule if the sender is Partner A and another processing rule if the sender is Partner B, you would extract the system attribute SenderID. Or if you want to use a particular processing rule when the receiver is Partner C and the total order amount is greater than \$10,000, you would extract the system attribute ReceiverID and the custom attribute Total_Order_Amount.
- To perform certain processing actions that require extracted attributes. For example, if you want to deliver a document to the receiver partner, you would extract the system attribute ReceiverID. If you want to verify the digital signature of an XML document, you would extract the system attributes SignedBody and Signature.
- To search for documents in My webMethods based on extracted attributes. For example, if you want to be able to find documents that were sent by Partner A for which the total order amount is greater than \$10,000, you would extract the system attribute SenderID and the custom attribute Total_Order_Amount.
- If you BAM-enable Trading Networks and you want to pass extracted attributes to Optimize for B2B for analysis and monitoring. For example, if you want to generate a report on the purchase order quantity for a particular sender from a particular receiver, you would extract the custom attribute PO_Quantity and the system attributes SenderID and ReceiverID.

Document Types

You can define document types for XML and flat file documents. For example, you might define document types for documents that represent purchase orders or acknowledgements. For XML documents, you define document types for variations of documents; for example, you might define purchase orders in cXML, OAG, and CBL.

A document type specifies criteria that an inbound document must meet to match the document type, document attributes to extract from documents that match, and

processing to perform for those documents. You define exactly one document type for each document you expect to receive or send.

If webMethods Module for EDI is installed on the same Integration Server that hosts Trading Networks, you can also define document types for EDI documents. For instructions, see *webMethods Module for EDI Installation and User's Guide*.

The explanation of document types in this section mentions the *pipeline*. The pipeline is an Integration Server data structure in which input and output values from services are maintained. The pipeline starts with the input to a service and collects inputs and outputs from subsequent services. When a service executes, it has access to all data in the pipeline.

XML Document Types

An XML document type defines the following:

Definition	Description
Document recognition criteria	<p>Content that an inbound XML document must have to be a match for the XML document type. You can specify one or more of the following:</p> <ul style="list-style-type: none"> ■ Root tag the document must have. ■ System or public identifier from the DOCTYPE declaration the document must have. ■ Nodes and, optionally, values the document must have. You define XML Query Language (XQL) queries that Trading Networks executes to find the nodes and values. ■ Pipeline variables (that is, variables inserted into the pipeline by the service that sends the document to Trading Networks) and, optionally, values the variables must have.
Attribute extraction	<p>System and custom attributes to extract from the document. You define XQL queries that Trading Networks executes to extract the attributes. You can also specify transformations for extracted attributes; for example, you might want to transform an extracted string value into all uppercase characters.</p>
Pre-processing actions	<p>A document type can specify one or more of these actions:</p> <ul style="list-style-type: none"> ■ Format as IS Document Type action, to transform the document into an IS document that can be parsed into an IData object. An IData object is a collection of name/value pairs on which services can operate. An IData object can contain any number of elements of any valid Java objects, including additional IData objects.

Definition	Description
	<ul style="list-style-type: none"> <li data-bbox="464 323 1321 491">■ Verify Digital Signature action, to make sure the signed body of the document arrived unchanged and the sender is who it claims to be. Trading Networks checks the sender by matching the certificate from the digital signature to the certificate in the partner's profile. <li data-bbox="464 512 1230 575">■ Validate Structure action, to validate the structure of the document against a specified schema. <li data-bbox="464 596 1354 806">■ Check for Duplicate Document action, to check whether Trading Networks has already received the document. Trading Networks saves the results of the check to the pipeline, so you can use the results in the Save Document to Database pre-processing action (below) to save only unique documents, and in processing rule actions. <li data-bbox="464 827 1354 953">■ Save Document to Database, to save the document to the Trading Networks database. For example, you save documents to the database when you want to make multiple attempts to deliver documents to partners. <li data-bbox="464 974 1321 1184">■ Indicate whether to use a processing rule to further process the document. You might not want to use a processing rule if you want to simply persist the document to the Trading Networks database, or to process the document using a business process instead of a processing rule. For information about business processes, see "Business Process Definition" on page 35.

Flat File Document Types

Flat file documents present complex hierarchical data in a record-based storage format which, unlike XML, does not embed structural information within the data. The Trading Networks definition of a flat file is any file or document whose format is non-describing; that is, a document that does not contain metadata. In other words, flat file data is externalized as a set of records (list of records containing fields and composites) without any structural information. Since the records are not structured in the document, the application receiving the flat file must know the structure of the flat file to read its content.

Since Trading Networks does not know the structure of a flat file document, it cannot extract values for attributes directly from documents. The entry points for flat files into Trading Networks, therefore, are *document gateway services*; that is, your trading partners send their flat files to document gateway services rather than directly to Trading Networks. A document gateway service does the following:

- Reads an inbound flat file document

- Places output variables such as the flat file document type for the document; the processing rule for the document; system document attributes such as SenderID, ReceiverID, and UserStatus; and custom attributes in the pipeline
- Passes control to Trading Networks

A flat file document type defines the following:

Definition	Description
Document recognition criteria	Content that an inbound flat file document must have to be a match for the document type. You can specify any of the variables the document gateway service placed in the pipeline, such as the name of the flat file document type to use; the SenderID, ReceiverID, or UserStatus system attributes; or custom attributes.
Attribute extraction	System and custom attributes to extract from the document. You can also specify transformations for extracted attributes; for example, you might want to transform an extracted string value into all uppercase characters.
Pre-processing actions	Same as for XML documents, except that the Format as IS Document Type action is not available for flat file documents (see "XML Document Types" on page 28).

Processing Rules

A processing rule specifies criteria that a document must meet to match the rule, and processing to perform for documents that match, as follows:

Definition	Description
Matching criteria	<p>Content that a document must have to match the processing rule. You can specify the following as criteria:</p> <ul style="list-style-type: none"> ■ Sender and receiver the document must have ■ Document type for the document ■ User status the document must have (for example, Needs Approval) ■ Whether errors were or were not encountered during document recognition or attribute extraction ■ Custom attributes the document must have (for example, Total_Order_Amount custom attribute value greater than \$10,000)

Definition	Description
Pre-processing actions	<p>The pre-processing actions you can define in a processing rule are the same as those you can define for document types (see "XML Document Types" on page 28). The difference is that you must specify what to do for each action. A processing rule can specify to:</p> <ul style="list-style-type: none"> ■ Perform the action that is defined in the document type. ■ Perform the action that you define in the rule. ■ Never perform the action at all.
Processing actions	<p>A document type can specify one or more of these actions:</p> <ul style="list-style-type: none"> ■ Execute a Service action, to execute a custom service. For example, you could execute a service to send the document to a back-end system for processing or to update data in an internal system with extracted attributes. You can use data that is in the pipeline in the service; for example, you could perform the Check for Duplicate Document pre-processing action and then execute one service if the document is a duplicate and another service if the document is unique. Trading Networks places the service results in the pipeline so you can use them in output templates or other processing actions. ■ Alert Email Message action, to send a message to a contact in the sender's or receiver's profile, the webMethods system administrator, or a specified email address when document processing is complete. For example, if the document is a purchase order, you could send a message to alert the person who approves purchase orders. You can include pipeline data in the message (for example, the document type for the document). ■ Change User Status, to assign a user status to the document for use when searching for documents or generating reports. For example, if the document is a purchase order, you could set the document's user status to Needs Approval. The person who approves purchase orders can search for documents whose user status is NeedsApproval. The user status is assigned to the User Status system attribute. ■ Deliver Document By action, to deliver the document to the receiving partner using a specified delivery method. ■ Respond action, to send a message back to the document sender (for example, an acknowledgement for a purchase order). You can include data from the pipeline in the message. For example, you could have the Execute a Service action asynchronously invoke a service that performs the appropriate processing for the purchase order, then return confirmation that the purchase order

Definition	Description
	was received and processed. You could then have the Respond With action return an acknowledgement that indicates that the purchase order was received.

Profiles

You add partners to your trading network by defining *profiles*. A profile contains information about a corporation that is a trading partner in your network. You define a profile for your corporation (called the Enterprise profile), and for each of the other corporations that are trading partners in your network.

Overview of Creating Partner Profiles

Trading Networks administrators can create partner profiles in two ways:

- Partner Administration page in My webMethods. The administrator obtains details about the partner from the business user who identified the partner as a prospective trading partner. The administrator then enters that information on various My webMethods pages on behalf of the partner. For information about creating a profile using this method, see ["Creating Profiles" on page 185](#).
- Partner onboarding process. For more information, see ["Overview of the Partner Onboarding Process" on page 33](#).

Trading Networks provides standard fields for profiles. You supply these types of information in standard fields:

- General information about the corporation, such as corporation name, address, and contacts
- Information about how to deliver documents to the corporation
- Certificate information for digital signing of documents, verifying digital signatures, encrypting and decrypting documents

You define *external ID types* for the types of identification used in documents that are exchanged in your network, and then specify the actual *external IDs* in profiles. For example, you could define the external ID type DUNS, and then in profiles, you could specify external IDs that contain the actual values of the D-U-N-S numbers in documents. Trading Networks uses external IDs to identify the sending and receiving partners within a document. For example, Trading Networks compares the value of the D-U-N-S number in a document to the external IDs in profiles. If it finds a profile with an external ID that matches the D-U-N-S number, it determines that the sender of the document is the partner with that profile. You can specify an unlimited number of external ID types. Each profile must include at least one external ID.

If you want to maintain additional information about your partners, you can define custom fields, called *extended fields*. For example, you might want to define extended fields for preferred shipping method, cost centers, or customer codes.

By default, many standard profile fields are required. When a profile field is required, all profiles on your system must have a value for it. You can make standard fields required or not required, and you can make extended profile fields required or not required.

You can associate profiles with *partner groups*. For example, you can create a partner group named Buyers to associate all of your partners who are buyers in your trading network. A partner can belong to more than one partner group. You can use partner groups in the sender or receiver criterion when matching processing rules to documents.

Overview of the Partner Onboarding Process

The partner onboarding process can help you automate the creation of partner profiles and trading partner agreements (TPAs). The administrator obtains details about the partner either from the business user who identified the partner as a prospective trading partner (by way of a comma-separated values, or .csv, file) or directly from the partner (by way of a questionnaire). Trading Networks automatically creates the partner's profile using the details provided in the spreadsheet or questionnaire.

Following is a high-level overview of the partner onboarding process:

1. The administrator uploads a spreadsheet into My webMethods that contains basic information about one or more partners. At a minimum, the spreadsheet can contain the partner's name and email address. The spreadsheet can also contain other standard profile, TPA, and document type information, such as partner address, preferred delivery method, external ID, and profile extended fields. The administrator can download the sample template provided, enter the information, save it, and upload the file.
2. Trading Networks creates the beginnings of a partner profile and TPA using the information in the spreadsheet. In My webMethods, Trading Networks sets the status of the partner to New (Pending Partners page) and the status of the TPA to Agreed (Trading Partner Agreements page).
3. If additional details are needed from the partner (for example, if the spreadsheet contained only the partner name and email address), the administrator does the following:
 - a. Creates a questionnaire template that identifies the additional information needed
 - b. Assigns the template to the partner
 - c. Configures settings for an invitation to be emailed to the partner
 - d. Grants the partner permission to access the partner onboarding pages in My webMethods
 - e. Sends the invitation email to the partner requesting him or her to log in to My webMethods to complete the questionnaire
4. Trading Networks sets the partner's status to Invited.

5. The partner completes the questionnaire. Trading Networks sets the partner's status to Responded.

Note: If the partner does not complete the questionnaire within a configured time limit, Trading Networks sets the partner's status to Expired. The administrator can restart the onboarding process for that partner by re-inviting the partner.

6. The administrator checks the details that the partner has submitted. If they are complete and in order, the administrator approves the partner. If the details are not in order, the administrator resends the questionnaire for the partner to enter again.
7. If the partner enters the questionnaire correctly this time, the administrator approves the partner. If not, the administrator rejects the partner and Trading Networks sets the partner's status to Rejected.
8. If the partner is approved, Trading Networks adds the information from the questionnaire to the partner's profile and TPA. The profile is changed from Pending to Approved.

For procedures related to the partner onboarding process, see ["Onboarding New Partners" on page 173](#).

TPAs

You can define trading partner agreements (TPAs) for pairs of partners. Each TPA contains specific information about how documents should be exchanged between two trading partners, as follows:

- The partner that represents the sender of the documents.
- The partner that represents the receiver of the documents.
- An agreement ID that identifies the type of TPA (for example, TPAs for the webMethods Module for EDI use the agreement ID EDITPA).
- The *TPA data* that contains the application-specific variables to use to tailor the processing of documents exchanged between the sender and receiver. You specify this data by defining an *IS document type*. An IS document type is an element in the Integration Server namespace that contains a set of fields that define the structure and type of data in an IS document, or IData object. For example, the webMethods Module for EDI ships with an IS document type (the `wm.b2b.edi.tn.TPA:EDITPA` IS document type) to use for TPAs for partners exchanging EDI documents. This IS document type contains a set of variables that are used for processing EDI documents.
- Optionally, an initialization service you supply to initialize the TPA data (for example, the webMethods Module for EDI supplies an initialization service to set the TPA values to its default values).

- Optionally, a validation service you supply to validate the data added to the IS document for the TPA.
- Optionally, an export service you supply to export the TPA data to an industry-standard format.

Trading Networks does not use TPAs for its own processing; rather, the data you specify in the TPA are available for your own use. For example, you can access the TPA information from services that are executed by a processing rule. Access to this information allows you to build a document exchange application that uses the TPA to tailor the exchange of documents between partners. Other webMethods products take advantage of the TPA feature in Trading Networks. For example, the webMethods ebXML Module uses the TPA feature to support ebXML Collaboration Protocol Agreements (CPAs).

The set of parameters can be different for different types of TPAs. For example, you might use TPAs for partners that exchange documents using ebXML that contain the parameters defined by the webMethods ebXML Module. Other partners might exchange documents using EDI, and for those partners you create TPAs that contain parameters defined by the webMethods Module for EDI.

The type of information a TPA contains is different than the type of information that Trading Networks maintains in profiles. A profile contains information about the partner that does not vary with each document being exchanged, such as company name and address, certificate information, delivery protocol parameters, and external IDs. TPAs are intended to contain transaction-dependent information (for example, configuration information to support specific types of documents being exchanged) that are specific to a group of transactions between the two trading partners (for example, digital signature or encryption to a message). TPAs augments profiles and offer a flexible way to process and manage the documents exchanged between two trading partners.

When you define a TPA, you set the agreement status. This status indicates the status of the TPA agreement between the receiver and sender. The status can be proposed, agreed, or disabled. webMethods products that use the TPA feature recognize these statuses. For example, if webMethods ebXML Module tries to use a TPA whose status is disabled, it acts as if there is no TPA. If you create an application that uses TPAs, it should check and honor the disabled status.

Business Process Definition

For some documents, you might require multi-step processing that involves interaction among systems, people, and trading partners. An example of such a processing is the fulfillment of a purchase order that includes a purchase order document, human interaction to determine whether to approve the purchase order, and either an order acknowledgement (ACK) document or an order negative acknowledgement (NACK) document. For such processing, you can define a *business process*. You can use a business process instead of or in addition to a processing rule.

You define business processes in Software AG Designer. Designer generates the run-time elements needed to run the business processes (for example, triggers and flow services) in Integration Server. The Process Engine, which runs on Integration Server, manages the execution of business processes.

To handle documents sent by Trading Networks, you must define the business process to handle a *conversation* of related documents that all contain the same *conversation ID*. The conversation ID is the Trading Networks ConversationID system attribute. When you want to process documents using a business process, the document types for the document must tell Trading Networks to extract this attribute. When this attribute has been extracted, Trading Networks knows to pass the document to the Process Engine after Trading Networks processing is complete.

For complete information on business processes, see the webMethods BPM documentation.

Document Delivery

To deliver a document to a receiving partner, you create a processing rule for the document type, and you define the Deliver Document By processing action in the processing rule. On the action, you identify the type of delivery to use, as follows:

Delivery Type	Description
Immediate delivery	Delivers the document directly to the receiving partner. Trading Networks tries to deliver the document only once unless you use <i>reliable delivery</i> to make repeated attempts.
Scheduled delivery	Places the document in a <i>queue</i> you define and delivers the document to the receiving partner at scheduled times. You use scheduled delivery when it is more efficient to deliver a batch of documents than to deliver each document as it arrives. For example, if you want to deliver documents through FTP, you might use scheduled delivery so you can open a connection, deliver the batch of documents, and then close the connection, rather than use immediate delivery, which requires you to open and close a connection for each document. Trading Networks tries to deliver the document only once unless you use <i>reliable delivery</i> to make repeated attempts.
Receiver's Preferred Protocol	A preferred protocol for receiving documents is set up when a partner profile is created.

Delivery Type	Description
Queue for polling	Places the document in the Trading Networks-provided queue. When a partner polls for documents, Trading Networks delivers all documents in the queue for which the partner is the receiver. Queue for polling is the default delivery method. If you do not set up immediate or scheduled delivery for a partner, Trading Networks uses queue for polling for that partner.

Immediate Delivery

Immediate delivery uses *immediate delivery methods*, *immediate delivery services*, and, if you use reliable delivery, *delivery tasks*. An immediate delivery method invokes an immediate delivery service that delivers a single document to the receiving partner. A delivery task is a task Trading Networks creates to keep track of its attempts to redeliver a document when you are using reliable delivery. The options for immediate delivery are as follows:

- Trading Networks provides standard FTP, FTPS, HTTP, HTTPS, SFTP, SMTP, and Web service delivery methods. You can create custom immediate delivery methods from these methods; you create the custom method in one profile, and then you can select the method in other profiles. When you create a custom method from a standard method, Trading Networks automatically creates the corresponding immediate delivery service. If that service does not meet your needs, you can create a custom service to replace it.
- Trading Networks provides built-in immediate delivery methods for primary and secondary email, FTP, FTPS, SFTP, HTTP, and HTTPS, and provides the corresponding built-in delivery services. If the services do not meet your needs, you can create custom services to replace them.
- You can create custom immediate delivery services. For example, you might want to create a custom service that delivers a message into a message queuing system. After you create a custom service, you register it with Trading Networks, and Trading Networks automatically creates the corresponding immediate delivery method.
- Trading Networks provides ActiveTransfer as a delivery method. By using the managed file transfer capabilities of ActiveTransfer, you can leverage the B2B capabilities of Trading Networks to manage and deliver documents to partners securely and efficiently as follows:
 - Manage partners in Trading Networks.
 - Set up complex endpoints (Virtual Folder System or VFS) for partners in ActiveTransfer.
 - Configure Trading Networks as a central partner management to use ActiveTransfer Server depending on whether ActiveTransfer is installed on a local, remote, or clustered environment, as the preferred document delivery

method. For more information, see ["Configuring Alias for ActiveTransfer on Remote Server" on page 84](#).

- Select an ActiveTransfer VFS path for the partners. Using the *VFSPath*, *senderPartnerID*, *receiverPartnerID*, and *VFSID* parameters, ActiveTransfer sends the BizDocument to the partners.
- View the status of the document delivery and the delivery details from ActiveTransfer on the **Transactions** and **Activity Log** pages.

Note: Trading Networks does not support the delivery of large documents to ActiveTransfer on a remote **Integration Server** instance. For more information about handling large documents in Trading Networks, see ["Large Document Handling" on page 329](#).

To use immediate delivery for a partner, define one or more immediate delivery methods in the partner's profile. You can then specify the method to use for a document on the Deliver Document By action.

You can designate one immediate delivery method as the partner's preferred protocol, and specify using the partner's preferred protocol for a document on the Deliver Document By action. The preferred protocol option is primarily for use with webMethods eStandards Modules.

To use reliable delivery, you specify the Save Document to Database pre-processing action on the processing rule. In the receiving partner's profile, you specify the number of times to try to deliver a document to the partner and the amount of time to wait between attempts. If Trading Networks retries the maximum number of times and delivery is unsuccessful, it changes the delivery task status to FAILED. If you correct the problem, you can restart the task and Trading Networks will start trying to deliver the document again.

Scheduled Delivery

Scheduled delivery for a partner uses a *scheduled delivery queue*, a *scheduled delivery service*, a *delivery schedule*, and delivery tasks.

There are two types of scheduled delivery queues:

- Private queues, each with a delivery schedule and delivery service you define. Each private queue delivers documents to only one partner.
- Public queues, each with its own delivery schedule and delivery service you define. Each public queue can deliver documents to multiple partners.

To use scheduled delivery for a partner, you indicate which queue to use in the partner's profile; that is, you specify the partner's private queue or one of the public queues. You then specify the queue to use for a document on the Deliver Document By processing action.

Trading Networks automatically uses reliable delivery with scheduled delivery. Trading Networks creates a delivery task that invokes the scheduled delivery service and places

the task in the specified queue, and that keeps track of Trading Networks's attempts to redeliver the document.

A scheduled delivery service delivers a batch of documents to the receiving partner. You can associate the same scheduled delivery service with multiple queues. The options for scheduled delivery services are as follows:

- Trading Networks provides one built-in scheduled delivery service that uses FTP to deliver documents to a single destination. The service opens a connection, delivers all documents, and then closes the connection.
- You can create scheduled delivery services and register them with Trading Networks.

At the times the queue schedule indicates, Trading Networks invokes the scheduled delivery service to deliver the documents in the queue. The service acts on all delivery tasks in the queue whose status is QUEUED. These QUEUED delivery tasks include all delivery tasks that are in the queue when the service is invoked and any new tasks that are added to the queue before the delivery service ends. The scheduled delivery service tries to deliver the document to the receiving partner and indicates whether the delivery was successful or not. The status of the delivery task is updated accordingly. If a document cannot be delivered, Trading Networks leaves the delivery task with the QUEUED status so the scheduled delivery service will try to deliver the document again at the next scheduled time.

Trading Networks automatically uses reliable delivery with scheduled delivery, but you must specify the Save Document to Database pre-processing action on the processing rule. In the receiving partner's profile, you specify the number of times to try to redeliver a document to the partner. If Trading Networks retries the maximum number of times and delivery is unsuccessful, it changes the task status to FAILED.

Queue Documents for Polling

Trading Networks uses queue for polling in these cases:

- When you set the Deliver Document By processing action for a document to queue for polling.
- When you set the Deliver Document By processing action for a document to Receiver's Preferred Protocol, and the receiving partner's preferred protocol is queue for polling.
- When it cannot find sufficient delivery information in the receiving partner's profile (for example, if the Deliver Document By processing action indicates using an immediate delivery method, but the profile does not contain delivery information for that method).

When using queue for polling, Trading Networks saves the document to its database and sets the document's processing status to POLLABLE. A receiving partner looks in your profile on its system to find the polling method to use and how often to poll. The receiving partner then polls for its documents, and your Trading Networks delivers them. The receiving partner returns a status for each document that indicates whether

the document was accepted or accepted with errors. Your Trading Networks updates the processing status of the document on your Trading Networks, setting it to ACCEPTED or ACCEPTED W/ ERRORS.

Document Processing

1. A partner sends a document, as follows:
 - The partner sends an XML document to the Trading Networkswm.tn:receive service. The service reads the document into the pipeline. If coded to do so, the service places the name of the document type and processing rule to use for the document on variables in the pipeline.
 - The partner sends a flat file document to a document gateway service. The gateway service reads the document and places variables you specify (for example, SenderID, ReceiverID, DocumentID, and ConversationID) in the pipeline. If coded to do so, the gateway service places the name of the document type and processing rule to use for the document on variables in the pipeline, and places specified custom document attributes and their values in the pipeline. The gateway service then invokes the wm.tn:receive service.
2. If Trading Networks does not find a document type name in the pipeline, it performs *document recognition*; that is, it compares the XML document to the criteria specified in enabled XML document types, or the flat file document to the criteria specified in enabled flat file document types. When it finds a match, Trading Networks creates a BizDocEnvelope from the document and places the BizDocEnvelope in the bizdoc variable in the pipeline. It then extracts the document attributes that are specified in the matching document type from the pipeline and adds them to the BizDocEnvelope, along with additional information required to route and process the document. The BizDocEnvelope represents a routeable Trading Networks transaction and conforms to the wm.tn.rec:BizDocEnvelope. The bizdoc variable adheres to the wm.tn.rec:BizDocEnvelope IS document type and is an instance of com.wm.app.tn.doc.BizDocEnvelope.

Trading Networks also places the sender and receiver variables in the pipeline. These variables contain the partners identified as the sender and receiver, respectively, in the document. The variables adhere to the wm.tn.rec:ProfileSummary IS document type and are instances of com.wm.app.tn.profile.ProfileSummary. To view the IS document type, open Software AG Designer and look in the WmTN package wm.tn.rec folder.
3. The next action depends on whether the pipeline contains a processing rule name and what the document type specifies, as follows:
 - If the document type specifies the use of a processing rule, Trading Networks continues with step 4.
 - If the pipeline specifies a processing rule name, Trading Networks skips to step 5.

- If the pipeline does not contain a processing rule name, and the document type specifies to *not* use a processing rule, Trading Networks performs any pre-processing actions that are specified in the document type in the order they are listed in "[Document Types](#)" on page 27. Trading Networks then skips to step 7.
4. Trading Networks compares the document to the criteria specified in all enabled processing rules, as follows:

If the rule specifies this criterion...	Trading Networks does this...
Sender, receiver, or both	Compares to the SenderID and ReceiverID system attributes in the document. If Trading Networks matches these SenderID and ReceiverID system attributes to a partner profile, it uses the corporation name and unit name specified in the profile as the name of the sender, the receiver, or both. Trading Networks compares the name of the partner to the partner or partners specified in the processing rule.
Document type	Compares to the document type for the document.
User status	Compares to the UserStatus system attribute for the document.
Occurrence or non-occurrence of errors during document recognition	Checks for errors in the pipeline.
Custom attributes	Compares to the attributes it extracted from the document.

5. Trading Networks performs the pre-processing actions specified by the matching processing rule in the order they are listed in "[Processing Rules](#)" on page 30.
6. Trading Networks performs the processing actions specified in the processing rule in the order they are listed below.

If this action is specified...	Trading Networks does this...
Execute a Service	Processing depends on how you invoke the action, as follows:

If this action is specified...	Trading Networks does this...
	<ul style="list-style-type: none"> <li data-bbox="537 359 1360 636">■ Synchronous: Invokes the service once, waits for the service to complete, and places the results in the pipeline. You can use the service results in output templates or other processing actions. If you do not use the Respond With action, Trading Networks returns the results of the service unmodified to the client that sent the document. If you do use that action, Trading Networks returns only the message you specify. <li data-bbox="537 653 1360 930">■ Asynchronous: Invokes the service once and does not place the results in the pipeline. If there are subsequent processing actions, Trading Networks immediately performs them; if there are not, it returns to the caller that sent the document for processing. If you do not use the Respond With action, Trading Networks returns no information to the caller. If you do use that action, Trading Networks returns the message you specify. <li data-bbox="537 947 1360 1224">■ Service execution task: Invokes the service and uses reliable execution to retry the service if it fails. Trading Networks creates a <i>service execution task</i> to keep track of its attempts to execute the service. If Trading Networks retries the maximum number of times and the service is still unsuccessful, it updates the task status to FAILED. The processing and information returned to the caller are the same as for Asynchronous.
Send an Alert Email	Sends the specified email to the specified person.
Change User Status	Assigns the specified user status to the document. For an example of using user status criteria during processing, see "Example of User Status in Document Processing" on page 337.
Deliver Document By	Delivers the document to the receiver that is identified in the document using the delivery method you specified in the processing rule.
Respond with	Returns the specified message to the caller that sent the document to be processed.

7. If you extracted the ConversationID system attribute from the document, Trading Networks passes the document to the Process Engine for processing by a business

process. If the Process Engine finds a running business process that uses the document's conversation ID, it passes the document to the running process. If Process Engine does not find such a running business process, it searches for a process model whose first step waits for the document and starts a new instance of the process model. The Process Engine logs the status of business processes so you can view and monitor their progress.

Error Logging

Document Recognition Errors

- To determine the document type to use for an XML document, Trading Networks looks at all enabled XML document types. To determine the document type to use for a flat file document, Trading Networks looks at all enabled flat file document types. Each document that passes through your system should match exactly one document type. If a document does not match any document type, or matches more than one document type, Trading Networks considers it an *unknown document type*. When Trading Networks encounters an unknown document type, it:
 - Cannot extract any attribute information from the document, since document types identify the attribute information to extract.
 - Tries to process the document using a processing rule that acts on unknown document types, if you define one.
 - For documents that match more than one document type, Trading Networks logs a message that identifies those document types to the Trading Networks *activity log*.
- Trading Networks flags the following during document recognition:
 - A sender or receiver error if the external ID of the sender or receiver in the document does not match the value of the external ID in any partner profile.
 - An attribute error if Trading Networks cannot extract an attribute that is marked as required in the document type.
 - For an XML document, a sender or receiver error if the XQL query in the document type cannot find the SenderID or ReceiverID system attribute in the document.
 - For an XML document, a general error if an XQL query in the document type fails. For example, a general error occurs when the document is identified as an XML document but contains invalid XML or is actually another format, such as plain text or EDI.

Document Processing Errors

- When a partner tries to send documents to a partner whose profile is disabled, Trading Networks logs an error message to the activity log, sets the processing status of the document to ABORTED, and sends an email message to the Trading

Networksadministrator. If the document was sent using HTTP, Trading Networks also responds to the sender with an HTTP 403 message.

- If Trading Networks encounters an error while performing a pre-processing action, it records the information to its activity log and continues to perform the other pre-processing actions. The same is true for processing actions.
- If you used the Check for Duplicate Document pre-processing action, Trading Networks might not be able to determine whether a document is unique if:
 - It could not match the document to a document type.
 - It previously received the same document but did not save it to the database.
 - It previously received the same document and saved it to the database, but:
 - The document type did not indicate to extract the DocumentID, SenderID, and ReceiverID system attributes.
 - For a flat file, the document gateway service did not place the DocumentID, SenderID, and ReceiverID system attributes in the pipeline.
- If you used the Deliver Document By processing action, Trading Networks might require additional information from the receiving partner's profile before it can deliver the document. For example, for immediately delivery methods, Trading Networks must look in the profile for the host name and port for the receiving partner's system. If Trading Networks cannot determine the required delivery information (for example, because the receiving partner's profile is disabled), it logs the error to the activity and queues the document for polling.
- If a document was sent by one of your partners but Trading Networks cannot determine the sender, Trading Networks considers the sender an *unknown partner*. If you have a processing rule that processes documents that are sent by the partner, Trading Networks will not select the rule. Trading Networks cannot determine the sender in these cases:
 - The document type is unknown, so Trading Networks cannot extract any document attributes, including the sender or receiver.
 - The document type does not indicate to extract the SenderID and ReceiverID system attributes, so Trading Networks cannot determine the document sender or receiver.
 - The value of the SenderID and ReceiverID system attributes do not match the value of the external ID type of any partner defined in a Trading Networks profile.
 - In an XML document type, the XQL query cannot find the SenderID and ReceiverID system attributes in the document.

Monitoring Transactions

An end-to-end business-to-business document exchange flow corresponds to a transaction in Trading Networks. Irrespective of the type of documents, delivery methods, trading partner agreements used, a transaction is a neutral runtime aspect of a business-to-business document exchange flow in Trading Networks. You can search, view, and monitor transactions using the simple and advanced search capabilities.

To monitor the intermediate states within a transaction at a granular level, you can define a set of user statuses and update these either using a service during the course of the transaction or add or delete user statuses through **Applications > Administration > Integration > B2B > B2B Settings > Transaction Settings > User Status Configuration**.

Trading Networks gives you visibility into your network to track run-time information about the documents that your Trading Networks system has sent or received, delivery and service execution tasks that have been run or started, and activity log entries relating to the server. You can use My webMethods to query run-time information that Trading Networks has saved to its database.

The following table lists the run-time information you can view, along with the My webMethods pages to use to view the information:

My webMethods Page	To View
Monitoring > Integration > B2B > Transactions	<p>Information about the documents that Trading Networks has saved:</p> <ul style="list-style-type: none"> ■ Attributes that have been extracted from documents ■ Content of documents that have passed through your system ■ Status of documents that Trading Networks is in the process of delivering ■ For each transaction you can view the following information: <ul style="list-style-type: none"> ■ Date Received ■ Document Type ■ Sender ■ Receiver ■ Processing Status ■ User Status ■ Related Documents

My webMethods Page	To View
	<ul style="list-style-type: none"> ■ Details ■ Action ■ Any task associated with the document exchange ■ Any comments that you have recorded for a transaction <p>In addition to viewing this information, Trading Networks also provides features you can use to resubmit and reprocess documents.</p>
Monitoring > Integration > B2B > Tasks	<p>The progress and status of each delivery task and service execution task. In addition to viewing this information, Trading Networks also provides features you can use stop, restart, and delete tasks.</p>
Monitoring > Integration > B2B > Activity Log	<p>Audit information of the activity that has occurred in your Trading Networks system.</p>

If you use the same query multiple times, you can save the query settings. When you want to use the same query again, select the saved query and run it. Queries you save in one of the user interfaces cannot be used in the other user interface on Trading Networks

Security

Managing, Authenticating, and Authorizing Trading Networks Users

Trading Networks manages Trading Networks users through webMethods central user management, as follows:

- When a My webMethods user who has Trading Networks administrator authority wants to perform an action that requires execution of a service on Integration Server. The user's My webMethods credentials are used to authenticate the user and authorize the request.
- When a My webMethods user wants to view Trading Networks data or perform Trading Networks actions, and Trading Networks services are invoked on Integration Server. The user's My webMethods credentials are used to authenticate the user and authorize the request.

- When a Trading Networks partner sends a document to Trading Networks, and a Trading Networks service is invoked. The partner can invoke the service using the credentials of a My webMethods user account.

When Integration Server receives a user name and password to authenticate, it first tries to authenticate the user using its own user account definitions. If the user is not defined in Integration Server user accounts, Integration Server determines whether the user account is defined in My webMethods central user management. If so, Integration Server checks whether the user supplied valid My webMethods credentials. When you use Trading Networks through My webMethods and other authentication methods such as client-side certificate authentication and third-party tools, My webMethods Server passes an authentication token to Integration Server.

To authorize a request, the Integration Server determines whether the user can access the requested Integration Server service. Access to Integration Server services is protected by Access Control Lists (ACLs). When using central user management, you can add My webMethods groups and roles that relate to Trading Networks to the Allowed list of ACLs.

Protecting Access to User Interfaces

To prevent unauthorized access to My webMethods, a user is authenticated with a user name and password.

Trading Networks actions that you access from My webMethods are protected by *role-based access*. That is, to view Trading Networks data (for example, information about a delivery task) and to perform actions against that data (for example, restart a task), you must be a member of a role to which that permission has been granted. There are two aspects to role-based access:

- *Data permissions*, which identify a data set of partner profiles, document types, transactions, tasks, and activity log entries along with the actions a role can perform against that Trading Networks data set. Using data-level security, you can grant roles the permissions to perform such actions as viewing, editing, or creating profiles.
- *General functional permissions*, which grant a role the permission to perform Trading Networks actions against other Trading Networks data. Using general functional permissions, you can grant roles the permissions to perform such tasks as managing public queues or partner groups, or deleting partner profiles.

Protecting Partner Profile Passwords

All passwords contained in partner profiles are securely managed by the Integration Server's Password Manager. When you create a partner profile, Trading Networks creates a handle for the password and passes both the password and its handle to the Password Manager. The Password Manager encrypts the password and stores the password and its handle in the IS repository. The Trading Networks database stores only the handle.

When you need to display or update an existing profile, Trading Networks reads the appropriate handle in its database and asks the Password Manager to return the password. The Password Manager obtains the password from the Integration Server repository, decrypts it, and returns it to Trading Networks. If the password is already cached in the Trading Networks database, this process is not necessary.

Note: Passwords used in scheduled delivery queues (public and private) are stored in the Trading Networks database in binary-encoded form (*not* in plain text). Trading Networks cannot encrypt passwords used in scheduled delivery queues; because trading partners are allowed to create custom scheduled delivery services, Trading Networks cannot anticipate which user-defined input variable might be a password.

Protecting Access to Trading Networks Processing

When trading partners want to connect to your Trading Networks system (for example, to send a document for processing), access can be protected through a user account (user name/password) or x.509v3 client certificates. A partner must have partner authority to access your Trading Networks system to exchange documents. When you define a profile for a partner, you can associate one or more My webMethods or Integration Server user accounts with a profile. Your partner can use the user accounts to access your system.

When your Trading Networks system needs to connect to a partner's system (for example, to deliver a document), it can use a user account (user name/password) or x.509v3 client certificates as credentials that the partner's system uses for authentication. If your partner requires authentication using user name/password, your Trading Networks system maintains the user name and password it needs to supply when connecting to that partner in the partner's profile on your system. If a partner requires authentication using client certificates, your Integration Server system maintains the client certificate it needs to supply when connecting with that partner.

Trading Networks protects access to the `wm.tn:receive` service using an Integration Server Access Control List (ACL). The protection ensures that only partners with Trading Networks administrative authority or partner authority can invoke this service. To invoke the `wm.tn:receive` service, the client must supply the user name and password of a valid My webMethods or Integration Server user account. When using a user account with Trading Networks administrative authority, Trading Networks always accepts and processes the document. However, you will typically not grant your partners administrative authority. Instead, they have user accounts that have Trading Networks partner authority.

When you create a profile for a partner, you can associate a user account with the profile, and therefore the partner. You can associate one or more My webMethods users with the profile; Trading Networks automatically gives partner authority to the My webMethods user. When using a user account with partner authority, Trading Networks makes sure that the user that invokes the `wm.tn:receive` service matches the sender specified within the document being sent. Trading Networks uses the sender identified within the document

to look up the sender's profile and makes sure that the profile is associated with the My webMethods or Integration Server user account that was used to send the document. If the user account is not associated with the sender's profile, Trading Networks does not process the document.

Certificates for Verifying, Signing, Encrypting, and Decrypting Documents and Authenticating Connections

Trading Networks certificate sets consist of sign/verify, encrypt/decrypt, and Secure Sockets Layer (SSL) authentication certificates. You can use a single set of certificates for all partners, or you can use a unique set of certificates for each sender/receiver pair (or selected pairs). For example, you can use one set of certificates for sending documents from A to B, and a different set of certificates for sending documents from C to A.

When you define your profile and the profiles of your trading partners, you specify the following kinds of certificates in the sender or receiver profiles:

Specify this certificate...	In this profile...	For this purpose...
Sign	Sender's profile	When you sign a document to send to a partner, Trading Networks looks at your profile to see if it contains the specific private key to use to sign the document.
Verify	Sender's profile	When a partner sends a document to you, Trading Networks looks at the sender's profile to see if it contains the specific public certificate to use to verify the document.
Encrypt	Receiver's profile	When you encrypt a document to send to a partner, Trading Networks looks at the receiver's profile to see if it contains the specific public certificate to use to encrypt the document.
Decrypt	Receiver's profile	When a partner sends an encrypted document to you, Trading Networks looks at your profile to see if it contains the specific private key to use to decrypt the document.
SSL	Sender's profile	This certificate represents the partner's authentication credentials when making an SSL connection with Integration Server.

Certificates associated with *partner profiles* are stored in separate files in the Trading Networks database. Certificates associated with *Enterprise profiles* are stored in keystore files on Integration Server.

Keystores consist of one or more pairs of private keys and signed certificates for their corresponding public keys. Each key pair is identified by a unique *key alias*. Keystores are identified by a unique *keystore alias*.

You create and edit keystore aliases for certificates associated with Enterprise profiles from the **Security > Keystore** panel in Integration Server Administrator. You create key aliases to identify specific keys within a keystore using a third-party certificate management tool.

For more information about keystores and certificates, see *webMethods Integration Server Administrator's Guide*.

Overlapping of Certificates

You can upload up to two active certificate sets each (referred to as the primary and secondary certificate sets) for sign/verify, encrypt/decrypt, and SSL certificate types, so that when one certificate set expires, Trading Networks can switch to the other one without interrupting the processing of documents. The certificate set that you add first is the *primary* certificate set. Trading Networks automatically switches from the primary certificate set to the secondary one when any of the following occurs:

- The primary certificate has expired and the secondary certificate has not expired.
- The receiver's sign/verify or SSL primary certificate set does not match the sender's sign/verify or SSL certificate set.

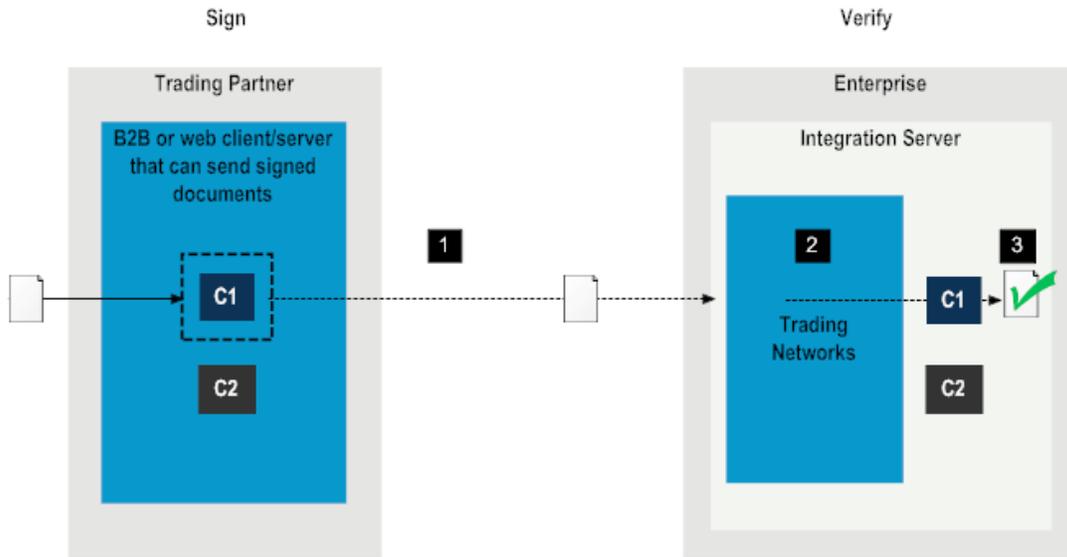
Trading Networks does not switch encryption/decryption certificates at the receiver's end. The receiver of the document must write a flow service that first obtains the certificate ID of the appropriate decrypt certificate, using the `wm.tn.security.getAllCertificateData` service, and then set that certificate as the primary one for that partner, using the `wm.tn.security.setPrimaryCertificate` service. Doing so ensures that the correct decryption certificate is retrieved for future transactions with that partner.

The following diagrams illustrate the sign/verify, encryption/decryption, and SSL scenarios:

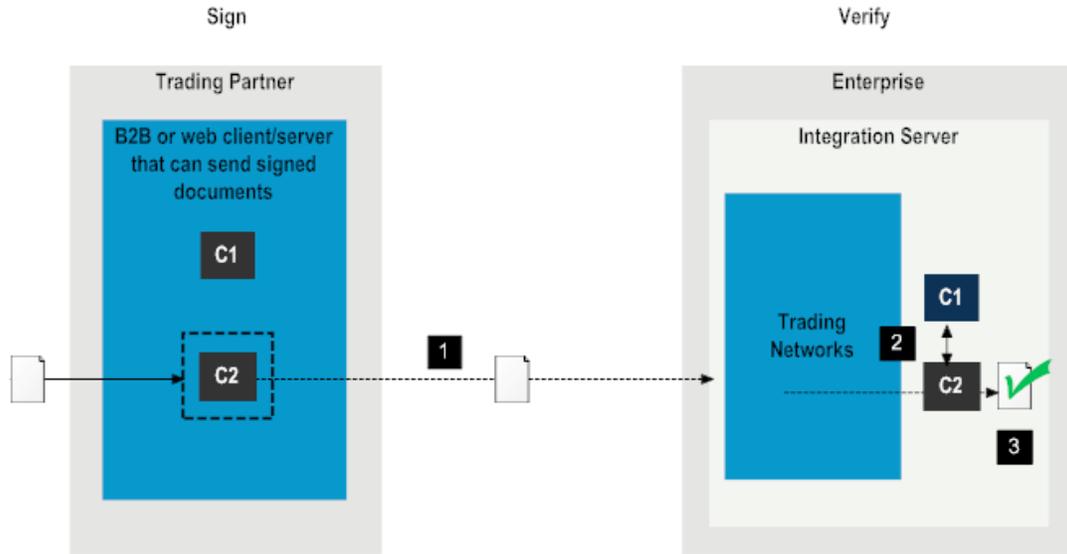
Sign-Verify Scenario 1

Enterprise: Certificate C1 and certificate C2 are valid; C1 is the primary certificate.

Trading Partner: Sends a document signed with certificate C1.



Step	Description
1	The trading partner sends a document signed with certificate C1 to the enterprise.
2	Trading Networks on the enterprise side retrieves the certificate C1 for the trading partner from Trading Networks.
3	Trading Networks on the enterprise side verifies the document with C1. Verification is successful.

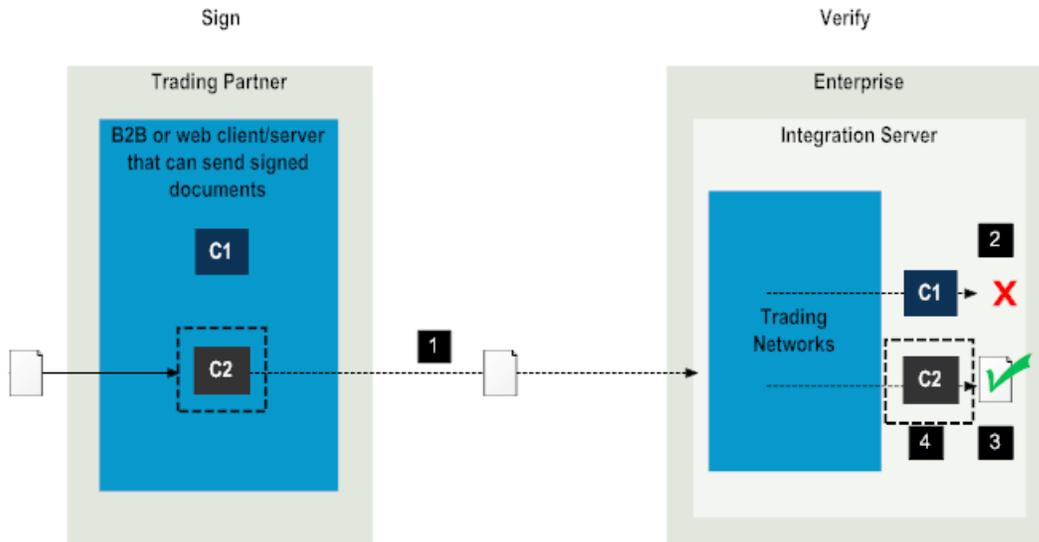


Step	Description
1	The trading partner sends a document signed with C2 certificate to the enterprise.
2	Trading Networks on the enterprise side switches the primary certificate to C2 and retrieves the certificate C2 because certificate C1 has expired.
3	Trading Networks on the enterprise side verifies the document with C2. Verification is successful.

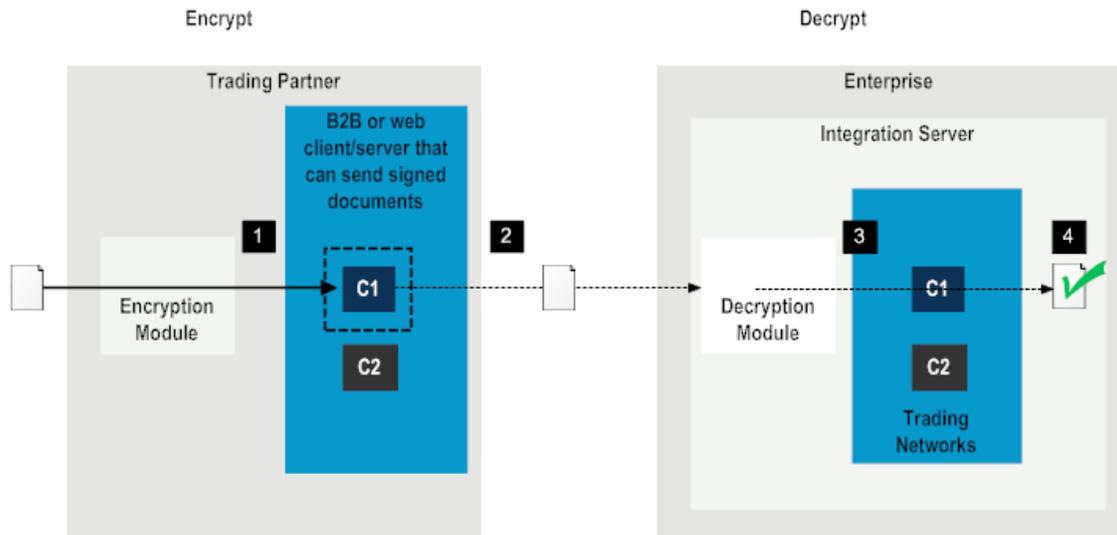
Sign-Verify Scenario 3

Enterprise: Certificate C1 and certificate C2 are valid; C1 is the primary certificate.

Trading Partner: Sends a document signed with certificate C2.



Step	Description
1	The trading partner sends a document signed with certificate C2 to the enterprise.
2	Trading Networks on the enterprise side retrieves the certificate C1 for the trading partner from Trading Networks and verifies the document with certificate C1. Verification fails as the document is signed with the certificate C2.
3	Trading Networks on the enterprise side retrieves the certificate C2 and verifies the document with C2. Verification is successful.
4	Trading Networks on the enterprise side sets the certificate C2 as the primary certificate for the trading partner.

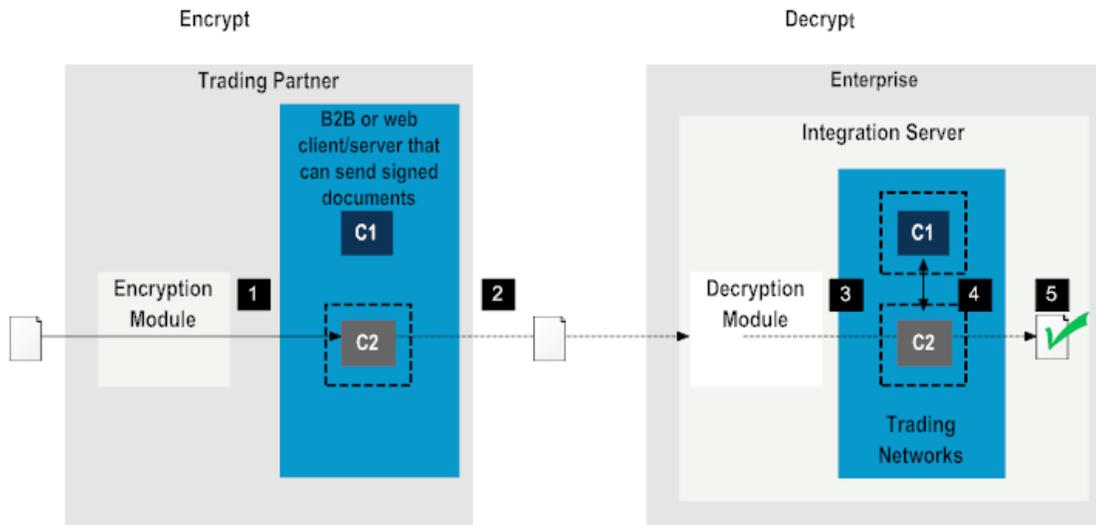


Step	Description
1	The encryption module on the trading partner requests the encryption certificate and gets the primary certificate C1.
2	The encryption module on the trading partner encrypts the document with the certificate C1 and sends the document to the enterprise.
3	The decryption module on the enterprise side requests Trading Networks for the decryption certificate and gets the certificate C1.
4	The decryption module on the enterprise side decrypts the document using certificate C1. Decryption is successful.

Encrypt-Decrypt Scenario 2

Enterprise: Certificate **C1** has expired and certificate **C2** is valid; **C1** is the **primary** certificate.

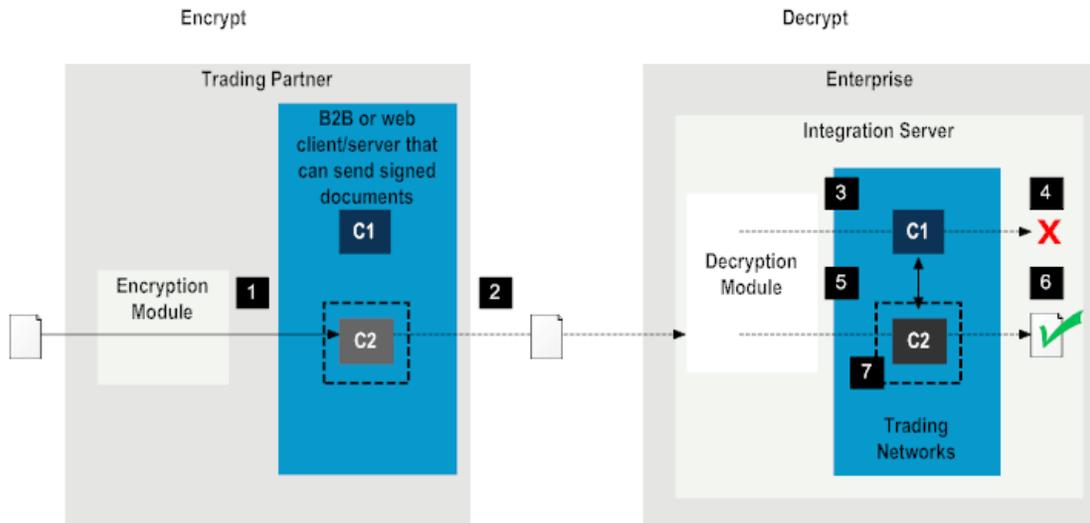
Trading Partner: Sends a document encrypted with certificate **C2**.



Step	Description
1	The encryption module on the trading partner requests the encryption certificate and gets the certificate C2 because C1 has expired.
2	The encryption module on the trading partner encrypts the document with the certificate C2 and sends the document to the enterprise.
3	The decryption module on the enterprise side requests Trading Networks for the decryption certificate.
4	Trading Networks on the enterprise side switches the primary certificate to C2 because certificate C1 has expired, and returns C2.
5	The decryption module on the enterprise decrypts the document using certificate C2. Decryption is successful.

Encrypt-Decrypt Scenario 3

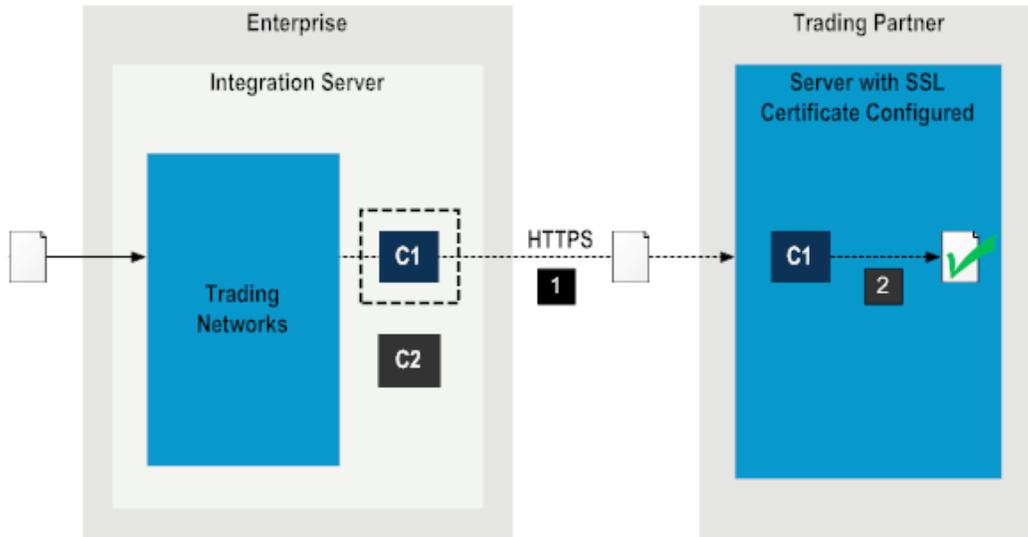
Enterprise: Certificate C1 and certificate C2 are valid; C1 is the primary certificate.
 Trading Partner: Sends a document encrypted with certificate C2.



Step	Description
1	The encryption module on the trading partner requests the encryption certificate and gets the certificate C2.
2	The encryption module on the trading partner encrypts the document with the certificate C2 and sends the document to the enterprise.
3	The decryption module on the enterprise side requests Trading Networks for the decryption certificate and gets the certificate C1.
4	The decryption module on the enterprise side decrypts the document using certificate C1. Decryption fails.
5	The decryption module requests for the secondary certificate C2 using the <code>wm.tn.security.getAllCertificateData</code> service.
6	The decryption module decrypts the document with certificate C2. Decryption is successful.
7	The decryption module calls the Trading Networks <code>setPrimaryCertificate</code> service to switch the primary certificate to C2.

SSL Scenario 1

Enterprise: Certificate **C1** and certificate **C2** are valid; **C1** is the **primary** certificate.
Trading Partner: Certificate **C1** is configured as the SSL certificate.

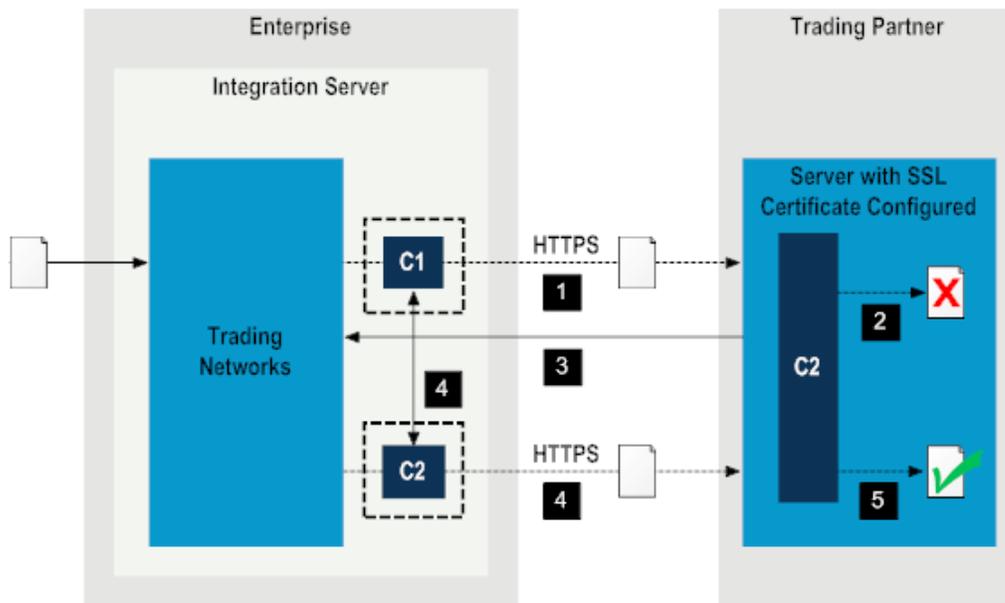


Step	Description
1	The enterprise sends a document to the trading partner over HTTPS using the private key from certificate C1.
2	The trading partner's server authenticates the document using the SSL certificate C1 configured on the server. Authentication is successful and the transaction is complete.

SSL Scenario 2

Enterprise: Certificate C1 and certificate C2 are valid; C1 is the primary certificate.

Trading Partner: Certificate C2 is configured as the SSL certificate.



Step	Description
1	The enterprise sends a secure document to the trading partner over HTTPS using the private key from certificate C1.
2	The trading partner's server authenticates the document using the SSL certificate C2 configured on the server. Authentication fails.
3	The trading partner's server sends an error message to the enterprise.
4	The enterprise switches the SSL certificate to C2 and resends the document to the trading partner over HTTPS.
5	The server on the trading partner authenticates the document using the SSL certificate C2 configured on the server. Authentication is successful. The transaction is complete.

Verifying Digital Signatures

Trading Networks supports x.509v3 certificates for verifying the digital signature of documents sent by a partner. Trading Networks verifies the digital signature to make sure the documents have arrived unchanged and the sender is who it claims to be. To verify the digital signature, you do the following:

- Save the partner's Verify certificate in the partner's profile. Trading Networks must have access to the partner's certificates. When you add a Verify certificate, Trading Networks stores the certificate in its database.

Note: If you include the private key in the certificate information, Trading Networks can also use this information to digitally sign documents on behalf of the partner. You might have the private key if the profile describes an internal group (for example, a department within your corporation).

- For XML documents, set up the document type to extract the SignedBody and Signature system attributes. The SignedBody attribute identifies the portion of the document that was digitally signed. The Signature attribute identifies the portion of the document that contains the digital signature. The signature must be a base64 encoded PKCS#7 detached digital signature and can contain information for one or more signers.
- Specify the Verify Digital Signature pre-processing action in the document type or processing rule.

When a partner sends a document to you, Trading Networks looks at the partner's profile to see if it contains the specific public certificate to use to verify the document. If Trading Networks finds a set of certificates to use for that specific receiver, it uses the appropriate certificate in that set. If Trading Networks does not find a set of certificates to use for that specific receiver, it uses the default set of certificates specified in the partner's profile.

To verify that the document arrived unchanged from the partner to you, Trading Networks invokes the Integration Server `pub.security.pkcs7:verify` built-in service. Trading Networks passes this service the value of the SignedBody and Signature system attributes that it extracted from the document.

Trading Networks can only verify information on itself because it does not have the certification/verification for the partner. Trading Networks makes sure that the CA that signed the certificate is included in the list of trusted CA certificates that the Integration Server maintains.

To ensure that the signed body has not changed, Trading Networks verifies the digital signature, which is the value of the Signature system attribute. To verify that the sender is who it claims to be, Trading Networks matches the certificate from the digital signature to the Verify certificate that Trading Networks has on file for the partner.

Digitally Signing Documents

Trading Networks supports x.509v3 certificates for digitally signing documents that you, the owner of the certificates, want to send to trading partners. To digitally sign a document, you invoke the `wm.tn.doc:sign` built-in service.

When you invoke this service, Trading Networks locates the sender and receiver to retrieve the correct signed certificate from the Trading Networks database. The owner of the certificate is the sender, and the receiver is the trading partner. You can set up Trading Networks to use different certificates for different partners.

You can also specify a default Sign certificate by providing the certificate information in the owner's profile. If a default Sign certificate is defined, then Trading Networks uses this default Sign certificate when a partner-specific Sign certificate is not available.

When you sign a document to send to a partner, Trading Networks looks at your profile to see if it contains the specific private key to use to sign the document. If Trading Networks finds a set of certificates to use for that specific receiver, it uses the appropriate certificate in that set. If Trading Networks does not find a set of certificates to use for that specific receiver, it uses the default set of certificates specified in your profile.

Encrypting and Decrypting Data

Trading Networks does not encrypt or decrypt documents. However, Trading Networks maintains x.509v3 certificates for other webMethods components, such as webMethods RosettaNet Module. These certificates are used for encrypting documents that are sent to partners and decrypting the encrypted documents that are received from partners.

If the other webMethods components, such as webMethods RosettaNet Module requires Encrypt certificates, save a partner's Encrypt certificate in the partner's profile. You can also add your own functionality that takes advantage of this certificate information. You can obtain the certification information by using built-in services.

Note: Trading Networks does not check to see if the CA that signed the Encrypt certificate is in the list of trusted CAs that the Integration Server maintains. If you include the private key in this certificate information, this certificate information can also be used to decrypt documents that were encrypted with the partner's public key. You might have the private key if the profile describes an internal group, for example a department within your corporation.

When you encrypt a document to send to a partner, Trading Networks looks at the partner's profile to see if it contains the specific public certificate to use to encrypt the document. If Trading Networks finds a set of certificates to use for that specific receiver, it uses the appropriate certificate in that set. If Trading Networks does not find a set of certificates to use for that specific receiver, it uses the default set of certificates specified in the partner's profile.

If the webMethods components require Decrypt certificates, save your Decrypt certificate in the owner's profile. Because you can store Decrypt certificates in the

owner's profile, you can set up alternate Decrypt certificates for different partners. You can also specify a default Decrypt certificate by providing the certificate information in the owner's profile. If a default Decrypt certificate is defined, then Trading Networks will use this default Decrypt certificate when a partner-specific Decrypt certificate is not available.

Note: Trading Networks does not check to see if the CA that signed the Decrypt certificate is in the list of trusted CAs that the Integration Server maintains.

When a partner sends an encrypted document to you, Trading Networks looks at your profile to see if it contains the specific private key to use to decrypt the document. If Trading Networks finds a set of certificates to use for that specific receiver, it uses the appropriate private key in that set. If Trading Networks does not find a set of certificates to use for that specific receiver, it uses the default set of private keys defined in the Default profile for partners.

Communicating Securely Using SSL

Because Trading Networks runs in Integration Server, it takes advantage of Integration Server features, such as its support of SSL for secure communications. To enable Trading Networks to act as an SSL client connecting to a remote secure server, you specify an SSL client certificate in your Enterprise profile or your partner's profile.

When using SSL connections that require client-side authentication, Trading Networks looks at the sender's profile to see whether it contains the specific private key to use to connect to the receiver (the remote secure server).

If Trading Networks...	Trading Networks uses...
Finds a set of certificates to use for that specific receiver	The private key from that certificate set.
Does not find a set of certificates to use for that specific receiver	The default private key specified in the sender's profile.
Does not find a default private key specified in the sender's profile	The default certificates specified in the Integration Server.

Run-Time Event Notifications

Run-time events occur on a continuous basis as Trading Networks processes a document or performs a post-processing task. Trading Networks supports run-time event notification using webMethods Event Routing. Event Routing is a framework that Software AG provides for applications to communicate using events. For more

information about using Event Routing, see *Communicating Between Software AG Products Using Event Routing*.

Trading Networks publishes two types of run-time events:

- **Transaction events.** Occur when Trading Networks processes a document.
- **Task events.** Occur when Trading Networks performs a post-processing task on a document (for example, send document).

To enable Trading Networks to publish run-time events, be sure to enable the event groups by setting the event properties for the corresponding event groups to `true`. For more information on event groups and event properties, see ["Event Groups" on page 286](#).

Software AG Intelligent Business Operations (IBO) platform-based applications can use these run-time events to generate static or real-time reports.

Caching

A cache is a small fast memory holding recently accessed data, designed to speed up subsequent access to the same data.

Ehcache is a standards-based caching API that enables applications to fetch frequently used data from an in-memory cache rather than having to retrieve it from a database or other back-end system.

Trading Networks uses the caching capabilities of Ehcache to provide:

- **Asset caching.** Trading Networks uses Ehcache to cache assets. While the assets are in cache, Trading Networks can quickly retrieve them for similar requests, rather than getting the details from Trading Networks database. For more information about asset caching, see ["Asset Caching" on page 264](#).
- **Query results caching.** Trading Networks uses Ehcache to store the results of a database query. It stores the number of rows you specify in a session object on the host Integration Server and stores the remaining rows in Ehcache. For more information about query results caching, see ["Query Results Caching" on page 264](#).

For an introduction to Ehcache and information for sizing Trading Networks system caches, see *Using Terracotta with webMethods Products*.

Dashboards and Charts

Trading Networks provides a set of dashboards and charts that give you a graphical view of partners and their transactions on a daily, weekly, and monthly basis. You can view the transaction volume summary, the transaction volume and total value trend for an attribute, successful and failed transactions, late FA count by partners and so on.

For more information on dashboards and charts, see ["Working with Dashboards and Charts" on page 289](#).

Database Partitioning

Database partitioning logically separates one set of data from the other, thereby enhancing performance especially while archiving data from Trading Networks production tables to archive tables. For more information on database partitioning, see ["Database Partitioning" on page 259](#).

2 Getting Started

■ Before Configuring Trading Networks	66
■ Starting the Host Integration Server	66
■ Shutting Down Integration Server	67
■ Planning Your Trading Network	67
■ Creating Services and Viewing Specifications	68
■ Additional Documentation	68
■ Sending Documents to Trading Networks	69
■ Summary of the Configuration Steps	71

Before Configuring Trading Networks

Before you start performing the configuration tasks described in this guide, make sure the following tasks are completed:

- Ensure that Trading Networks is installed. For details, see *Using the Software AG Installer*.
- When in a cluster setup, ensure that all Trading Networks nodes in the Trading Networks installation run the same webMethods Trading Networks version, with the same fixes applied.
- Ensure that you create a database component for Trading Networks and create a JDBC pool and point to the Trading Networks database.
- Ensure that central user management is configured in Integration Server to provide My webMethods users with access to Trading Networks.
- Ensure that the correct SAML resolver location is specified for every Integration Server instance that communicates with My webMethods Server through the WmTN package. The default SAML resolver URL is `https://localhost:8585/services/SAML`. Integration Server stores this URL in the `watt.server.auth.samlResolver` server configuration parameter. For information about changing the SAML resolver URL, see *webMethods Integration Server Administrator's Guide*.
- Ensure that My webMethods Server and Integration Server are started, in that order.

Starting the Host Integration Server

Starting Integration Server on Microsoft Windows

To start Integration Server on Microsoft Windows

1. On the Windows **Start** menu, select **Programs > webMethods > Servers > webMethods Integration Server**.
2. Make sure the Integration Server package WmPublic is enabled. For instructions, see *webMethods Integration Server Administrator's Guide*.

Starting Integration Server on UNIX

To start Integration Server on UNIX

1. Log in as a non-root user.

Important: Running the startup script as root might reduce the security of your system.

2. Execute the `server.sh` script. The script is located in the *Integration Server_directory / instances/instance_name /bin* directory.
3. Make sure the Integration Server package WmPublic is enabled. For instructions, see *webMethods Integration Server Administrator's Guide*.

Starting Integration Server from the Command Line

If you start Integration Server from the command line, you can override certain configuration settings, such as where to write the log file for the session. It also lets you start Integration Server in “debug” mode, so you can record or display Integration Server activity. For complete information, see *webMethods Integration Server Administrator's Guide*.

After you start the Integration Server, make sure the Integration Server package WmPublic is enabled. For instructions, see *webMethods Integration Server Administrator's Guide*.

Shutting Down Integration Server

When you shut down the host Integration Server, Trading Networks automatically closes its session on Integration Server.

To shut down Integration Server

1. Open Integration Server Administrator by opening an internet browser and pointing it to the Integration Server host and port.

For example, if Integration Server were running on port 4040 on a machine called QUICKSILVER, you would type `http://QUICKSILVER:4040`.
2. In the upper right corner of any Integration Server Administrator screen, click **Shutdown and Restart**.
3. Choose whether to shut down immediately (thus ending all active sessions immediately), to wait *number* minutes, or to wait until all active client sessions are complete. For instructions on how to view active sessions, see *webMethods Integration Server Administrator's Guide*.
4. Click **Shutdown**.

Planning Your Trading Network

This chapter describes the tasks to plan for and to set up your trading network. The general order of tasks is:

1. Identify the types of documents Trading Networks handles.
2. Define any custom document attributes you require.

3. Define the document types for the documents Trading Networks handles.
4. If desired, define processing rules to process the documents.
5. If desired, define business processes to process the documents.
6. Define how to deliver documents to partners.
7. Create the Enterprise profile for your corporation and a partner profile for each of your trading partners. Review the standard profile fields and define any extended profile fields that you require.

You might want to perform tasks 1 through 6 in a development environment. You can use the development environment to test whether documents are processing as you intended. After you ensure that your trading network is defined properly, you can export your Trading Networks assets to a production environment.

You can then create your profile and the partner profiles in your production environment.

You can BAM-enable Trading Networks using webMethods Optimize for B2B; that is, you can analyze and monitor Trading Networks transactions for business purposes. For instructions, see *webMethods Optimize for B2B User's Guide*.

Creating Services and Viewing Specifications

You can use Software AG Designer or your own Java development environment to create services. If you use your own Java development environment, you must use a jcode utility to put the Java code into the Integration Server namespace; that is, to place it in a package and folder. The jcode utility is provided. For instructions on creating services and using the jcode utility, see *webMethods Service Development Help*.

Do not save your services in the WmTN package or in any other webMethods package that contains webMethods software. If you do, when you upgrade the webMethods software, your service will be lost.

To view specifications for Trading Networks built-in services, open Software AG Designer and look in the WmTN package. For complete information about the specifications, see *webMethods Trading Networks Built-In Services Reference*.

Additional Documentation

For information about...	See...
Trading Networks built-in services and service specifications	<i>webMethods Trading Networks Built-In Services Reference</i>

For information about...	See...
IS document types that relate to business documents, document types, and document attributes, including the BizDocEnvelope	<i>webMethods Trading Networks Built-In Services Reference</i>
Creating services and IS document types	<i>webMethods Service Development Help</i>
Integration Server built-in services	<i>webMethods Integration Server Built-In Services Reference</i>

Sending Documents to Trading Networks

Documents can be sent to Trading Networks as follows:

- A trading partner that does not have its own Trading Networks can send a document to your Trading Networks.
- A trading partner's Trading Networks can deliver a document to your Trading Networks, or vice versa.
- Your Trading Networks can process a document and then return the document to your system for additional processing.

Sending Documents to Your Trading Networks from Another System

If a trading partner does not have its own Trading Networks, the partner must use an application that communicates with the host Integration Server to send documents to Trading Networks. This application is called an *IS client*. Clients can be Java, C/C++, or browser based.

If sending an XML document, the client must invoke the `wm.tn:receive` service. If sending a flat file document, the client must invoke a document gateway service, and the gateway service must invoke the `wm.tn:receive` service.

Clients can send documents to Trading Networks service using the methods below.

Method	Client
Post an XML or flat file document to the Integration Server through HTTP, HTTPS, FTP, SFTP, FTPS, or SMTP.	Any type except browser based

Method	Client
Submit an XML document in the <i>\$xmldata</i> variable.	Any type
Send an XML or flat file document by invoking a Web service you publish.	Any type

You can create a Web service using any Integration Server service. The service you use to create the Web service must internally invoke the `wm.tn:receive` service.

For instructions on...	See...
Creating and publishing Web services	<i>Web Services Developer's Guide</i>
Creating clients	<i>webMethods Service Development Help</i>
Passing XML data to services	<i>webMethods Service Development Help</i>
Security for this type of document exchange	"Protecting Access to Trading Networks Processing" on page 48

Sending Documents from One Trading Networks to Another

A trading partner can have its Trading Networks process a document and then use a processing rule to send the document to your Trading Networks. In this situation, the partner's Trading Networks acts as a client to your Trading Networks. Conversely, you can have your Trading Networks process a document and then use a processing rule to send the document to a partner's system. In this situation, your Trading Networks becomes the client to your partner's Trading Networks.

To send a document from one Trading Networks to another, you use a processing rule action, as follows:

Processing Action	Description
Execute a Service	Invokes a service that forwards the document to the target Trading Networks. For an XML document, the service must invoke the <code>wm.tn:receive</code> service. For a flat file document, the service must invoke a document gateway service, and the gateway service must invoke the <code>wm.tn:receive</code> service.

Processing Action	Description
Deliver Document By	Delivers the document to the partner identified as the receiver in the document using the specified delivery method.

Sending a Document Back to Your Trading Networks

You can have your Trading Networks process a document and then send the document back to itself for additional processing. For example, a partner sends a document to your Trading Networks in cXML format. The receiving partner needs the document in CBL format. Your Trading Networks performs processing that converts the cXML document to CBL format. Your Trading Networks then sends the CBL document back to itself so the CBL document's processing rule can deliver the document to the receiving partner.

To return an XML document to the same Trading Networks, use the `wm.tn.doc.xml:routeXml` service rather than the `wm.tn:receive` service. For a flat file document, use the `wm.tn.doc.ff:routeFlatFile` service rather than the document gateway service.

Note: Trading Networks does not check the identity of the sender in this case against the IS user account that invokes the `wm.tn.doc.xml:routeXml` or `wm.tn.doc.ff:routeFlatFile` service. That is, Trading Networks does not verify that the user invoking these services has Trading Networks partner authority or that the sender identified in the document is associated with the partner that sent the document.

Summary of the Configuration Steps

After you install Trading Networks, you must perform the following high-level steps to configure Trading Networks.

Note: Command Central allows you to configure Trading Networks server properties. See ["Using Command Central to Manage Trading Networks" on page 75](#)

1. Add an Trading Networks instance to My webMethods Server if the Integration Server that hosts Trading Networks is running on a machine other than localhost on port 5555.
2. Replace the default SSL certificate for Trading Networks instances used in production environments if you are using ActiveTransfer Server along with Trading Networks.
3. Set up the central user management in My webMethods Server.
4. Create server ports and configure settings for specific Trading Networks instances.

5. Create an enterprise profile.
6. Create partner profiles for partners in Trading Networks to exchange business documents. Import or onboard partners as necessary. See ["Creating Profiles" on page 185](#)
7. Grant access permissions to partners, and other non-administrator users access to specific assets in My webMethods Server user interface. See ["Granting Access to Trading Networks " on page 91](#)
8. Define custom attributes to recognize documents and set the preferred delivery method for each of the partners. See ["Create Custom Attributes" on page 104](#)
9. Install appropriate eStandards Modules to enable document exchange between partners.
10. Define document types. The two default document types available for document exchange are XML and flat file types. See ["Creating XML Document Types" on page 105](#) and ["Creating Flat File Document Types" on page 121](#).
11. If there are specific attributes which need special processing, you can enhance the processing of the document by creating a new rule or modifying it instead of using the default processing rule. See ["Creating Processing Rules" on page 157](#)
12. Define or customize Trading Partner Agreements. See ["Working with TPAs" on page 211](#)
13. Identify assets. See ["Asset Definition" on page 25](#)
14. Define jobs that, when triggered, cause Trading Networks to perform a specified set of actions. These events ensure the continuous exchange of business documents.
15. Monitor Trading Networks transactions by searching, viewing, and troubleshooting failed transactions.
16. Perform asset migration activities. See [" Migrating Assets" on page 227](#)
17. Import, export, or deploy assets across environments. See the *webMethods Deployer User's Guide* for more information.
18. Create dashboards and perform analytics. See ["Working with Dashboards and Charts" on page 289](#)
19. Perform administrator activities such as:
 - Manage caching. See ["Caching Assets and Query Results" on page 263](#)
 - Define parameters to handle large documents. See ["Large Document Handling" on page 329](#)
 - Manage security. See ["Security" on page 46](#)
 - Manage database partitioning. See ["Partitioning Trading Networks Database" on page 260](#)
 - Manage document delivery. See ["Document Delivery" on page 36](#)

- Perform document archival and deletion. See "[Archiving and Deleting Documents](#)" on page 247
- Exchange business documents by associating Trading Networks with an webMethods API Gateway instance. See "[Managing Partner Access to APIs](#)" on page 216

After completing these activities, Trading Networks configuration is complete. You can monitor transactions using the Activity Log.

3 Using Command Central to Manage Trading Networks

Command Central provides you with a centralized management capability that enables you to manage the webMethods suite of products.

You can perform the following activities for Trading Networks using Command Central:

- Manage Trading Networks servers through Command Central by navigating to **Home > Instances > MWS_instance_name > Trading Networks > Configuration** tab.
- Manage Trading Networks properties through Command Central by navigating to **Home > Instances > IS_instance_name > Trading Networks > Configuration** tab > **General Properties > Trading Networks properties**.
- Manage Trading Networks logs through the Integration Server that is hosting the Trading Networks instance.
- Manage fixes by configuring the appropriate source repository for Trading Networks.
- Perform template-based provisioning of Trading Networks instances.

For details on how to perform the Command Central operations, see *Software AG Command Central Help*.

Important: The Trading Networks Server aliases that are created with special characters and white spaces on My webMethods server do not get listed on **Home > Instance > MWS_instance_name > Trading Networks > Configuration** tab on Command Central.

4 Configuring Trading Networks

■ Overview	78
■ Configuring Database Settings	78
■ Configuring Trading Networks for MySQL Community Edition 5.7	80
■ Configuring Trading Networks for MySQL Enterprise Edition 5.7	82
■ Configuring Trading Networks for a Clustered Environment	82
■ Configuring E-mail Settings for Document Delivery	84
■ Configuring Alias for ActiveTransfer on Remote Server	84
■ Configuring Task Settings	85
■ Configuring User Status	86
■ Configuring Trading Networks to Work with My webMethods	87
■ Configuring Resubmission and Reprocessing Settings	89

Overview

To see a list of the Trading Networks configuration properties and a description of each, see ["Configuration Properties" on page 297](#).

Configuring Database Settings

When you install Trading Networks, you specify these database connection parameters:

- Type of database (for example, Oracle), and URL that specifies the location of the database.

Note: To get optimal performance from Trading Networks, you need to add the `WireProtocolMode=1` parameter to the JDBC URL when Oracle is the database. Sample URL: `jdbc:wm:oracle://localhost:1521;serviceName=orcl;WireProtocolMode=1`.

- If required, user name and password for Trading Networks to connect to the database.
- Any additional database properties for Trading Networks to supply when connecting to its database.

You can change these database connection parameters or others, such as minimum and maximum connections. For instructions, see *Installing Software AG Products*.

Configuring a High Availability Environment

To meet your high availability requirements, you may need multiple instances of Trading Networks.

Clustered Trading Networks instances must share a single Trading Networks database. When a Trading Networks database is shared:

- All configuration data and run-time data (for example, documents received from partners, or the activity log) is shared among all Trading Networks instances in the cluster.
- Changes you make to Trading Networks objects (for example, profiles, custom document attributes, document types, and processing rules), are automatically shared among all Trading Networks instances in the cluster.
- If you modify Trading Networks properties in My webMethods or by invoking the `wm.tn.admin:setProperties` service, the changes are automatically replicated to all instances of Trading Networks in the cluster. If you want to modify Trading Networks properties by directly editing the `properties.cnf` file, you must manually make the changes to the `properties.cnf` file for each Trading Networks

instance. The `properties.cnf` file is in the `Integration Server_directory\instances\instance_name\packages\WmTN\config` directory.

Multiple Trading Networks clusters can share a single Trading Networks database. Clustered Trading Networks instances can also share a single database. Non-clustered Trading Networks instances cannot share a single Trading Networks database. You need to create a Trading Networks database for each non-clustered Trading Networks instance. When a Trading Networks database is not shared:

- Changes you make to Trading Networks assets or properties are not shared among Trading Networks instances. If you want to keep the instances in sync, you must make the changes to each instance.
- Trading Networks instances do not share run-time data.
- You cannot relate documents sent to and processed by one Trading Networks instance to documents sent by and processed by other instances. For example, an Acknowledgement document that is sent to instance 1 cannot be related to the Confirmation document if the Confirmation document is sent to instance 2.

Configuring a Cluster to Share a Database

To configure a cluster to share a database

1. Set the `tn.cluster.sync.remoteAliases` property for each Trading Networks instance in the cluster.
2. Make sure the system properties in the `Integration Server_directory\instances\instance_name\config\server.cnf` file for all Integration Server instances in the cluster are identical.
3. Make sure the services you need Trading Networks to execute (for example, for delivery services or services invoked from processing rules) is available on all Integration Servers in the cluster.

Configuring Multiple Clusters Instances to Share a Database

To configure multiple clusters instances to share a database

1. Set the `tn.cluster.sync.remoteAliases` property for each Trading Networks instance that will share the database.
2. Define a remote alias for each Trading Networks instance that will share the database. For instructions, see *webMethods Integration Server Administrator's Guide*.
3. Make sure the system properties in the `Integration Server_directory\instances\instance_name\config\server.cnf` file for all Integration Server instances in the cluster are identical.
4. Make sure the services you need Trading Networks to execute (for example, delivery services or services invoked from processing rules) is available on all Integration Servers.

Note: You can use a load balancer to “cluster” Trading Networks instances if your application code is stateless and you are not using stateful processes (for example, long-running conversations).

Configure a Non-Clustered Instance to Each Have Its Own Database

To configure a non-clustered instance to have its own database

1. If you want each Trading Networks instance to have identical versions of Trading Networks objects, use the Trading Networks import/export feature to copy the data from one instance to another. For instructions, see "[Migrating Assets](#)" on page 227.
2. If you want all Trading Networks instances to have identical configurations, make sure that:
 - All Trading Networks properties in the *Integration Server_directory*\instances*instance_name*\packages\WmTN\config\properties.cnf file are identical. You can view and set these properties from the **Administration > Integration > B2B Settings > Configure Properties** page of My webMethods.
 - All system properties in the *Integration Server_directory*\instances*instance_name*\config\server.cnf files are identical.
3. If you want all Trading Networks instances to be able to handle all inbound documents, make sure the services you need Trading Networks to execute (for example, delivery services or services invoked from processing rules) is available on all Integration Servers.

Configuring Trading Networks for MySQL Community Edition 5.7

Before you configure MySQL Community Edition 5.7 with Trading Networks, the required JDBC database drivers should be configured in Integration Server and Database Component Configurator. For more information, see *webMethods Integration Server Administrator's Guide* and *Installing Software AG Products* respectively.

Note: You can configure Trading Networks for MySQL Community Edition 5.7 using Trading Networks Server 10.1 Fix 1 and above.

Important: The default value for `max_allowed_packet` is 4MB for MySQL. If this value is less than the document size, the document insertion into the database fails. To avoid this error, contact your MySQL database administrator to update the value of `max_allowed_packet` value to the largest possible document size to enable document processing and insertion into the database. The protocol limit for `max_allowed_packet` is 1GB. The value must be a multiple of 1024; non-multiples are rounded down to the nearest multiple.

To configure Trading Networks for MySQL Community Edition 5.7

1. Download the JDBC driver for MySQL Community Edition 5.7 and place it in the following directory:

Integration_Server_Installation_directory \common\lib\ext

2. To configure the MySQL Community Edition driver on Windows operating system:
 - a. Go to the *Integration_Server_Installation_directory* \packages\WmTN\bin directory, open the setcp.bat file, and add the MySQL Community Edition driver jar to the classpath, as follows:

```
if exist "%INSTALLDIR%\common\lib\ext\" (
    set CLASSPATH=%CLASSPATH%;"%INSTALLDIR%\
common\lib\ext\"
)
```

For example:

```
if exist "%INSTALLDIR%\common\lib\ext\mysql-connector-
java-5.1.41.jar" (
    set CLASSPATH=%CLASSPATH%;"%INSTALLDIR%\common\lib\ext
\mysql-connector-java-5.1.41.jar"
)
```

- b. Go to the *Integration_Server_Instance_directory* \packages\WmTN\bin directory, open the setcp.bat file, and add the MySQL Community Edition driver jar to the classpath, as follows:

```
if exist "%INSTALLDIR%\common\lib\ext\" (
    set CLASSPATH=%CLASSPATH%;"%INSTALLDIR%\
common\lib\ext\"
)
```

For example:

```
if exist "%INSTALLDIR%\common\lib\ext\mysql-connector-
java-5.1.41.jar" (
    set CLASSPATH=%CLASSPATH%;"%INSTALLDIR%\common\lib\ext
\mysql-connector-java-5.1.41.jar"
)
```

3. To configure the MySQL Community Edition driver on operating systems not based on Windows:

- a. Go to the *Integration_Server_Installation_directory* \packages\WmTN\bin directory, open the setcp.sh file, and add the MySQL Community Edition driver jar to the classpath, as follows:

```
if [ -f ${INSTALLDIR}/common/lib/ext/<MySQL Community Edition
driver name> ]; then
CLASSPATH=${CLASSPATH}:${INSTALLDIR}/common/lib/ext/<MySQL
Community Edition driver name>
fi
```

For example:

```
if [ -f ${INSTALLDIR}/common/lib/ext/mysql-connector-
java-5.1.41.jar ]; then
CLASSPATH=${CLASSPATH}:${INSTALLDIR}/common/lib/ext/mysql-
```

```
connector-java-5.1.41.jar
fi
```

- b. Go to the *Integration_Server_Instance_directory* \packages\WmTN\bin directory, open the setcp.sh file, and add the MySQL Community Edition driver jar to the classpath, as follows:

```
if [ -f ${INSTALLDIR}/common/lib/ext/<MySQL Community Edition
driver name> ]; then
CLASSPATH=${CLASSPATH}:${INSTALLDIR}/common/lib/ext/<MySQL
Community Edition driver name>
fi
```

For example:

```
if [ -f ${INSTALLDIR}/common/lib/ext/mysql-connector-
java-5.1.41.jar ]; then
CLASSPATH=${CLASSPATH}:${INSTALLDIR}/common/lib/ext/mysql-
connector-java-5.1.41.jar
fi
```

Configuring Trading Networks for MySQL Enterprise Edition 5.7

Important: The default value for `max_allowed_packet` is 4MB for MySQL. If this value is less than the document size, the document insertion into the database fails. To avoid this error, contact your MySQL database administrator to update the value of `max_allowed_packet` value to the largest possible document size to enable document processing and insertion into the database. The protocol limit for `max_allowed_packet` is 1GB. The value must be a multiple of 1024; non-multiples are rounded down to the nearest multiple.

You can configure Trading Networks for MySQL Enterprise Edition 5.7 using Trading Networks Server 10.1 and above.

Configuring Trading Networks for a Clustered Environment

When you cluster Trading Networks instances, you can configure the instances so that data cached in memory, user accounts, and properties are synchronized on all instances in the cluster. However, the Integration Servers that host the Trading Networks instances need not be in cluster. For more information on clustering Integration Servers, refer *webMethods Integration Server Clustering Guide*.

Use the following procedure to configure Trading Networks for a clustered environment.

Important: When you save the properties you edit in this procedure, Trading Networks will update all properties, including those that you did not edit. In a clustered environment, this will cause Trading Networks to synchronize all properties. If you want the Trading Networks instances

in the cluster to maintain different values for some properties, set the `tn.cluster.notifyPropertyChange` property to `false`, and maintain the properties for each server individually.

To configure Trading Networks for a clustered environment

1. On each Integration Server in the cluster, create a remote server alias for the other Integration Servers in the cluster. For instructions, see *webMethods Integration Server Administrator's Guide*.

Important: If you have already created a remote server alias so the Integration Server is part of a cluster without Trading Networks, you must create an additional remote server alias for the Integration Server that is specifically for the Trading Networks cluster.

2. To update the Trading Networks properties.cnf file for each Trading Networks instance in the cluster, do the following:
 - a. In My webMethods: **Administration > Integration > B2B Settings > Configure Properties**.
 - b. Add or edit these properties:

Property	Setting
<code>tn.cluster.sync.remoteAliases</code>	Remote server aliases you created in step 1.
<code>tn.query.threshold</code>	Number of rows of query results for Trading Networks to store in the session object. When clustered, Trading Networks should not store query results in the session object; set this property to "-1". for example: <code>tn.query.threshold=-1</code>

- c. Set these properties to `true`:

Set this property to true...	To notify all servers in the cluster about...
<code>tn.cluster.notifyCacheChange</code>	Change to data cached in memory.
<code>tn.cluster.notifyProfileAddUser</code>	New user accounts.
<code>tn.cluster.notifyPropertyChange</code>	Change to Trading Networks properties.

- d. Click **Save**.
3. Trading Networks data cached in memory, user accounts, and properties are automatically synchronized on all Integration Servers in the cluster. To turn one of these synchronizations off for a Trading Networks instance, modify the appropriate

property in the Trading Networks properties.cnf file. For example, if you want to manually control user account creation, turn off the `tn.cluster.notifyProfileAddUser` property.

Configuring E-mail Settings for Document Delivery

You can configure the From e-mail address to use when you deliver documents using an E-mail immediate delivery method.

To configure e-mail settings for document delivery

1. In My webMethods: **Administration > Integration > B2B Settings > Configure Properties**.
2. Add or edit the `tn.mail.from` property to the From e-mail address to use. For example, to use the From e-mail address `TNAdmin@ifc.com`, specify `tn.mail.from=TNAdmin@ifc.com`.
3. Add or edit the `tn.mail.subject` property to the subject line to use. For example, to use the subject line "Document from IFC," specify `tn.mail.subject=Document from IFC`.
4. Click **Save**.

Configuring Alias for ActiveTransfer on Remote Server

Configure the Trading Networks system property settings for an ActiveTransfer remote server alias when ActiveTransfer is installed on a remote Integration Server instance. This configuration adds ActiveTransfer as a delivery method in Trading Networks for document delivery. In a Trading Networks clustered environment, you can configure the Trading Networks system property for each instance of Trading Networks with the aliases of all the ActiveTransfer Server that are installed either on a local or remote Integration Server instance. This enables Trading Networks to share a bizDocument to an ActiveTransfer Server based on its availability using the round robin mechanism.

To configure the Trading Networks system property

1. In My webMethods: **Administration > Integration > B2B Settings > Configure Properties**.
2. In **Trading Networks Configuration Properties**, set the "`tn.delivery.active.transfer.server.aliases`" on page 304 to the Trading Networks system property with the ActiveTransfer remote server alias. You can configure multiple ActiveTransfer remote server aliases using comma separated values.

- Note:**
- If you configure a remote server alias for ActiveTransfer, then Trading Networks ignores the local installation of ActiveTransfer, if any.
 - ActiveTransfer remote server aliases provide notifications when a partner profile is added, updated, or deleted.

- ActiveTransfer remote server aliases should be part of the same ActiveTransfer Server Group.
- Trading Networks does not support the delivery of large documents to ActiveTransfer on a remote Integration Server instance. For more information about handling large documents in Trading Networks, see "[Large Document Handling](#)" on page 329.
- For more information about configuring remote Trading Networks Server in ActiveTransfer, see the ActiveTransfer documentation.

Configuring Task Settings

Configure the settings for the delivery tasks used with the Deliver Document By processing action and the reliable execution tasks used with the Execute a Service processing action.

To configure task settings

1. In My webMethods: **Administration > Integration > B2B Settings > Configure Properties**.
2. The reliable task manager thread becomes idle when all tasks have completed or failed or are in the wait period between attempts. Add or edit the `tn.task.sweepTime` property to specify the number of seconds the thread should remain idle before checking for delivery or service execution tasks that failed on their last attempt and that have not yet reached their retry limit.

Note: The thread wakes up immediately when a new task arrives.

Typically, the wait between retries value is governed by the following:

Type of task	Wait between retries governed by...
Delivery	The Wait between retries value in the receiving partner's profile.
Service execution	The <code>tn.task.ttw</code> property.

However, Trading Networks could wait as long as the `tn.task.sweepTime`. If the task manager thread is idle when the wait between retries expires, the task manager thread will not wake up to process the task until the `tn.task.sweepTime` expires.

3. Add or edit the properties below to configure the number of times Trading Networks is to try to execute services that use reliable execution and how long Trading Networks is to wait between attempts to re-execute a service that failed.

Property	Setting
<code>tn.task.maxRetries</code>	Number of times to try to re-execute a service. Specify an integer.
<code>tn.task.ttw</code>	Number of milliseconds to wait before making the first attempt to re-execute a service that has failed.
<code>tn.task.retryFactor</code>	Factor to use to determine how long to wait before making subsequent attempts to re-execute the service. Trading Networks calculates the time to wait by multiplying the last wait time by this factor. Specify a whole number greater than zero.

- By default, Trading Networks uses the Administrator user account to invoke services. Add or edit the property to indicate the user account to use when invoking the service.
- If you want to be notified when a task fails, add or set the `tn.task.notifyFailure` property to `true`. Trading Networks will publish an IS document with the format defined by the `wm.tn.rec:TaskFailure` IS document type.

Note: You can create an Integration Server trigger that subscribes to documents of this IS document type. For instructions, see the *Publish-Subscribe Developer's Guide*.

Configuring User Status

Configure the user status attribute for a transaction with predefined values. Once you configure the user status values in the **Transaction Settings** page, note that the user can only use these configured values to specify the user status for a transaction in My webMethods.

To configure user status

- In My webMethods: **Administration > Integration > B2B Settings > Transaction Settings**.
- In the **User Status Configuration** panel, perform the following actions:
 - To add a user status: Click **Add User Status**. In the Add User Status dialog box, specify a user status value in the **User Status** field and click **OK**.
 - To delete a user status: Select the corresponding check box of the user status and click **Delete**.

Configuring Trading Networks to Work with My webMethods

To use the My webMethods interface, you use the *central user management* feature. Central user management allows Integration Server to use the credentials of a My webMethods user account to authenticate users and to determine whether users are authorized for an action.

Configuring Integration Server to Use Central User Management

Configure Integration Server to use central user management. For instructions, see *webMethods Integration Server Administrator's Guide*.

Specifying the SAML Resolver Location

For Integration Server to accept an authentication token from My webMethods Server, you must indicate the location of the *SAML resolver* in My webMethods Server. By default, Integration Server uses `http://localhost:8585/services/SAML` as the SAML resolver.

To specify the SAML resolver location

1. In Integration Server Administrator, go to **Settings > Resources**.
2. Under **Single Sign On with My webMethods Server**, set the value of **MWS SAML Resolver URL** in this format:

```
http://mws_host_name_or_IP_address:mws_port/services/SAML
```

The value you provide in the **MWS SAML Resolver URL** field is assigned to the `watt.server.auth.samlResolver` watt property.

3. Repeat these steps for every Integration Server instance that will communicate with My webMethods Server.

Connecting Trading Networks Servers to My webMethods Server

You can connect one My webMethods Server to multiple Trading Networks Server instances. You can connect Trading Networks Server instances to a cluster of My webMethods Server instances.

My webMethods users with access to multiple Trading Networks Server instances can select the instance on which to act.

Note: For every user that logs on to My webMethods Server, My webMethods Server creates a session that expires only after the user logs off. If a user logs on with the user credentials of another user who has already logged on, and makes changes to the Trading Networks Server configuration, those changes do not affect the other user until the session expires. If the user deletes a

Trading Networks Server that the other user is connecting to, then the other user is not affected by the changes until the user logs out and logs in again.

To connect a Trading Networks Server to My webMethods Server

1. In My webMethods: **Administration > My webMethods > System Settings > TN Servers.**
2. Click **Add TN Configuration.**
3. Fill in these fields:

Field	Specify
TN Config Name	Name of the Trading Networks Server instance to connect to.
TN Host	Host name or IP address of the Integration Server that hosts the Trading Networks Server instance.
TN Port	Port number for the Trading Networks Server instance.
Use Secure Connection	Whether to secure communication between My webMethods and the Trading Networks Server instance.

Note: Using the My webMethods Server database as a JMS provider is removed. My webMethods Server now uses Universal Messaging Server as the JMS provider to optimize the performance of a My webMethods Server cluster. Therefore, while upgrading to Trading Networks 10.0, ensure that you configure the Universal Messaging Server as a JMS provider. Otherwise, the Partner Onboarding and Functional Permissions functionalities in Trading Networks may not work properly. For more information, see *Administering My webMethods Server*.

Configuring the Timeout for Trading Networks Web Services

When a My webMethods user issues a request that requires access to Trading Networks data, My webMethods Server executes a Trading Networks Web service on the Integration Server to perform the requested action. You configure how long My webMethods Server waits for a response from the service before timing out the requested action.

To configure the timeout

1. In My webMethods: **Monitoring > Integration > B2B > User Preferences.**
2. In the **Web Service Timeout** field in the General panel, specify the number of seconds to wait for a response.

Configuring Resubmission and Reprocessing Settings

You can configure limits for resubmitting and reprocessing documents.

To configure resubmission and reprocessing settings

1. In My webMethods: **Administration > My webMethods > System Settings > TN Servers**.
2. In the General panel, fill in these fields:

Field	Specify
Edit for Resubmit Maximum Size	Maximum size in kilobytes of documents that My webMethods users can edit and resubmit. The default is 32 kilobytes.
Reprocessing/ Resubmission Limit	Maximum number of documents that My webMethods users can resubmit or reprocess at one time. The default is 50 documents.

5 Granting Access to Trading Networks

■ Overview	92
■ Defining Roles	92
■ Adding My webMethods Users to the TN Administrators Role	93
■ Granting Users Authority to Connect Trading Networks Server Instances to My webMethods	93
■ Granting Users Access to Trading Networks Server Instances	94
■ Granting Users Authority to Execute Trading Networks Services	94
■ Granting Users General Functional Permissions	94
■ Granting Users Data Permissions	96
■ Removing References to Users	101

Overview

You grant access to Trading Networks functionality and data in My webMethods based on the role membership of My webMethods users. By default, My webMethods users have no access to Trading Networks data. To grant access, you create My webMethods user accounts for Trading Networks users, define roles, make users member of those roles, and then grant permissions to roles.

The My webMethods roles that relate to Trading Networks are as follows:

Role	Description
TN Administrators	These users can view all Trading Networks data and perform all Trading Networks actions. You manually add My webMethods users to this role.
TN MWS Users	Trading Networks automatically adds My webMethods users who have general functional permissions or data permissions to this role.
TN Partners	Trading Networks automatically adds My webMethods users who are associated with partner profiles to this role.
TN Onboarding Users	Trading Networks automatically adds My webMethods users who receive the questionnaire as a part of the onboarding process to this role.

When you first start Integration Server and Trading Networks Server after configuring central user management, Trading Networks automatically adds the My webMethods TN Administrators, TN MWS Users, and TN Partners roles to the Allowed list of the TNAdministrators, TNMWSUsers, and TNPartners ACLs, respectively.

Note: If you later want to force Trading Networks to add the roles to the ACLs again, remove the `watt.WmTN.mwsMap` property from the Integration Server `server.cnf` file and restart Trading Networks.

Defining Roles

Create My webMethods user accounts for Trading Networks users. Define the roles to which to grant Trading Networks access and add My webMethods users to the roles. For instructions, see *Administering My webMethods Server*.

Adding My webMethods Users to the TN Administrators Role

To grant a My webMethods user administrator authority to Trading Networks, you add the user to the TN Administrators role. You can do either of the following:

- Add a group or role to which the user belongs to the TN Administrator role. For instructions, see *Administering My webMethods Server*.
- Add the user to the TN Administrators role. For instructions, see the procedure below.

To add a user to the TN Administrator role

1. In My webMethods: **Administration > System-Wide > User Management > Roles**.
2. On the Roles page, click  **Edit** in the row for the **TN Administrators** role.
3. Click the **Members** tab.
4. Click **Edit Members**.
5. Use the SELECT PRINCIPAL(S) dialog to add the My webMethods user to the role. For help with this step, see information about editing the membership of a static role in *Administering My webMethods Server*.
6. Click **Apply** and then click **Save**.

Granting Users Authority to Connect Trading Networks Server Instances to My webMethods

To grant a user authority to connect Trading Networks Server instances to My webMethods Server, you grant access to the **Administration > My webMethods > System Settings > TN Servers** page.

To grant a user authority to connect Trading Networks Server instances to My webMethods

1. In My webMethods: **Administration > System-Wide > Permissions Management**.
2. In the **Advanced** tab on the Search panel, click **webMethods Applications** in the **Resource Type** list and click **Search**.
3. Move **webMethods Applications** from the **Found** list to the **Selected** list and click **Next**.
4. On the Edit Permissions panel, click **Add Users/Groups/Roles**, search for the users, groups, or roles you want, move the ones to which to grant permission to the **Selected** list, and then click **Add**.
5. Click the **Grant** check box for the **Administration > My webMethods > System Settings > TN Servers**, and then click **OK**.

Granting Users Access to Trading Networks Server Instances

You can grant a user access to one or more Trading Networks Server instances.

To grant a user access to a Trading Networks Server instance

1. In My webMethods: **Administration > My webMethods > System Settings > TN Servers**.
2. Click the **Permissions** icon for a Trading Networks Server instance to which to grant access.
3. In the Select Roles dialog box, type a string that exists in the names of the roles you want and then click **Search**.
4. Move the roles to which to grant access from the **Available** list to the **Selected** list.
5. Click **Apply** and then click **Save**.

Granting Users Authority to Execute Trading Networks Services

When a My webMethods user issues a request that requires access to Trading Networks data, My webMethods Server executes a Trading Networks service on the Integration Server to perform the requested action, and then displays the results in My webMethods. To execute the service, My webMethods Server uses the credentials of the logged-in My webMethods user. The My webMethods user must therefore have authority in Integration Server to execute the service.

In Integration Server, Trading Networks services are protected by the TNMWSUsers ACL. For My webMethods users to execute the services, the My webMethods TN MWS Users role must be in the Allowed list of the TNMWSUsers ACL, and is by default. However, you must still grant the user's role the appropriate general functional permissions and data permissions. For instructions, see "[Granting Users General Functional Permissions](#)" on page 94 and "[Granting Users Data Permissions](#)" on page 96.

Granting Users General Functional Permissions

General functional permission grant authority to perform Trading Networks functions, such as restarting Trading Networks delivery tasks.

To grant users general functional permissions

1. In My webMethods: **Administration > System-Wide > B2B Permissions > General Functional Permissions**.

2. Click  **EDIT** in the row for a general functional permission. Permissions are as follows:

<u>Functional permission</u>	<u>Actions this permission grants</u>
Manage Public Queues	Add, edit, and delete public queues.
Manage Partner Groups	Add, edit, and delete partner groups.
Delete Partner Profile	Delete partner profiles.
Manage Extended Profile Fields	Add, edit, and delete extended fields in profiles.
Add External ID Types	Create external ID types for use in profiles.
Show SQL	View the SQL that Trading Networks uses to search the database for data.
View User Preferences	View user preferences.
Edit User Preferences	Update user preferences.
View B2B properties	Configuration properties to view Trading Networks Server.
Edit B2B properties	Configuration properties to edit Trading Networks Server.
Query Expiring Certificates	Search for partner certificates that have expired or that will expire soon.
Manage TN Archiving	Add, edit, and delete the criteria and schedules for archiving and deleting documents.
Manage Delivery Methods	Add, edit, and delete immediate delivery methods created from the delivery methods that Trading Networks provides (for example, HTTP, FTP, FTPS). If this permission is not assigned, the user can only use the built-in immediate delivery methods and the immediate delivery methods that have been added by users who have this permission.

Functional permission	Actions this permission grants
Submit Documents to TN	Submit XML documents to your Trading Networks for processing. If you have webMethods Module for EDI installed, you can also submit EDI documents.
View Transaction Settings	View transaction settings configurations.
Edit Transaction Settings	Update transaction settings configurations.

- In the Select Roles dialog box, type a string that exists in the names of the roles you want and click **Search**.
- Move the roles to which to grant permission from the **Available** list to the **Selected** list.
- Click **Apply**.

Granting Users Data Permissions

Data permissions grant access to Trading Networks partner profiles, document types, processing rules, transactions, tasks, and activity log entries. To indicate the partners, document types, and processing rules to which to control access, you define a data set. You then grant data permissions to the data set using roles.

Defining a Data Set

You define a data set by specifying partners, document types, and processing rules to include or exclude from the data set. When defining data sets, consider how you want to provide access. For example:

- You might want to create a single data set that identifies all Trading Networks data so that you can assign various roles different access to all Trading Networks data.
- You might want to allow one role to only view all the data and the other role to view all the data and also resubmit and reprocess transactions. In this case, you would define two data sets.
- You might want to allow all access rights to a role, but only to a limited set of data. In this case, you would define multiple data sets, each identifying the data on which a role can act.

To define a data set

- In My webMethods: **Administration > System-Wide > B2B Permissions > Data Permissions**.
- Click **Add Data Set**.

3. In the **Dataset Name** field, type the name for the data set.
4. Click the **Dataset** tab.
5. In the Partners panel, specify the partners to include in the data set.
 - To include all partners, click **All partners**.
 - To include or exclude a subset of partners, select **Custom selection** and do the following:
 - i. Click  **Edit** above the **Partners** box.
 - ii. In the Select Members dialog box, search for the partners to include or exclude in the data set.
 - iii. Move the partners to add to the data set from the **Available Partners** the **Selected Partners** list.
 - iv. Click **OK**.
 - v. If you want to make the partner list relative to the logged-in My webMethods user, select **Self Partners**. For information about associating a My webMethods user with a partner, see "[Associate My webMethods User Accounts with a Profile](#)" on page 203.
 - vi. If you want to include the selected partners in the data set, click **Include**. If you want to exclude the selected partners from the data set and include all other partners, click **Exclude**.
6. In the Document Types panel, select the document types to include in the data set.
 - To include all document types, select **All document types**.
 - To include or exclude a subset of the document types, select **Selected document types** and do the following:
 - i. Click  **Edit** above the **Document Types** box.
 - ii. In the Select Document Types dialog box, type a string that exists in the names of the document types you want and click **Search**.
 - iii. Move the document types to include in the data set from the **Available** list to the **Selected** list.
 - iv. Click **OK**.
 - v. If you want to include the selected document types in the data set, click **Include**. If you want to exclude the selected document types from the data set and include all other document types, click **Exclude**.
7. In the Processing Rules panel, select the processing rules to include in the data set.
 - To include all processing rules, select **All processing rules**.
 - To include or exclude a subset of the processing rules, select **Selected processing rules** and follow the same steps as the previous step, but for processing rules.

8. Click **Save**.

Granting Data Permissions to a Data Set

You must grant data permissions to act on a data set to at least one role.

When you grant data permissions to act on more than one data set to a role, the users belonging to that role get the data permissions granted in all those data sets. Be careful when assigning multiple data sets to a role, especially if one of the data sets contains the **Exclude** option for the partners, document types, or the processing rules. Trading Networks limits the partner profiles, TPAs, processing rules, transactions, document types, document attributes, tasks, and activity log entries in the data set based on the partners, document types, and processing rules you specify, as follows:

- For partner profiles and TPAs in the data set, access is based on and limited to the partners specified in the data set definition. For document types in the data set, access is based on and limited to the document types specified in the data set definition. For processing rules in the data set, access is based on and limited to the processing rules specified in the data set definition.
- For transactions in the data set, access is based on partners and document types specified in the data set definition, and is limited to transactions where both of the following are true:
 - The sender or receiver of the transactions is associated with a partner specified in the data set.
 - The transactions use a document type specified in the data set.
- For tasks (delivery and/or service execution) in the data set, access is based on the partners and document types specified in the data set definition, and is limited to tasks where both of the following are true:
 - The receiver of the transaction is associated with a partner specified in the data set.
 - The transaction uses a document type specified in the data set.
- For activity log entries, access is based on partners, document types, and processing rules specified in the data set definition, and is limited to activity log entries associated with the following:
 - Events of transactions where the sender or receiver is associated with a partner specified in the data sets, and these transactions use a document type specified in the data set.
 - Profile administration audit trails of partners specified in the data set.
 - Document type audit trails of document types specified in the data set.
 - Processing rule audit trails of processing rules specified in the dataset.

- Activity logs that are not associated with any of the transactions, document type administration audit trails, processing rule administration audit trails, or profile administration audit trails.

To grant data permissions to a data set

1. In My webMethods: **Administration > System-Wide > B2B Permissions > Data Permissions**.
2. In the Data Set Name column, click the data set to which to grant data permissions.
3. Click the **Permissions** tab.
4. To add a role for the data set, click **Add Role** and do the following:
 - a. In the Select Roles dialog box, type a string that exists in the names of the roles to add and click **Search**.
 - b. Move the roles to which to grant data permissions from the **Available** list to the **Selected** list.
 - c. Click **Apply**.
5. To grant data permissions to a role:
 - a. Click  **Edit** in the row that contains the role to which to grant data permissions.
 - b. In the Edit Permissions dialog box, select the check boxes for the permissions to assign to the role and clear the check boxes for the permissions you do not want the role to have. Most data permissions are self-explanatory. Permissions that require additional explanation are covered below.

Note: Only administrators can perform actions on the Enterprise profile.

<u>Data Permission</u>	<u>Notes</u>
Create Profile	<p>Edit Profile Summary and Edit External IDs are granted automatically after you save and re-open the data set.</p> <p>Note: By default, profiles created by this user are added to all data sets to which the user has this permission. If the data set has the All Partners option selected, the profile of the partner will always belong to that data set, and the My webMethods+ user does not need to explicitly add it to the data set.</p>
View Transactions	View Profile, View Document Type, and View Document Attributes are granted automatically.
View Activity Log	View Profile and View Document Type data are granted automatically.

Data Permission	Notes
View Task	View Profile and View Document are granted automatically.
Create Document Types	View Document Types data permission is granted automatically. Note: By default, document types created by this user are added to all data sets to which the user has this permission. If the data set has the All document types option selected, the new document types will always belong to that data set, and the user need not explicitly add it to the data set.
View Document Types	View Document Attributes data permission is granted automatically.
Manage Attributes	View, create, edit, enable, and disable document attributes.
Create Trading Partner Agreements	Includes permission to delete TPAs.
Create Processing Rules	Includes permission to delete processing rules. Note: By default, processing rules created by this user are added to all data sets to which the user has this data permission. If the data set has the All processing rules option selected, the processing rule of the partners will always belong to that data set, and the user need not explicitly add it to the data set.
Manage Partner Access to APIs	Add APIs and view APIs data permissions are granted to manage API access to partners.
Note:	Even if a user does not have the Create Document Types data permission to a particular document type, Trading Networks might list that document type in search results or as related documents in some My webMethods pages. However, the user will not be able to view or edit the details of the documents without those data permission.

Setting Partner Permissions

Trading Networks lists all the data sets for which the My webMethods user has the Create Profile permission. You can add a partner to a data set by selecting the check box for the data set. If the check box for the data set name is not selected, the partner does not belong to the data set. To add the partner to all the data sets, select the **Rule Name** check box (Integration > B2B > Partner Administration > Partner Profiles > Partner Details > Permissions tab).

The partner must be in at least one data set where the My webMethods user has Create Profile permission. If the My webMethods user tries to remove the partner from all data sets, and that partner does not belong to any data set that has the **All Partners** option selected, an error message is displayed and the changes are not saved.

Trading Networks lists only those data sets that have the role in the data set that has the Create Profile permission and has the **Custom Selection** option selected. It does not list the data sets if the role has the **All Partners** option selected even if the role has the Create Profile permission.

Administrators can edit the data sets using the **Dataset** tab.

Removing a Data Set and Its Data Permissions

You can remove a data set and all its data permissions.

1. In My webMethods: **Administration > System-Wide > B2B Permissions > Data Permissions**.
2. Select the data set to remove.
3. Click **Delete**.

Removing References to Users

To remove a user from Trading Networks, you must delete the user from the Users page and remove the references associated with the user from the corresponding **Partner Profiles**.

To remove references to a user from Trading Networks

1. In My webMethods: **Administration > System-Wide > B2B Permissions > Data Permissions > User Management > Users**, select the user.
2. Click **Delete**.
3. Go to **Administration > Integration > Partner Administration > Partner Profiles** and select the partner profile in which the user is referenced.
4. In the **Users** tab of the **Partner Profile**, select the user and click **Remove**.
5. Click **Save** or **Save and Close**.

For more information on removing user data from Integration Server that hosts Trading Networks, see *webMethods Integration Server Administrator's Guide*. And for more information on General Data Protection Regulation (GDPR) settings, see *Administering My webMethods Server*.

6 Creating Custom Document Attributes

- Create Custom Attributes 104

Create Custom Attributes

Create custom attributes for all the types of documents you expect to receive; attributes are not associated with a specific document type. Later, when you define your document types, you will specify the attributes to use as criteria and to extract.

To create custom attributes

1. In My webMethods: **Administration > Integration > B2B > Document Types Administration > Document Attributes**.
2. In the Document Attributes Search Result panel, do one of the following:
 - To create a custom attribute by providing new information, click **Add**.
 - To create a custom attribute by copying an existing attribute and modifying it, click  in the row for the attribute to copy.
3. Give the attribute a unique name and, optionally, a description, and specify the attribute's data type. Then click **OK**.

Note: There are certain words reserved in Trading Networks that you need to avoid while creating custom attributes because they might be a part of system attributes, database columns, to name a few. Few words that you must avoid include DocID, DocTimestamp, TypeName, SenderID, SenderCorp, SenderUnit, ReceiverID, ReceiverCorp, ReceiverUnit, RoutingStatus, UserStatus, NativeID, GroupID, ConversationID, Comments, SenderProfileGroup, ReceiverProfileGroup, DocTypeID, NS Name.

7 Creating XML Document Types

- Before You Begin 106
- Creating an XML Document Type 108
- Editing an XML Document Types 118
- Testing XML Document Types 118

Before You Begin

Obtaining XML Document Samples

When you create an XML document type, you specify XQL queries. The queries do the following:

- Search for specific document nodes to determine whether the document matches the document type.
- Locate attributes you want to extract from the document.

If you have a sample document, you can simplify the creation of the XQL queries; you can highlight the node in the sample document and Trading Networks will form the query for you. If you do not have a sample document, you must manually create the XQL queries.

Obtain a sample document for each variation of XML document type you expect Trading Networks to receive and place the samples in a directory you can access from My webMethods.

Creating IS Schemas

Create an XML schema for each variation of XML document type you want to validate. For instructions, see *webMethods Service Development Help*.

Creating IS Document Types

Create an IS document type for each variation of XML document type that you want to transform into IS documents (IData objects).

Creating Custom Transformations for Extracted Attributes

You can create your own services to transform extracted attributes. Base each service on the appropriate specification as listed below.

Data Type	Specification to Use
DATETIME or DATETIME LIST	wm.tn.rec:DateAttributeTransformService
STRING or STRING LIST	wm.tn.rec:StringAttributeTransformService
NUMBER or NUMBER LIST	wm.tn.rec:NumberAttributeTransformService.

Trading Networks passes these input variables:

Input Value	Meaning
<i>values</i>	One or more values for your service to transform. The values might be null or an empty (zero length) string. For an array data type (for example, DATETIME LIST), any element in the <i>values</i> array can be null or an empty string. Make sure your custom transformation service can handle null values or empty strings.
<i>isArray</i>	Whether <i>values</i> contains a single value (<code>false</code>) to transform or multiple values (<code>true</code>) to transform.
<i>arg</i>	Optional arguments you define.

Trading Networks expects you to return the transformed values in the *newValues* output variable:

Return...	For...
<i>Date List</i>	DATETIME or DATETIME LIST
<i>String List</i>	STRING or STRING LIST
<i>Number List</i>	NUMBER or NUMBER LIST

If the input variable *isArray* is `false`, your service should transform the single item and save just a single item in the output variable *newValues*.

If the input variable *isArray* is `true`, your service should loop through the items in the *values* variable and transform each one. It should store each as an item in the output variable *newValues*. For example, if you are transforming string data, *newValues* is a *Date List* of the transformed *Strings*.

Creating a Service for the Check for Duplicate Document Action

In the Check for Duplicate Document pre-processing action, Trading Networks offers options to check based on Document ID, sender, receiver, or document type. If you want to check based on other attributes, you can use a custom service. Create the service using the `wm.tn.rec:DupCheckService` specification.

Note: Custom duplicate checking services can affect performance, depending on the service logic.

Creating an XML Document Type

Creating and Naming the Document Type

To create and name the document type

1. In My webMethods: **Administration > Integration > B2B > Document Types administration > Document Types**.
2. Do one of the following:
 - To create the document type by providing new information, click **Add Document Types** and select **XML**.
 - To create the document type by copying an existing type and modifying it, click  **Copy** on the row for the XML document type to copy.
 - To create the document type based on a transaction for which the document type is Unknown, do the following:
 - i. In My webMethods: **Monitoring > Integration > B2B > Transactions**.
 - ii. Click  **Show Actions** on the row for the Unknown document to use.
 - iii. Select **Create Document Type**.
 - iv. In the Create Document Type dialog box, click **XML**.
3. In the Document Types Details page, provide a name and description, and indicate whether to enable or disable the document type. When a document type is disabled, Trading Networks does not match inbound documents against it. Do not enable the document type until you are done creating it.

Specifying Criteria for Matching Inbound Documents to the Document Type

You can define XML document types to be very general or very specific. For example, you can define an XML document type that recognizes OAG documents, or OAG PROCESS_PO_004 documents, or OAG PROCESS_PO_004 documents from a specific sender.

Click the Identify tab. If you are using a sample document, identify the sample document in the **File** field and click **Upload**. Trading Networks displays a tree view of the sample XML document in the left panel.

Note: If the sample XML document contains a large number of subnodes, My webMethods might take a long time to load the tree view.

For a document to match a document type, specify at least one criterion to match. In case of more than one criteria, it must match all the identification criteria you specify. If you specify the Root tag and DOCTYPE criteria, Trading Networks matches inbound XML documents to those criteria first. If these match, Trading Networks also checks the identifying XQL queries and values, and the pipeline matching. Matching on XQL queries and its values is case-sensitive.

Specifying the Root Tag

You can specify the value that inbound documents must have in the root tag to match the document type.

- If you are using a sample document, click **Set** next to the **Root Tag** field. Trading Networks populates the fields with the value of the root tag of the sample document.
- If you are not using a sample document, type the root tag value. For example, to identify a document that has the document type declaration `<!DOCTYPE cXML SYSTEM "cXML.dtd">`, you would enter the following:

```
<?xml version="1.0"?>
<orderInfo>
```

Specifying the DOCTYPE Identifier

You can specify the system or public identifier from the DOCTYPE declaration that inbound documents must have to match the document type.

- If you are using a sample document, click **Set** next to the **DOCTYPE** field. Trading Networks populates the field with the value that follows the SYSTEM or PUBLIC literal string in the document type declaration (DOCTYPE) in the sample document. If both literal strings exist, Trading Networks uses the value that follows the SYSTEM literal string.
- If you are not using a sample document, type the value that follows the SYSTEM literal string. For example, to identify a document that has the DOCTYPE declaration `<!DOCTYPE cXML SYSTEM "cXML.dtd">`, you would type `cXML.dtd`.

Defining Identifying Queries

You can specify that certain nodes must be present for inbound documents to match the document type. You must also specify the values those nodes must have. To do so, you define identifying XQL queries.

For example, the identifying query below specifies that the OrderRequest tag must be present:

```
/cXML[0]/Request[0]/OrderRequest[0]
```

The identifying query below specifies that the Identity tag within the Credential tag that is within the Sender tag must be present, and the Identity tag must evaluate to XYZ Steel Company:

```
/cXML[0]/Header[0]/Sender[0]/Credential[0]/XYZ Steel Company
```

For complete information about identifying queries, see ["XQL Reference" on page 321](#).

- If you are using a sample document, select in the left panel the node for which to define an identifying query, click  **More Actions** next to the **Query** field, and select **Add Query**. In the Add Identifying Query dialog box, Trading Networks fills in the XQL query and value for the node.

Note: The node is applicable for DATETIME, STRING, or NUMBER data types. If you want to use DATETIME LIST, STRING LIST, NUMBER (with Average, Minimum, Maximum, and Sum), or NUMBERLIST data types, update the query so that it identifies an array of values.

- If you are not using a sample document, click **Add** in the Identifying Queries panel on the right. In the Add Identifying Query dialog box, type the query and, optionally, the value.

Specify Pipeline Variables

You can specify pipeline variables that must be present in inbound documents to match the document type. Criteria for matching inbound documents to the document type are name/value pairs, where values are optional. If you do not specify a value, the document matches the document type if the pipeline has a variable that matches the specified name, regardless of the variable's value. If you specify a value, the variable must have the specified value. The variables are inserted into the pipeline by the service that sends the document to Trading Networks.

To add a pipeline variable, click **Add** in the Pipeline Matching panel on the right. In the Add Pipeline Variable dialog box, type the variable name and optionally, the value.

Specifying Document Attributes to Extract

In your document type, specify the system and custom document attribute values to extract from documents that match this document type. If you want to transform extracted attributes, you can use built-in or custom transformations.

Click the **Extract** tab.

- If you are using a sample document, select in the left panel the node for which to define an XML query, click  **More Actions** next to the **Query** field, and select **Add Query**. In the Add Extracted Attribute dialog box, Trading Networks fills in the XQL query for the node. Also provide the information in the sections below in the dialog box.

Note: This node is applicable for DATETIME, STRING, or NUMBER data types. If you want to use DATETIME LIST, STRING LIST, NUMBER (with Average, Minimum, Maximum, and Sum), or NUMBERLIST data types, update the query so that it identifies an array of values.

- If you are not using a sample document, click **Add** in the right panel. In the Add Extracted Attribute dialog box, type the XQL query. Also provide the information in the sections below in the dialog box.

Specifying the Attributes

SenderID and ReceiverID System Attributes

When Trading Networks has a value for SenderID or ReceiverID, you can perform these tasks:

- Use the sender or receiver to determine whether a document is a duplicate (Check for Duplicate Document pre-processing action).
- Use the sender or receiver as criteria for determining the processing rule to use for the document.
- Deliver a document.
- Search for the document in My webMethods based on sender or receiver.
- Have a business process perform steps based on sender or receiver.

If you extract the SenderID or ReceiverID system attributes from a document, you must specify how the sender or receiver is identified in the document. For example, if you want to extract ReceiverID from the XML document and the ReceiverID is identified in the document by a D-U-N-S number, define the XQL query to identify the location of the receiver information and select the external ID type DUNS. For example:

```

XQL Query for ReceiverID = /PurchaseOrder[0]/OrderHeader[0]/DUNS[0]
<PurchaseOrder>
  <OrderHeader>
    <CorpName>XYZ Steel Company</CorpName>
    <DUNS>123456789</DUNS>
  </OrderHeader>
  .
  .
  .
</PurchaseOrder>

```

The XQL query points to the <DUNS> tag. Trading Networks stores the value 123456789 for the ReceiverID attribute and associates it with the external ID type DUNS.

Be sure to select an external ID type that your partners specify in their profiles. You can choose from a list of external ID types. If the external ID type you need is not in the list or if you have agreed on a non-standard format with a trading partner, select the external ID type **Mutually Defined** and add an appropriate ID value.

DocumentID

Extract this system attribute if you want to use the document ID to determine whether a document is a duplicate (Check for Duplicate Document pre-processing action), and if you want to search for documents of this document type in My webMethods based on document ID.

UserStatus

Extract this system attribute if you want to specify user status as a criterion for determining the processing rule to use for the document, and if you want to search for documents of this document type in My webMethods based on user status.

GroupID

Extract this system attribute if you want to search for documents of this document type in My webMethods based on group ID.

ConversationID

Extract this attribute if you want to pass the document to a business process.

If you do not want to pass the document to a business process, add a variable to the pipeline called *prtIgnoreDocument* with the value `true`. If *ConversationID* is extracted and *prtIgnoreDocument* is set to `true`, it will not be passed to the process model.

SignedBody and Signature

Extract these attributes if you want to verify the digital signature of a document.

Custom Attributes

Extract any custom attribute you want to use later.

Indicating Whether the Attribute is Required

By default, when Trading Networks cannot extract an attribute, it continues processing, and, if processing completes successfully, sets the processing status to DONE. It does not log an error to the activity log. If you want Trading Networks to log an error to the activity log when it cannot extract an attribute, you can designate that attribute as required for extraction. If processing completes successfully, Trading Networks sets the processing status to DONE W/ERRORS.

Configure the document type's processing rule to handle the error as follows:

- Use the processing rule's recognition error criterion to trigger the processing rule based on whether it contains errors.
- Use the processing rule's Execute a Service processing action to inspect the *bizdoc/Errors* pipeline variable for errors.

Transforming Extracted Attributes

You can have Trading Networks transform extracted attributes before storing them in the Trading Networks database.

Note: You cannot transform the SenderID or ReceiverID attributes.

Built-In Transformations for DATETIME and DATETIME LIST Data Type Attributes

If you extract an attribute that has the data type DATETIME or DATETIME LIST, you must identify the date format. You can choose a built-in common date/time format or you can type the format you need using a pattern string based on the "Time Format Syntax" described for the `java.text.SimpleDateFormat` class. Trading Networks extracts the value of the date and uses it and the pattern you specify to decode the value and convert it to the format that Trading Networks requires to store the date in the BizDocEnvelope.

Built-In Transformations for STRING or STRING LIST Data Type Attributes

If you extract an attribute that has the data type STRING or STRINGLIST, you can transform the string value using one of the built-in transformations below.

Format	Description
Uppercase	Transforms the extracted string attribute value to all uppercase.
Lowercase	Transforms the extracted string attribute to all lowercase.
String Substitution	<p>Substitutes extracted values with a pattern you specify. Trading Networks uses the <code>java.text.MessageFormat</code> class to perform this transformation.</p> <p>For example, you might specify the pattern <code>Items Ordered: {0}, {1}, {2}</code>, and specify an XQL query for the attribute that extracts this array of values:</p> <pre>0 Cellular phone 1 Belt clip 2 Rapid mobile charger</pre> <p>Trading Networks would store the following for the value of the attribute:</p> <pre>Items Ordered: Cellular phone, Belt clip, Rapid mobile charger</pre>

Format	Description
	If you place more arguments in the pattern than there are extracted values, the string stored in the database for the attribute will contain the extra arguments. If you specify fewer arguments than there are extracted values, the string will contain only the values for the number of arguments.

Built-In Transformations for NUMBER Data Type Attributes)

If you extract an attribute that has the data type NUMBER and contains an array of numbers, you can use transform the array into a single value using one of the built-in transformations below.

Format	Description
Average	Calculates the average value of all the numbers in the array.
Minimum	Calculates the smallest number in the array.
Maximum	Calculates the largest number in the array.
Sum	Calculates the sum of all the numbers in the array.
No format	Stores the first value of the array as the value of the attribute.

Note: When Trading Networks extracts a NUMBER or NUMBER LIST from a document, it uses the number parsing behavior of `java.lang.Number`. For example, if the NUMBER or NUMBER LIST contains the value `100zzz`, Trading Networks interprets the value as `100`, instead of throwing an error as it would if the value were `zzz100`. If you want to redefine this parsing behavior, you can write a custom transformation service.

Suppose your gateway service extracts the quantity of each item of a purchase order, and you want to sum the quantities of all items. Suppose you defined the following attributes in `TN_parms` and in the document type:

TN_parms Attributes	Document Type
<code>/TN_parms/itemNumber</code>	<code>itemNumber NUMBER</code>
<code>/TN_parms/quantity</code>	<code>quantity NUMBER</code>
<code>/TN_parms/totalQuantity</code>	<code>totalQuantity NUMBER</code> <code>summedValue NUMBER LIST</code>

If you select the **Sum** built-in transformation for *totalQuantity*, it will sum its values and produce a single number. For example, if the value of *totalQuantity* is this array:

```
/TN_parms/totalQuantity = {2, 3, 4, 5}
```

The value of *summedValue* will be:

```
/TN_parms/summedValue = 2 + 3 + 4 + 5
```

Trading Networks will store the value 14 for *totalQuantity*.

Custom Transformation Services

Click  and select the custom transformation service to use. To locate a service, specify the **Package Name** or **NSName** and click **Search** or browse the services in the listed packages. Specify the inputs for the selected service.

Specifying the Namespace Mapping

If an XML document uses namespaces, the elements in that document might be prefixed with a string. When you create XQL queries to identify elements within the document, the XQL queries must include the prefix. If XML documents use equivalent namespaces but have different prefixes, you must define *namespace mappings* for Trading Networks to correctly locate the nodes identified by the XQL queries. Namespace mappings identify all prefixes that identify the same namespace (that is, point to the same URI).

Include a namespace mapping for each prefix/URI combination that you expect to receive in XML documents from your trading partners. If Trading Networks receives a document that uses a prefix that is not defined in the namespace mappings table, it performs a literal match of the XQL queries against the document.

If you are using a sample document, you can have Trading Networks populate the namespace mappings table with the prefix and URIs from the `xmlns` attributes that define namespaces in the sample document. If you are not using a sample document, you can define namespace mappings manually.

Trading Networks uses the literal string `prefix0` to identify the default namespace. To define a mapping for the default namespace, which does not use a prefix, use the literal string `prefix0` for the prefix. Do not use `prefix0` as the prefix in any document.

To create the namespace mapping

1. Click the **Namespaces** tab.
2. Do one of the following:
 - If you are using a sample document, click  **More Actions** next to the **Query** field and click **Append Namespace Mapping**.

Note: If you need to replace existing namespace mappings, use **Set Namespace Mapping**.

- If you are not using a sample document, do the following:

- i. Click **Add** in the right panel.
 - ii. In the **Prefix** field, type the prefix for the namespace. If the mapping is for the default namespace, specify `prefix0`.
 - iii. In the **URI** field, type the URI for the namespace.
3. Click **OK**.

Defining Pre-Processing Actions for the Document Type

Click the **Options** tab and define the options and actions.

Format as an IS Document Type Option

This option invokes the `wm.tn.doc.xml:bizdocToRecord` built-in service.

If you know the fully qualified name of the IS document type to use, type the name. Otherwise, to locate the IS document type, specify the **Package Name** or **NSName** and click **Search** or browse the IS document types in the listed packages.

Enable Processing Rule Routing Option

If you want to process documents using pre-processing and processing actions defined in a processing rule, enable processing rule routing. If you want to process documents using only pre-processing actions defined in the document type, disable processing rule routing.

Note: If you do not want to process documents using a processing rule, disabling routing can improve performance by avoiding the search for processing rules.

Verify Digital Signature Pre-Processing Action

This action executes a verification service to verify the digital signatures of documents.

To use this action, the document type must specify extraction of the `SignedBody` and `Signature` system attributes, and the signature must be a PKCS#7 detached signature of the signed body. In addition, the profile for each partner whose digital signature you want to verify must specify a certificate. Trading Networks makes sure the signed body has not changed by verifying the digital signature. To verify that the sender is who it claims to be, Trading Networks matches the certificate from the digital signature to the certificate in its database for the sender.

Validate Structure Pre-Processing Action

This action invokes the `Integration Serverpub.schema.validate` built-in service.

If you know the folder that contains the IS schema to use, click  and select the IS schema. To locate a schema, specify the **Package Name** or **NSName** and click **Search** or browse the schemas in the listed packages.

Important: If the structure of a document is not valid, the validation service can return many error messages. To limit the number of error messages, set the `tn.doc.validate.max_errs` property.

Check for Duplicate Document Pre-Processing Action

Trading Networks provides options to check for a document with the same Document ID; the same Document ID and sender; the same document ID, sender, and receiver; or the same document ID, sender, and document type. In the extraction specifications, you must specify extracting the necessary attributes for the option you select.

If you want to use a custom service that performs the check based on other attributes, click **Select** and search or browse for the service.

Note: Custom duplicate checking services can affect performance, depending on the service logic.

Save Document to Database Pre-Processing Action

You must save a document to the database when you want to:

- Deliver the document using reliable delivery with immediate or scheduled delivery.
- Deliver the document using queue for polling.
- Pass a document to a business process.
- Send a document back to the beginning as a “new” document (for example, because the document did not match any defined document type). This is called resubmitting the document.
- Send a document back through processing rules (for example, because the document was processed by the wrong rule). This is called reprocessing the document.

You can choose to save all documents, no documents, or only unique documents. If you want to save only unique documents, you must also specify the Check for Duplicate Document pre-processing action.

Option	Description
Content	Save document content if you want to: <ul style="list-style-type: none"> ■ View the content in My webMethods when monitoring transactions. ■ Resubmit or reprocess the document. ■ Use the Deliver Document By processing action with reliable delivery.

Option	Description
Attributes	<p>Save custom attributes if you want to view attributes in My webMethods when monitoring transactions, search for documents based on custom attributes, or reprocess documents.</p> <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p>Important If you reprocess a document without saving custom attributes, you might get unexpected results. Because the attributes are not saved, the document will not match processing rules that use extended criteria. Instead, the document will match another processing rule, such as the default processing rule, and Trading Networks will perform the processing actions defined in that rule.</p> </div>
Activity Log	Record activity log entries when monitoring transactions.

Editing an XML Document Types

While editing an XML document type, if another user tries to edit at the same time, Trading Networks ensures that data written by you is not over written by the other user and vice versa. While saving the document type, Trading Networks ensures that you are working on the latest version of data, and notify you if another transaction has modified the document type. In that case, you have to cancel the changes you have made, and reopen the document type for edit.

To edit an XML document type

1. In My webMethods: **Monitoring > Integration > B2B > Transactions**.
2. If you want to edit the document type based on a transaction, click  **Show Actions** on the row for the transaction to use and then select **Edit Document Type**.
3. Edit the fields as described in "[Creating an XML Document Type](#)" on page 108.
4. Save the document type.

Testing XML Document Types

If you have samples of documents Trading Networks will process, you can test to determine whether each document:

- Matches exactly one XML document type. This result means the XML document types are set up correctly for the document.

- Does not match any XML document type. Trading Networks would consider such a document an unknown document type. Create a document type or update an existing document type to identify the document.
- Matches more than one XML document type. Trading Networks would consider the document an unknown document type because it cannot determine which document type to use. You can either update the XML document types so that the document matches exactly one XML document type, or you can configure Trading Networks to use the first document type that matches an inbound document by setting the `tn.chart.fetchMaxRows` property to `false`.

Note:

- Trading Networks does not test documents against disabled document types.
- Trading Networks does not actually process the document you are testing. That is, Trading Networks does not perform any pre-processing or processing actions on the document.

To test a document type

1. Place the sample documents you want to test in your file system.
2. In My webMethods: **Administration > Integration > B2B > Document Types administration > Document Types**.
3. Click the **Advanced** tab.
4. Click the **Test** tab.
5. Click **Browse** and select the sample document to test.
6. Click **Test**.

Trading Networks displays all document types that match the sample document.

8 Creating Flat File Document Types

- Before You Begin 122
- Creating a Document Gateway Service 123
- Creating a Flat File Document Type 127
- Editing a Flat File Document Type 135

Before You Begin

Registering Flat File Content Types

By default, Trading Networks considers inbound documents with the “text/plain” content type to be flat file documents. You can register other content types as flat file documents using the [tn.ff.contenttypes](#) property.

Creating Flat File Schemas

Create a flat file schema for each flat file document you want to validate. For more information about flat file schemas and parsing, see the *Flat File Schema Developer’s Guide*.

Creating Custom Transformations for Extracted Attributes

You can create your own services to transform extracted attributes. Base each service on the appropriate specification as listed below.

Data Type	Specification to Use
DATETIME or DATETIME LIST	wm.tn.rec:DateAttributeTransformService
STRING or STRING LIST	wm.tn.rec:StringAttributeTransformService
NUMBER or NUMBER LIST	wm.tn.rec:NumberAttributeTransformService.

Trading Networks passes these input variables:

Input Value	Meaning
<i>values</i>	One or more values for your service to transform. The values might be null or an empty (zero length) string. For an array data type (for example, DATETIME LIST), any element in the <i>values</i> array can be null or an empty string. Make sure your custom transformation service can handle null values or empty strings.
<i>isArray</i>	Whether <i>values</i> contains a single value (<code>false</code>) to transform or multiple values (<code>true</code>) to transform.
<i>arg</i>	Optional arguments you define.

Trading Networks expects you to return the transformed values in the *newValues* output variable:

Return...	For...
<i>Date List</i>	DATETIME or DATETIME LIST
<i>String List</i>	STRING or STRING LIST
<i>Number List</i>	NUMBER or NUMBER LIST

If the input variable *isArray* is `false`, your service should transform the single item and save just a single item in the output variable *newValues*.

If the input variable *isArray* is `true`, your service should loop through the items in the *values* variable and transform each one. It should store each as an item in the output variable *newValues*. For example, if you are transforming string data, *newValues* is a *Date List* of the transformed *Strings*.

Creating a Service for the Verify Signature Action

You can use a custom service to verify the digital signatures of inbound flat file documents. The `wm.tn.rec:BizDocVerificationService` specification defines the service input and output variables.

Creating a Service for the Check for Duplicate Document Action

In the Check for Duplicate Document pre-processing action, Trading Networks offers options to check based on Document ID, sender, receiver, or document type. If you want to check based on other attributes, you can use a custom service. Create the service using the `wm.tn.rec:DupCheckService` specification.

Note: Custom duplicate checking services can affect performance, depending on the service logic.

Creating a Document Gateway Service

You can create a document gateway service using Java or flow. The `wm.tn.rec:GatewayService` specification defines the service input and output variables.

Create the service using the `wm.tn.rec:BizDocVerificationService` specification for document verification.

- The input variable *ffdata* is the document sent by the partner through the gateway service in the form of an input stream.
- The output variables are *ffdata* (an object) and *TN_parms* (an IS document).

Note: You must provide your partners with the names of the document gateway service or services for their document types.

Reading the Input Stream

If the document gateway service is a Java service, use the methods of the `java.io.InputStream` class to read bytes from the input stream contained in *ffdata*. See your Java documentation for descriptions of these methods.

Important: Remember to use `InputStream`'s `mark` and `reset` methods so that Trading Networks will be able to read from the beginning of the stream. If `InputStream` does not support marking, make sure you transform `InputStream` to support marking. Otherwise, Trading Networks will lose the bytes read by the gateway service when it tries to save the document content to the `BizDocEnvelope`.

If the document gateway service is a flow service, use the services in the `pub.io` folder in the Integration Server `WmPublic` package. The `WmPublic` package must be enabled.

You can parse the input stream and convert it into an IS document from which you can map data to the *TN_parms* pipeline variable. In the document gateway service, invoke the `pub.flatfile:convertToValues` service in the `WmFlatFile` package. This service invokes the flat file parser to validate the document's structure and content, based on the flat file schema you identify as input to the service. Trading Networks can save or throw away the parsed contents. For information about the `pub.flatfile:convertToValues` service, see the *Flat File Schema Developer's Guide*. Alternatively, you can convert the document using the `Validate Structure` pre-processing action. Do not convert a document in both the document gateway service and using the `Validate Structure` action. If you convert a document in the gateway service, it will remain in the pipeline.

Note: Parse the document only if absolutely necessary. Parsing can be time-consuming and return very large outputs that consume much memory.

Specifying and Passing the Outputs

A document gateway service can return the output variables *ffdata* (an object) and *TN_parms* (an IS document). *TN_parms* can contain system attributes, system variables, and custom attributes.

DoctypeID or DoctypeName

This variable is a string, and is optional. It is the Trading Networks-generated internal identifier or name of the flat file document type to use for the document. Specify the document type if you want to avoid the overhead of searching for a document type. For example, you would specify:

```
/TN_parms/DoctypeID = "5066kt00tttva37r00000009"  
/TN_parms/DoctypeName = "xCBL_PO"
```

To view the names and IDs of your document types, open Software AG Designer and invoke the `wm.tn.doctype:list` service.

Note: Because document type IDs cannot be changed, your gateway service will be more stable and efficient if you use *DoctypeID* rather than *DoctypeName*. If you specify both variables, *DoctypeID* is used.

SenderID and ReceiverID

These variables are strings, and are required. The *SenderID* variable is an external ID value specified in the sender's profile (for example, the sender's D-U-N-S number). The *ReceiverID* is an external ID specified in the receiver's profile.

DocumentID

This variable is a string, and is optional. It is the identifier of the flat file document. Specify this variable if you want to use the document ID to determine whether a document is a duplicate (Check for Duplicate Document pre-processing action), or if you want to search for documents in My webMethods based on document ID.

GroupID

This variable is a string, and is optional. If you expect to receive more documents in the same group, make sure the value you specify for *GroupID* is reproducible for those documents. The same *GroupID* is being used in all the transactions.

Specify this variable if you want to search for documents of this document type in My webMethods based on group ID.

ConversationID

This variable is a string, and is optional. Specify this variable if you want to pass the document to a business process. If you expect to receive more documents in the same conversation, make sure the value you specify for *ConversationID* is reproducible for those documents. The same *ConversationID* is being used in all the transactions.

If you do not want to pass the document to a business process, add a variable to the pipeline called *prtIgnoreDocument* with the value `true`.

UserStatus

This variable is a string, and is optional. Specify this variable if you want to specify user status as a criterion for determining the processing rule to use for the document, and if you want to search for documents of this document type in My webMethods based on user status.

processingRuleID* or *processingRuleName

This variable is a string, and is optional. It is the Trading Networks-generated internal identifier or name of the processing rule to use for the document. Specify the rule if you want to avoid the overhead of searching for a document type.

To view the names and IDs of your processing rules, open Software AG Designer and invoke the `wm.tn.route:list` built-in service.

Note: Because process rule IDs cannot be changed, your gateway service will be more stable and efficient if you use *processingRuleID* rather than *processingRuleName*. If you specify both variables, *processingRuleID* is used.

\$receiveSvc

This variable is a string, and is optional. It is the fully qualified name of the document gateway service you are creating. Save the name of the gateway service if you want to be able to resubmit the flat file document.

\$contentType* and *\$contentEncoding

These variables are strings, and are optional. It is the content type and encoding to use for the flat file document. You can specify either or both variables.

If your partners send documents directly to your document gateway service (as opposed to through an intervening service), Trading Networks tries to determine the content type and content encoding of the submitted flat file document.

If you are using the gateway service directly, you need to specify content type and content encoding, else it will assume a default value, which is `application/x-wmidatabin` for content type and `UTF8` for content encoding. If you are directly sending the document to the gateway service, you should specify the content type and content encoding. You can provide the content type and content encoding to make sure Trading Networks uses accurate values.

For example, if a document with a content type of `text/plain` and an encoding of `SJIS` is submitted to a service and that service passes the document into the pipeline to your document gateway service, the original content type and content encoding will be lost. The pipeline has a content type of `application/x-wmidatabin` and a content encoding of `UTF8`. In this case, your gateway service or the service that invoked it should specify `text/plain` for `TN_parms/$contentType` and `SJIS` for `TN_parms/$contentEncoding`.

Note: When you send a document to Trading Networks, if the header of the document contains an encoding attribute with charset value (such as ANSI), Trading Networks uses the same charset to convert the document into bytes. If no *encoding* is set in the header, Trading Networks uses default UTF-8 charset to convert it into bytes.

Custom Attributes

Specify all custom document attributes to extract. Make sure the name and case you use for the *TN_parms* pipeline variable exactly match the attribute name as it is defined in the document type.

Pass the Outputs to Trading Networks

Invoke the `wm.tn.doc.xml:routeXML` or `wm.tn.doc.ff:routeFlatFile` built-in service to pass *ffdata* and *TN_parms* to Trading Networks.

Creating a Flat File Document Type

Creating and Naming the Document Type

To create and name the document type

1. In My webMethods: **Administration > Integration > B2B > Document Types Administration**.
2. Do one of the following:
 - To create the document type by providing new information, click **Add Document Types** and click **Flat File**.
 - To create the document type by copying an existing type and modifying it, click  **Copy** on the row for the flat file document type to copy.
 - To create the document type based on a transaction for which the document type is Unknown, do the following:
 - i. In My webMethods: **Monitoring > Integration > B2B > Transactions**.
 - ii. Click  **Show Actions** on the row for the Unknown document to use.
 - iii. Click **Create Document Type**.
 - iv. In the Create Document Type dialog box, click **Flat File**.
3. In the Document Types Details page, provide a name and description, and indicate whether to enable or disable the document type. When a document type is disabled, Trading Networks does not match inbound documents against it. Do not enable the document type until you are done creating it.

Specifying Criteria for Matching Inbound Documents to the Document Type

Criteria for matching inbound documents to the document type are key/value pairs, where values are optional. If you do not specify a value, the document matches the document type if *TN_parms* has a variable that matches the specified key, regardless of the variable's value. If you specify a value, the variable must have the specified value.

Note: You must specify at least one key or Trading Networks will not use the document type.

Click the **Configure** tab. In the Pipeline Matching panel, click **Add** to specify the key and, optionally, the value for each *TN_parms* variable to use to match inbound documents to the document type.

Specifying Attributes to Extract

In your document type, specify the system and custom attributes to extract from the *TN_parms* variable for documents that match this document type. If you want to transform extracted attributes, you can use built-in or custom transformations.

In the Extracted Attributes panel, click **Add** and provide information in the Add Extracted Attribute dialog box.

Specifying the Attributes

SenderID and ReceiverID System Attributes

When Trading Networks has a value for SenderID and ReceiverID, you can perform these tasks:

- Use the sender or receiver to determine whether a document is a duplicate (Check for Duplicate Document pre-processing action).
- Use the sender or receiver as criteria for determining the processing rule to use for the document.
- Deliver a document.
- Search for the document in My webMethods based on sender or receiver.
- Have a business process perform steps based on sender or receiver.

On the Extract panel, select the external ID type (for example, **DUNS**) for each attribute so that Trading Networks can match the external ID in *TN_parms/SenderID* and *TN_parms/ReceiverID* with the external IDs in the profiles. This is how Trading Networks verifies the identities of the sender and receiver. When Trading Networks finds a match, it transforms the external ID in *TN_parms/SenderID* and *TN_parms/ReceiverID* into its internal IDs before storing them in the BizDocEnvelope.

Be sure to select an external ID type that your partners specify in their profiles. You can choose from a list of external ID types. If the external ID type you need is not in the list or if you have agreed on a non-standard format with a trading partner, select the external ID type **Mutually Defined** and add an appropriate ID value.

You can use the **Retrieve SenderID from session** option as an alternative way to set the value of the SenderID system attribute. When this option is set, rather than looking into *TN_parms\SenderID* for the value to use, Trading Networks uses the connected user. This means that when the client sends the document to Trading Networks, it has to present credentials (an Integration Server user name and password). That user name/password pair directly correlates to a Trading Networks partner. When you create a profile, you associate one or more My webMethods user accounts with the profile. For more information on associating My webMethods user accounts, see "[Associate My webMethods User Accounts with a Profile](#)" on page 203. So, if **Retrieve SenderID from session** is set to true in a document type, Trading Networks determines the sender of the document based on the user account you associate with the profile, finds the profile (and therefore the internal ID), and puts that value in the BizDocEnvelope.

Important: When you resubmit a flat file document with the **Retrieve SenderID from session** option selected, the resubmit might fail. If you set up Sender criteria in a processing rule, the first time the document arrives, the criteria matches a processing rule based on the SenderID value that is determined through the **Retrieve SenderID from session** option. However, if you resubmit the document, that same processing rule will not be selected based on the **Retrieve SenderID from session** option. This is because when you resubmit the document from My webMethods, the user account that you used to log in to My webMethods (for example, Administrator), will not match the Sender criteria of the processing rule. It could end up matching another processing rule and then go through completely different processing.

DocumentID

Extract this system attribute if you want to use the document ID to determine whether a document is a duplicate (Check for Duplicate Document pre-processing action), and if you want to search for documents of this document type in My webMethods based on document ID.

UserStatus

Extract this system attribute if you want to specify user status as a criterion for determining the processing rule to use for the document, and if you want to search for documents of this document type in My webMethods based on user status.

GroupID

Extract this system attribute if you want to search for documents of this document type in My webMethods based on group ID.

ConversationID

Extract this attribute if you want to pass the document to a business process. If you do not want to pass the document to BPM, add a variable to the pipeline called `prtIgnoreDocument` with the value of `true`. If *ConversationID* is extracted and *prtIgnoreDocument* variable is set to `true`, then it will not be passed to the process model.

Custom Attributes

Extract any custom attribute you want to use later.

Indicating Whether the Attribute is Required

By default, when Trading Networks cannot extract an attribute, it continues processing, and, if processing completes successfully, sets the processing status to DONE. It does not log an error to the activity log. If you want Trading Networks to log an error to the activity log when it cannot extract an attribute, you can designate that attribute as required for extraction. If processing completes successfully, Trading Networks sets the processing status to DONE W/ERRORS.

Configure the document type's processing rule to handle the error as follows:

- Use the processing rule's recognition error criterion to trigger the processing rule based on whether it contains errors.
- Use the processing rule's Execute a Service processing action to inspect the *bizdoc/Errors* pipeline variable for errors.

Transforming Extracted Attributes

You can have Trading Networks transform extracted attributes before storing them in the Trading Networks database.

Note: You cannot transform the SenderID or ReceiverID attributes.

Built-In Transformations for DATETIME and DATETIME LIST Data Type Attributes

If you extract an attribute that has the data type DATETIME or DATETIME LIST, you must identify the date format. You can choose a built-in common date/time format or you can type the format you need using a pattern string based on the "Time Format Syntax" described for the `java.text.SimpleDateFormat` class. Trading Networks extracts the value of the date and uses it and the pattern you specify to decode the value and convert it to the format that Trading Networks requires to store the date in the `BizDocEnvelope`.

Built-In Transformations for STRING or STRING LIST Data Type Attributes

If you extract an attribute that has the data type STRING or STRINGLIST, you can transform the string value using one of the built-in transformations below.

Format	Description
Uppercase	Transforms the extracted string attribute value to all uppercase.
Lowercase	Transforms the extracted string attribute to all lowercase.
String Substitution	<p>Substitutes extracted values with a pattern you specify. Trading Networks uses the <code>java.text.MessageFormat</code> class to perform this transformation.</p> <p>Suppose your gateway service extracts the names of three items in a purchase order and places them in an attribute named <code>itemName</code> in <code>TN_parms</code> as follows:</p> <pre>/TN_parms/itemName = {"WidgetA", "WidgetB", "WidgetC"}</pre> <p>Your document type would have to include the following:</p> <pre>itemName STRING LIST</pre> <p>Suppose you want to replace the three names with a single string that contains the three names. Use this field to specify the following string substitution pattern:</p> <pre>"Item Name: {0}, {1}, {2}"</pre> <p>Trading Networks would store the following for the value of <code>itemName</code>:</p> <pre>Item Name: WidgetA, WidgetB, WidgetC</pre> <p>If you place more arguments in the pattern than there are extracted values, the string stored in the database for the attribute will contain the extra arguments. If you specify fewer arguments than there are extracted values, the string will contain only the values for the number of arguments.</p>

Built-In Transformations for NUMBER Data Type Attributes

If you extract an attribute that has the data type `NUMBER` and contains an array of numbers, you can use transform the array into a single value using one of the built-in transformations below.

Format	Description
Average	Calculates the average value of all the numbers in the array.
Minimum	Calculates the smallest number in the array.

Format	Description
Maximum	Calculates the largest number in the array.
Sum	Calculates the sum of all the numbers in the array.
No format	Stores the first value of the array as the value of the attribute.

Note: When Trading Networks extracts a NUMBER or NUMBER LIST from a document, it uses the number parsing behavior of `java.lang.Number`. For example, if the NUMBER or NUMBER LIST contains the value `100zzz`, Trading Networks interprets the value as 100, instead of throwing an error as it would if the value were `zzz100`. If you want to redefine this parsing behavior, you can write a custom transformation service.

Suppose your gateway service extracts the quantity of each item of a purchase order, and you want to sum the quantities of all items. Suppose you defined the following attributes in `TN_parms` and in the document type:

<code>TN_parms</code> Attributes	Document Type
<code>/TN_parms/itemNumber</code>	<code>itemNumber NUMBER</code>
<code>/TN_parms/quantity</code>	<code>quantity NUMBER</code>
<code>/TN_parms/totalQuantity</code>	<code>totalQuantity NUMBER</code> <code>summedValue NUMBER LIST</code>

If you select the **Sum** built-in transformation for `totalQuantity`, it will sum its values and produce a single number. For example, if the value of `totalQuantity` is this array:

```
/TN_parms/totalQuantity = {2, 3, 4, 5}
```

The value of `summedValue` will be:

```
/TN_parms/summedValue = 2 + 3 + 4 + 5
```

Trading Networks will store the value 14 for `totalQuantity`.

Custom Transformation Services

To select a custom transformation service, click  and select the service to use. To locate a service, specify the **Package Name** and click **Search** or browse the services in the listed packages. Specify the inputs for the selected service.

Defining Pre-Processing Actions for the Document Type

Click the **Options** tab and define the options and actions.

Enable Processing Rule Routing Option

If you want to process documents using pre-processing and processing actions defined in a processing rule, enable processing rule routing. If you want to process documents using only pre-processing actions defined in the document type, disable processing rule routing.

Note: If you do not want to process documents using a processing rule, disabling routing can improve performance by avoiding the search for processing rules.

Verify Digital Signature Pre-Processing Action

Click  next to the **Verify Digital Signature** field and then select the service to use.

Validate Structure Pre-Processing Action

This action invokes a flat file parser to validate the inbound document's structure and content against a flat file schema, and to convert the document into an IS document you can map to the *TN_parms* pipeline variable. Alternatively, you can perform the same validation and conversion in your document gateway service. Do not convert a document in both the document gateway service and using the Validate Structure action. If you convert a document in the gateway service, it will remain in the pipeline.

Note: Parse the document only if absolutely necessary. Parsing can be time-consuming and return very large outputs that consume much memory.

If you know the fully qualified name of the flat file (parsing) schema to use, type it in the **Parsing schema** field. Otherwise, to locate the flat file schema, click . In the Input for Flat File Validation dialog box, select the parsing settings. If you want to save the parsed contents, select the **keepResults** option.

This action invokes `pub.flatfile:convertToValues` service in the `WmFlatFile` package, which in turn invokes the flat file parser. For information about the service, see the *Flat File Schema Developer's Guide*.

Check for Duplicate Document Pre-Processing Action

Trading Networks provides options to check for a document with the same Document ID; the same Document ID and sender; the same document ID, sender, and receiver; or the same document ID, sender, and document type. In the extraction specifications, you must specify extracting the necessary attributes for the option you select.

If you want to use a custom service that performs the check based on other attributes, click **Select** and search or browse for the service.

Note: Custom duplicate checking services can affect performance, depending on the service logic.

Save Document to Database Pre-Processing Action

You must save a document to the database when you want to:

- Deliver the document using reliable delivery with immediate or scheduled delivery.
- Deliver the document using queue for polling.
- Pass a document to a business process.
- Send a document back to the beginning as a “new” document (for example, because the document did not match any defined document type). This is called resubmitting the document.
- Send a document back through processing rules (for example, because the document was processed by the wrong rule). This is called reprocessing the document.

You can choose to save all documents, no documents, or only unique documents. If you want to save only unique documents, you must also specify the Check for Duplicate Document pre-processing action.

Option	Description
Content	<p>Save document content if you want to:</p> <ul style="list-style-type: none"> ■ View the content in My webMethods when monitoring transactions. ■ Resubmit or reprocess the document. ■ Use the Deliver Document By processing action with reliable delivery.
Attributes	<p>Save custom attributes if you want to view attributes in My webMethods when monitoring transactions, search for documents based on custom attributes, or reprocess documents.</p> <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;"> <p>Important If you reprocess a document without saving custom attributes, you might get unexpected results. Because the attributes are not saved, the document will not match processing rules that use extended criteria. Instead, the document will match another processing rule, such as the default processing rule, and Trading Networks will perform the processing actions defined in that rule.</p> </div>
Activity Log	Record activity log entries when monitoring transactions.

Editing a Flat File Document Type

While editing a flat file document type, if another user tries to edit at the same time, Trading Networks ensures that data written by you is not over written by the other user and vice versa. While saving the document type, Trading Networks ensures that you are working on the latest version of data, and notify you if another transaction has modified the document type. In that case, you have to cancel the changes you have made, and reopen the document type for edit.

To edit a flat file document type

1. My webMethods: **Administration > Integration > B2B > Document Types Administration > Document Types.**
2. Click  **Edit** for the document type to edit.
3. If you want to edit the document type based on a transaction, click  **Show Actions** on the row for the document to use and then select **Edit Document Type**.
4. Edit the fields as described in "[Creating a Flat File Document Type](#)" on page 127.
5. Save the document type.

9 Preparing for Document Delivery

■ Overview	138
■ Creating Custom Immediate Delivery Services	138
■ Enabling Document Delivery Using a Web Service	141
■ Creating Custom Scheduled Delivery Services	143
■ Adding a Public Queue to Trading Networks	152
■ Setting Up the Queue for Polling for webMethods for Partners	154

Overview

Before you specify document deliver in processing rules or partner profiles, you might need to perform the tasks in this chapter.

Creating Custom Immediate Delivery Services

You can create a custom immediate delivery service to do the following:

- Replace services provided with Trading Networks built-in immediate delivery methods.
- Replace services automatically created by Trading Networks when you create custom immediate delivery methods from built-in immediate delivery methods.
- Deliver documents in ways not provided with Trading Networks.

Creating a Custom Immediate Delivery Service

Your immediate delivery service must implement, at a minimum, the input and output variables defined in the `wm.tn.rec:DeliveryServiceSignature` specification. The input value defined by the specification is the document to deliver. The output values in the specification indicate the outcome of the service. Your service must provide values for the `serviceOutput` output variable. The `serviceOutput` variable is an IS document (IData object) that contains these String variables:

For this variable in the <code>serviceOutput..</code>	Specify...
<code>status</code>	Whether the service executed successfully or unsuccessfully. Your service must provide the value <code>success</code> or <code>fail</code> .
<code>statusMessage</code>	Optional. Delivery-specific message about the outcome of the delivery.
<code>transportTime</code>	Optional. Number of milliseconds it took for the service to deliver the document.
<code>output</code>	Optional. Return information from the delivery service (for example, response bytes received from an HTTP post).

Retrieving the Document Content to Deliver

To obtain the document content to deliver, use the `wm.tn.doc:getDeliveryContent` service. You pass this service the `BizDocEnvelope` that contains the document and the service returns the delivery content.

You can use the `wm.tn.doc:getDeliveryContent` service regardless of whether Trading Networks considers the document large or not. This service determines whether to handle the document as large. If the document is not considered large, the service returns the delivery content as a byte array. If the document is considered large, the service returns the delivery content as a Java `InputStream` object.

Note: Trading Networks cannot handle large documents that are delivered using the Web service delivery method.

Registering the New Delivery Service

Execute the `wm.tn.delivery:registerService` built-in service to register your new delivery service. One of the input variables of the `wm.tn.delivery:registerService` built-in service is `serviceName`. The value you specify for `serviceName` becomes the name of the delivery option that My webMethods displays in lists that contain all delivery options. You can execute the `wm.tn.delivery:registerService` service from Software AG Designer, as explained below.

If you later need to change the parameters you specify in this procedure, remove the service using the `wm.tn.delivery:removeService` service and then repeat this procedure.

To register the new delivery service

1. In Software AG Designer, select the `wm.tn.delivery:registerService` service from the Navigation Panel.
2. Select **Test > Run**. Software AG Designer displays this dialog box:

The screenshot shows a dialog box titled "Input for 'registerService'". It contains the following fields and controls:

- serviceName**: Message Queue
- host**: localhost
- port**: 5555
- user**: (empty)
- password**: (empty)
- ifc**: TNCustomize.deliveryServices
- svc**: messageQueue
- scheduled**: false (dropdown menu)
- Include empty values for String Types
- Buttons: OK, Cancel, Load, Save, Help

3. Fill in the input parameters as follows:

Parameter	Entry
serviceName	Unique name to assign to the delivery service and therefore the delivery method (for example, <code>Message Queue</code>).
host	<p>Host name of the Integration Server that hosts Trading Networks. If the delivery service resides on the local machine, specify <code>localhost</code> or leave host blank; do not specify the host name for the local machine.</p> <p>To invoke a delivery service on a remote server, Trading Networks opens an HTTP connection and posts the document to be delivered.</p> <p>Important When you specify a host name, Trading Networks assumes the delivery service is on a remote server. If the delivery service is on the local machine, this consumes resources unnecessarily.</p>
port	Port for the Integration Server that hosts Trading Networks. If you leave host blank, do not specify a value for port.

Parameter	Entry
user	User name of a user account that has the authority to execute the <code>wm.tn.delivery:registerService</code> service. If host is <code>localhost</code> or blank, Trading Networks ignores user .
password	Password for the user account identified in user . If host is <code>localhost</code> or blank, Trading Networks ignores password . Note: The password is securely managed by the Integration Server's Password Manager.
ifc	Fully-qualified folder name of the new delivery service (for example, <code>TNCustomize.deliveryServices</code>).
service	Service name of the new delivery service (for example, <code>messageQueue</code>).
scheduled	Whether the service you are registering is an immediate (click false) or scheduled delivery service (click true).

4. Click **OK**.

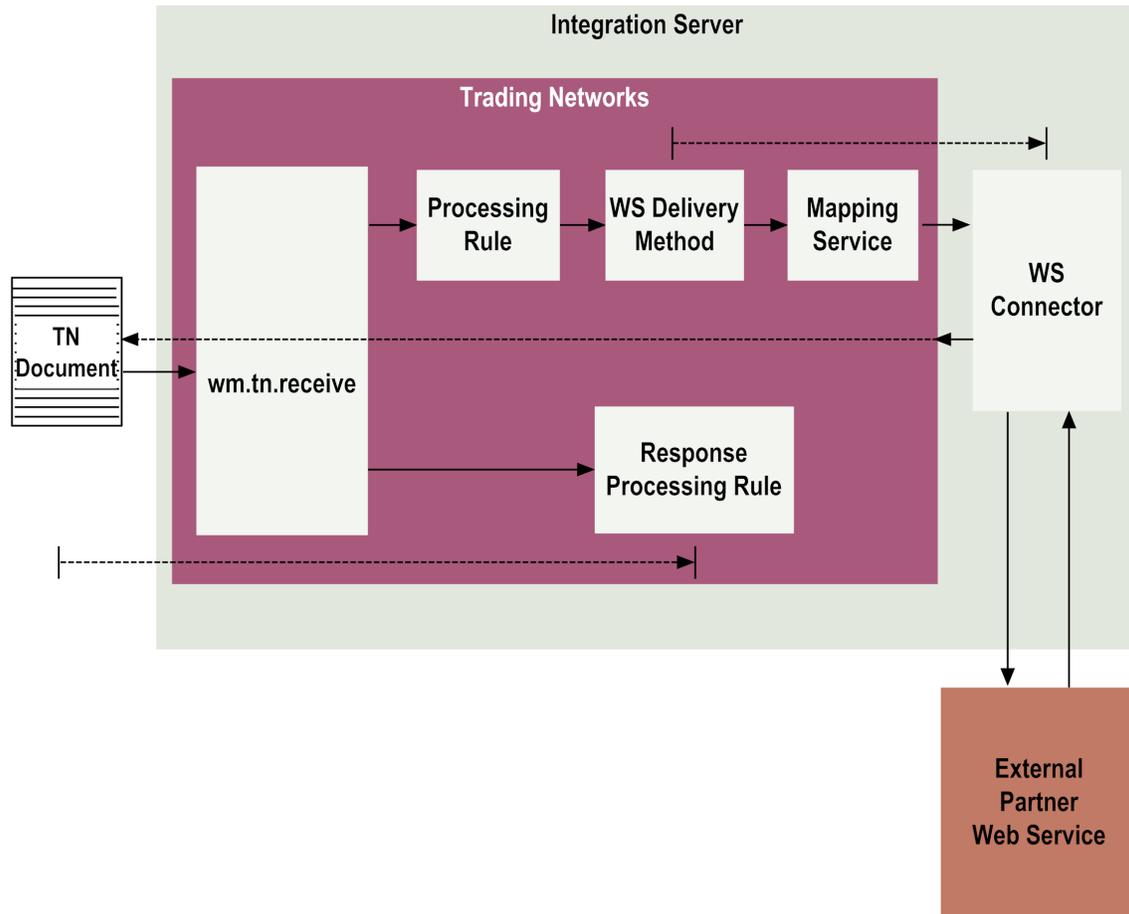
Enabling Document Delivery Using a Web Service

To deliver a document to a partner using a Web service (that is, by invoking a Web service at the partner's end), you create a Web service delivery method. Trading Networks uses the Web services feature provided by Integration Server to deliver documents using a Web service.

Trading Networks uses a Web service connector to invoke the Web service at the partner's end. To pass the data from Trading Networks to the Web service connector based on the connector's input data structure, you create a mapping service. The mapping service maps the pipeline variables to the required input values of the Web service connector. The output of the service must have the same document structure as that of the Web service connector's input document structure. In most cases, the input of the mapping service is the `BizDocEnvelope`, so the mapping maps the `bizdoc` attributes to the input parameters of the Web service connector.

After the Web service connector invokes the Web service at the partner's end, it obtains the response from the partner. The response obtained can be added to the `bizdoc` or sent as a separate document to Trading Networks. If Trading Networks receives the response as a separate document, it processes that document just like any other document it receives.

The following diagram illustrates delivery using a Web service delivery method.



Create a Web service connector for every Web service you want to invoke. Using Software AG Designer, create a consumer Web service descriptor from the remote Web service. For instructions, see the *Web Services Developer's Guide*. Integration Server automatically creates the Web service connector from the descriptor.

Create a mapping service that describes how to pass the data from Trading Networks to the Web service connector. Create a Java service or a flow service whose output has the same document structure as the Web service connector's input document. Map the pipeline variables with the required input values of the Web service connector. In most cases, you might want to map the bizdoc attributes with the relevant input parameters of the Web service connector. In such cases, the input for the mapping service is the BizDocEnvelope.

You will create the Web service delivery method by providing the Web service connector and mapping service details in the receiving partner's profile (see ["Adding an Immediate Delivery Method to a Profile" on page 191](#)). You can add the Web service response to the bizdoc, or you can create a document out of the response. If you want to create a document, define a document type for the document so that Trading Networks recognizes the document, and define a processing rule for that document type.

Creating Custom Scheduled Delivery Services

Creating a Custom Scheduled Delivery Service

A scheduled delivery service must act on a batch of delivery tasks. Your scheduled delivery service must retrieve each delivery task and act on it individually. It is the responsibility of the scheduled delivery service to retrieve each delivery task, deliver the document associated with the delivery task, and update the status of the delivery task.

To perform these functions, your scheduled delivery service might have to take other actions as well. For example, your scheduled delivery service might need to open a connection to a remote machine before it delivers documents and close the connection after it has delivered all the documents. Trading Networks provides a single scheduled delivery service, the `wm.tn.transport:batchFtp` service, which you can use as a reference implementation to write your own scheduled delivery schedule. (The `wm.tn.transport:batchFtp` service opens a connection, delivers all the documents, and then closes the connection.)

Scheduled Delivery Service Inputs

The table below lists the input variables your scheduled delivery service should expect.

Variable	Data Type	Description
<i>queue</i>	String	Name of the queue that contains the delivery tasks on which the service is to act. Your service will use this value as input to the <code>wm.tn.queuing:getQueuedTask</code> built-in service to retrieve a delivery task. For more information, see "Required Logic to Include in the Scheduled Delivery Service" on page 144.
other variables	n/a	Any other input variables your scheduled delivery service requires. These typically depend on the transport protocol your service will use to deliver the documents. For example, for a service that uses FTP, an additional input variable would be the directory in which to place the documents that are being delivered.

Scheduled Delivery Service Outputs

The table below lists the output variables your scheduled delivery service must return.

Variable	Data Type	Description
<i>logMsg</i>	String	Text that describes the result of the scheduled delivery service. This is not the status of acting on a single delivery task in the queue; it is the result of the scheduled delivery service as a whole.

Retrieving the Document Content to Deliver

To obtain the document content to deliver, use the `wm.tn.doc:getDeliveryContent` service. You pass this service the `BizDocEnvelope` that contains the document and the service returns the delivery content.

You can use the `wm.tn.doc:getDeliveryContent` service regardless of whether Trading Networks considers the document large or not. This service determines whether to handle the document as large. If the document is not considered large, the service returns the delivery content as a byte array. If the document is considered large, the service returns the delivery content as a Java `InputStream` object.

Note: Trading Networks cannot handle large documents that are delivered using the Web service delivery method.

Required Logic to Include in the Scheduled Delivery Service

The following table lists actions that you must include in all scheduled delivery services and the built-in service to use to accomplish each action.

Action	Built-In Service to Use
Retrieve a delivery task from a queue	<p><code>wm.tn.queuing:getQueuedTask</code> service.</p> <p>When you invoke this service, you pass it the name of the queue from which you want a delivery task. The service returns the oldest task in the queue; the task includes the document to deliver.</p> <p>When you retrieve a task from the queue, the task is not removed from the queue. Instead, the status of the delivery task is updated from <code>QUEUED</code> to <code>DELIVERING</code>.</p> <p>How the delivery task is represented depends on how your scheduled delivery service is written:</p> <ul style="list-style-type: none"> ■ Flow service: IS document (<code>IData</code> object) that has a structure defined by the <code>wm.tn.rec:Task</code> IS document type.

Action	Built-In Service to Use
	<ul style="list-style-type: none"> ■ Java Service: <code>com.wm.app.tn.delivery.GuaranteedJob</code> object (see the <i>webMethods Trading Networks Java API Reference</i>). <p>Typically, your scheduled delivery service should invoke the <code>wm.tn.queuing:getQueuedTask</code> service and process delivery tasks until the queue is empty (that is, until you have processed each of the QUEUED delivery tasks in the queue). When the queue is empty, the <code>wm.tn.queuing:getQueuedTask</code> service returns null for the <i>task</i> variable.</p> <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;"> <p>Note: The <code>wm.tn.queuing:getQueuedTask</code> service places the <i>timeDequeued</i> variable in the pipeline. Trading Networks uses this variable to calculate the amount of time it takes to transmit the document. Do not update or drop this variable from the pipeline.</p> </div>
Deliver the document	<p>One of the built-in services in the <code>pub.client</code> folder, or one of your own custom transport services.</p> <p>After your scheduled delivery service retrieves a delivery task, it should deliver the document that is returned with the task information.</p>
Update the delivery task	<p><code>wm.tn.queuing:updateQueuedTask</code> service. the service uses the returned <i>status</i> to update the delivery task status:</p> <ul style="list-style-type: none"> ■ If <i>status</i> is <code>success</code>, the service updates the task status to <code>DONE</code>. ■ If <i>status</i> is <code>fail</code>, the service updates the task's retry count. <ul style="list-style-type: none"> ■ If the maximum retry count has not been reached, the service updates the task status to <code>QUEUED</code>. The next time the queue schedule invokes your delivery service, the delivery task will try again to deliver the document. ■ If the maximum retry count has been reached, the service updates the task status to <code>FAILED</code>. <p>The <code>wm.tn.queuing:updateQueuedTask</code> service also takes care of logging the results of the delivery to the activity log.</p>

Typical Logic to Include in the Scheduled Delivery Service

A typical scheduled delivery service might deliver all the batched documents to a single destination. The `wm.tn.transport:batchFtp` built-in scheduled delivery service behaves in this manner and is a good sample to refer to when creating your own scheduled delivery

service. For more information, see "[The `wm.tn.transport:batchFtp` Built-in Service](#)" on page 147.

The following shows steps you might want to include in your scheduled delivery service to deliver all documents to the same destination.

```
// Retrieve a delivery task from the queue:

If the queue is empty
  Exit the service

Open a connection to the remote server

While queue is not empty
  // Try to deliver the document associated with the delivery task
  // Update the task with the results of the delivery (e.g., success or fail)
  // Retrieve the next delivery task from the queue

Close the connection to the remote server
```

The actions in *italics* are items that all scheduled delivery services should perform.

The above shows only one way to create the service. You might have different requirements for your service. For example, you might want to combine several documents and transmit them as one large document or might want to limit the number of documents sent to a destination in a single invocation of your service.

As shown above, the scheduled delivery service might need to make a connection to a remote server. The typical approach is to open the connection after you have determined that the queue has delivery tasks in it (that is, the queue is not empty). You can determine whether the queue is empty by invoking the `wm.tn.queuing:getQueuedTask` built-in service. If this service returns null, the queue is empty. After your scheduled delivery service has acted on the delivery tasks in the queue, your service can close the connection to the remote server.

Some transport protocols require additional actions after opening a connection and before you begin delivering documents. For example, for the FTP protocol, after performing the `LOGIN` to connect to the remote server, a `cd` (change directory) is performed to navigate to the appropriate directory.

Handling Exceptions

When the queue schedule dictates, the `wm.tn.queuing:deliverBatch` service is invoked. The `wm.tn.queuing:deliverBatch` built-in service, in turn, invokes your scheduled delivery service and passes your service information about the queue.

This `wm.tn.queuing:deliverBatch` service catches any exceptions thrown by your scheduled delivery service or the underlying transport service that your services uses to deliver the document. In response to an exception, the `wm.tn.queuing:deliverBatch` service increments the retry count for the current delivery task and logs the exception to the activity log.

If your scheduled delivery service needs to take action in response to an exception, use the `pub.flow:getLastError` service to detect whether an exception was thrown.

If you are creating your scheduled delivery service using the flow language, and you want to exit your scheduled delivery service due to a failure, use the following flow operation:

```
EXIT $flow and signal FAILURE
```

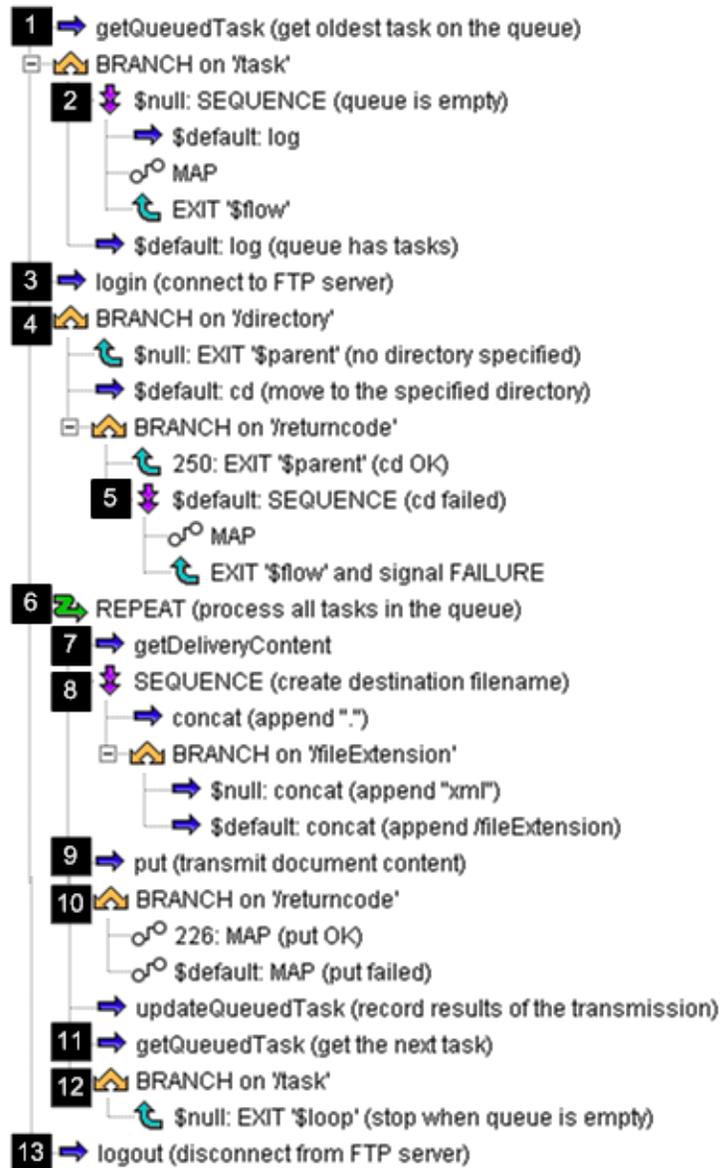
The above flow operation causes the `wm.tn.queuing:deliverBatch` service to increment the current delivery task's retry count and log the exception to the activity log. You provide the information that Trading Networks logs to the activity log. To provide this information, be sure you specify meaningful information in:

- The failure-message property of the EXIT flow operation.
- The `logMsg` variable in Pipeline out; map meaningful text to this variable.

The `wm.tn.transport:batchFtp` Built-in Service

Trading Networks provides one built-in scheduled delivery service, `wm.tn.transport:batchFtp`, which is shown below. This service delivers a batch of documents using the FTP protocol to a single destination. It is written in the flow language. Use this service as a model if you are creating your own custom scheduled delivery service.

The wm.tn.transport:batchFtp Service



Flow Operations	Description
1	The INVOKE flow operation invokes the <code>wm.tn.queuing:getQueuedTask</code> built-in service to retrieve the first delivery task from the queue. The service returns the task information in the <code>task</code> variable.
2	If the variable, <code>task</code> , is null, the queue is empty. In this case, map text to the output variable, <code>logMsg</code> , and exit the service.

Flow Operations	Description
3	If the queue contains delivery tasks, attempt to log in to the remote FTP server.
4	An input to the <code>wm.tn.transport:batchFtp</code> service is the variable, <i>directory</i> . If this variable was specified, <code>cd</code> to the specified directory.
5	If the <code>cd</code> to the specified directory fails, map a message to the output variable, <i>logMsg</i> , and exit the service.
6	The REPEAT flow operation causes the service to loop over the delivery tasks in the queue.
7	The INVOKE flow operation invokes the <code>wm.tn.doc:getDeliveryContent</code> built-in service to retrieve the document content to be delivered.
8	<p>The operations in the SEQUENCE flow operation form the file name to use for the document being delivered. The file name will be <i>internalID.ext</i>, where <i>internalID</i> is the Trading Networks generated internal ID for the document and <i>ext</i> is the file extension.</p> <p>The first INVOKE flow operation in the sequence invokes the <code>pub.string:concat</code> service to append a period to the internal ID. The BRANCH flow operation determines the file extension. An input to the <code>wm.tn.transport:batchFtp</code> service is the variable, <i>fileExtension</i>. If this variable was not specified, the file extension defaults to “xml”. Otherwise, the file extension is the value specified for the <i>fileExtension</i> variable.</p>
9	This INVOKE flow operation invokes the <code>pub.client.ftp:put</code> service that transmits the file using FTP.
10	<p>If the return code from the <code>pub.string:concat</code> service is 226, the file was transmitted successfully. In this case, the MAP flow operation sets the value of the <i>status</i> variable in Pipeline Out to “success.”</p> <p>For any other return code the other MAP flow operation sets the value of the <i>status</i> variable in Pipeline Out to “fail.”</p>
11	The first INVOKE flow operation invokes the <code>wm.tn.queuing:updateQueuedTask</code> built-in service to update the status of the delivery task. One of the inputs to the service is <i>status</i> that was set to either “success” or “fail”.

Flow Operations	Description
	The second INVOKE flow operation invokes the <code>wm.tn.queuing:getQueuedTask</code> service to retrieve the next delivery task from the delivery queue.
12	If the <code>wm.tn.queuing:getQueuedTask</code> service returned a delivery task, continue in the REPEAT loop. If the service returned null, the queue is empty; exit the REPEAT loop.
13	After looping through all delivery tasks, log out of the remote FTP server.

The `wm.tn.transport:batchFtp` service does not do any exception handling. All transport-level exceptions cause the current invocation of the `wm.tn.transport:batchFtp` service to terminate, and the exception will be handled by its caller, the `wm.tn.queuing:deliverBatch` built-in service. The `wm.tn.transport:batchFtp` service will be invoked again for the delivery queue, according to the queue's schedule.

Registering the New Delivery Service

Execute the `wm.tn.delivery:registerService` built-in service to register your new delivery service. One of the input variables of the `wm.tn.delivery:registerService` built-in service is `serviceName`. The value you specify for `serviceName` becomes the name of the delivery option that My webMethods displays in lists that contain all delivery options. You can execute the `wm.tn.delivery:registerService` service from Software AG Designer, as explained below.

If you later need to change the parameters you specify in this procedure, remove the service using the `wm.tn.delivery:removeService` service and then repeat this procedure.

To register the new delivery service

1. In Software AG Designer, select the `wm.tn.delivery:registerService` service from the Navigation Panel.
2. Select **Test > Run**. Software AG Designer displays this dialog box:

3. Fill in the input parameters as follows:

Parameter	Entry
serviceName	Unique name to assign to the delivery service and therefore the delivery method (for example, <code>Message Queue</code>).
host	<p>Host name of the Integration Server that hosts Trading Networks. If the delivery service resides on the local machine, specify <code>localhost</code> or leave host blank; do not specify the host name for the local machine.</p> <p>To invoke a delivery service on a remote server, Trading Networks opens an HTTP connection and posts the document to be delivered.</p> <p>Important When you specify a host name, Trading Networks assumes the delivery service is on a remote server. If the delivery service is on the local machine, this consumes resources unnecessarily.</p>
port	Port for the Integration Server that hosts Trading Networks. If you leave host blank, do not specify a value for port.

Parameter	Entry
user	User name of a user account that has the authority to execute the <code>wm.tn.delivery:registerService</code> service. If host is <code>localhost</code> or blank, Trading Networks ignores user .
password	Password for the user account identified in user . If host is <code>localhost</code> or blank, Trading Networks ignores password . Note: The password is securely managed by the Integration Server's Password Manager.
ifc	Fully-qualified folder name of the new delivery service (for example, <code>TNCustomize.deliveryServices</code>).
service	Service name of the new delivery service (for example, <code>messageQueue</code>).
scheduled	Whether the service you are registering is an immediate (click false) or scheduled delivery service (click true).

4. Click **OK**.

Adding a Public Queue to Trading Networks

To add a public queue to Trading Networks

1. In My webMethods: **Administration > Integration > B2B > Public Queues**.
2. Click **Add Queue**.
3. In the **Name** field, specify the name for the queue.
4. Do the following:
 - a. In the **Delivery Service** field, select the scheduled delivery service to use with the queue, then click the **Inputs** tab and provide the inputs to the scheduled delivery service.

If you want the input parameters to include fields that you leave blank, select **Include empty values for string types?**. If you select this option and you leave a field blank, Trading Networks passes an empty string for the value. If you do not select this option and you leave a field blank, the field is not included in the inputs.
 - b. From the **State** list, assign a state to the queue, as follows:

State	Meaning
Enable	Place delivery tasks in the queue and deliver documents.
Disable	Do not place new delivery tasks in the queue and do not deliver documents. All delivery tasks that are already in the queue remain in the queue. You might disable the queue if you need to fix the scheduled delivery service, or if a processing rule has been placed on the wrong queue and you need to change the rule or reassign tasks to another queue. If the partner profile is disabled, the private queue is disabled.
Drain	Do not place new delivery tasks in the queue but deliver documents in the queue. You might drain a queue if you want to delete it, or perform maintenance changes (for example, change queue settings, update delivery services).
Suspend Delivery	Place new delivery tasks in the queue but do not deliver documents. You might suspend a queue if a trading partner is temporarily unable to accept documents.

You can also change the status on the **Public Queues** page. For this, select the checkbox beside the queue, and select a status from **Change Status**. You can do this to change the status of multiple queues at once.

Note: When the state of a queue is **Disable** or **Drain**, delivery fails because Trading Networks cannot place the delivery task in the queue. Trading Networks sets the delivery task status to FAILED and logs a message to the activity log.

- c. Click the **Schedule** tab. In the **Process Queue** list, select how often to invoke the scheduled delivery service to deliver the documents in the queue, then fill in the fields below.

Important: Only use My webMethods to change the delivery schedule and other queue settings for a scheduled delivery service. Do not try to use Integration Server Administrator to change the schedule, or the data passed to the service will be lost and the service will not be able to deliver documents.

Selection	Action
Run Once	<p>Invokes the scheduled delivery service one time only, at the date and time you specify. This option is typically used in a development environment.</p> <p>Note: If the scheduled date and time occurs when the host Integration Server is not running, Integration Server throws an exception when you next start Integration Server.</p>
Fixed Interval	<p>Invokes the scheduled delivery service at the fixed interval you specify (for example, every 3 hours). The scheduled delivery service executes for the first time immediately after Trading Networks adds the delivery task.</p> <p>If you want to wait for the scheduled delivery service to finish running before starting the next scheduled invocation of the service, select the Do not overlap task check box.</p> <p>Note: When you add or update a fixed interval schedule for a queue whose state is Enabled or Drained, Trading Networks invokes the service immediately. Subsequent invocations are governed by the fixed interval.</p>
Hourly	Invokes the scheduled delivery service at the dates and times you specify.
Daily	
Weekly	Use the Date Range fields to specify the first and last days to invoke the service. Use the Process Queue fields to specify the exact times to invoke the service.
Monthly	
Yearly	If you want to wait for the scheduled delivery service to finish running before starting the next scheduled invocation of the service, select the Do not overlap task check box.

- Click **Save & Close**.

Setting Up the Queue for Polling for webMethods for Partners

This section describes how to set up the queue for polling for webMethods for Partners.

To set up the queue for polling for webMethods for Partners

1. For setting up queues for polling, Trading Networks uses license files. Make sure the partner Integration Server that hosts webMethods for Partners uses the partner license key, and the hub Integration Server that hosts Trading Networks uses the hub license key.
2. Go to the `partner_Integration_Server_directory\instances\instance_name\config` directory and open the `server.cnf` file.
3. Make sure the `watt.server.partner` property is set to the host name for the hub Integration Server. For example, if the host name for the hub Integration Server is `localhost`, the property would be `watt.server.partner=localhost`.

If the `watt.server.partner` property is missing or incorrect, add or update the property. To do so, shut down the partner Integration Server, add or update the property, and restart the partner Integration Server.

Note: If the `watt.server.partner` property is missing, the partner Integration Server cannot request the documents queued for it on the hub Integration Server. If the property is set incorrectly, the partner Integration Server will throw an exception when it tries to request queued documents.

10

Creating Processing Rules

■ Overview	158
■ Before You Begin	158
■ Creating a Processing Rule	160
■ Defining a Processing Rule Based on a Saved Document	169
■ Editing the Default Processing Rule	170
■ Reordering Processing Rules	170
■ Testing the Order of Processing Rules	171
■ Deleting a Processing Rule	171

Overview

If you are working in a clustered environment, all Trading Networks servers must be set up for cluster synchronization. For instructions, see "[Configuring Trading Networks for a Clustered Environment](#)" on page 82.

The order in which you list your processing rules on the Processing Rules page is important. If a document type enables processing rule routing, and does not identify a specific rule to use to process documents, Trading Networks matches each inbound document against processing rules in the order the rules are listed on that page, and uses the first processing rule that matches the document. Order rules so that rules with specific criteria are above rules with more general criteria. Also, set up a default processing rule for Trading Networks to use when a document does not match any of the other processing rules.

Before You Begin

Creating a Service for the Check for Duplicate Document Action

In the Check for Duplicate Document pre-processing action, Trading Networks offers options to check based on Document ID, sender, receiver, or document type. If you want to check based on other attributes, you can use a custom service. Create the service using the `wm.tn.rec:DupCheckService` specification.

Note: Custom duplicate checking services can affect performance, depending on the service logic.

Creating a Service for the Execute a Service Action

In the Execute a Service processing action, you can invoke a custom service synchronously, asynchronously, or using a service execution task.

If you invoke the service synchronously or asynchronously, the service must at minimum implement the input variables defined in the `wm.tn.rec:ProcessingService` specification. The input values are the IS documents (IData objects) that are in the pipeline during processing. The specification does not identify any outputs.

If you invoke the service asynchronously using a service execution task, the service must at minimum implement the input and output variables defined in the `wm.tn.rec:ReliableProcessingService` specification. The input values are the IS documents (IData objects) that are in the pipeline during processing. The service must provide values for the `serviceOutput` output variable, which is an IS document (IData object) that contains these String variables:

String Variable	Value
<i>status</i>	Whether the service: <ul style="list-style-type: none"> ■ Executed successfully. When the service returns <code>success</code>, Trading Networks updates the status of the service execution task to DONE. ■ Failed. When the service returns <code>fail</code> and the maximum retry limit has been met, Trading Networks updates the status of the service execution task to FAILED. If the maximum retry limit has not been reached, the status of the service execution task remains PENDING and Trading Networks tries to re-execute the service after the wait period elapses.
<i>statusMessage</i>	Optional. Message that contains more information about the outcome of the service. If the service fails, for example, you should provide an error message that explains the failure.
<i>transportTime</i>	Optional. Number of milliseconds it took the service to execute.

If the service requires input values in addition to those in the specification, you might want to hardcode some of the values when you set up the processing rule. For example, you might create a service that can add, update, or delete an entry for a partner in your back-end system. Input to the service is a flag that indicates the action to take. You could create three separate processing rules that all execute this service, but set the input values for the flag differently for each rule. For the rule that adds an entry, set the input value to indicate that the service should add an entry. For the rule that updates an entry, set the input value to indicate the service should update an entry, and so on.

Using Alert Email Messages

If you want to use the Alert e-Mail Message processing action, you must configure Trading Networks to send email messages. You do this by specifying the address of the Simple Mail Transfer Protocol (SMTP) server through which Integration Server is to issue email messages.

To specify the address of the SMTP server

1. In Integration Server Administrator: **Settings > Resources > Edit Resource Settings**.
2. Under **Email Notification**, set the **SMTP Server** field to the domain name or IP address of the SMTP server to use.
3. If you want to send the email to the webMethods administrator, specify an email address in the **Internal Email** field.
4. Save your changes.

Performing Tasks for the Deliver Document By Action

If you want to use the Delivery Document By action, see ["Preparing for Document Delivery" on page 137](#) for tasks you might need to perform.

Creating a Processing Rule

Creating and Naming the Rule

To create and name the rule

1. In My webMethods: **Administration > Integration > B2B > Processing Rules**.
2. Do one of the following:
 - To create the rule by providing new information, select the row near which to place the new rule, click **Add Processing Rule**, and then click **Above**, **Below**, or **Last**. If you use **Last**, reposition the Default rule so it remains in the last position in the list.

Note: Trading Networks adds the new rule to the full, unfiltered list that shows all processing rules. If you are working with a filtered list of processing rules, adding a new processing rule might affect the rules that are not displayed on your query results table. Ask a Trading Networks administrator to check the order in the full, unfiltered list, or test the rule order to make sure it is working as you intend (see ["Testing the Order of Processing Rules" on page 171](#)).

- To create the rule by copying an existing rule and modifying it, click  in the row for the rule to copy.
3. On the Processing Rule Details page, provide a name and description for the rule, and indicate whether to enable or disable the rule. When a rule is disabled, Trading Networks does not match documents to it or use it to process documents. Do not enable the rule until you are done creating it.

Specifying Criteria for the Rule

Click the **Criteria** tab and fill in the fields to identify the documents to process using this rule.

Specifying Sender and Receiver Criteria

You can specify sender, receiver, or both as criteria for matching inbound documents to the processing rule. For **Sender**, select one of the following:

Select...	To use this rule for documents...
Any	Regardless of sender.
Enterprise	That identify your corporation as sender.
Unknown	For which the sender is an unknown partner.
Selected	That identify specified partner profiles and partner groups as sender. Click  Select Partner(s) and select partners, or click  Select Partner Group(s) and select groups.

For **Receiver**, select from the same options, but for receiver.

Specifying Document Type Criteria

You can specify document type as a criterion for matching inbound documents to the processing rule. For example, you could specify that documents that match the cXML Purchase Order document type should use this processing rule.

Select...	To use this rule for documents...
Any	Regardless of document type.
Unknown	That are an unknown document type.
Selected	That match specified document types. Click  Select Document Types and select document types.

Specifying User Status Criteria

You can specify user status as a criterion for matching inbound documents to the processing rule. You can specify in the document type to extract the UserStatus system attribute or, if you do not, the initial value for this attribute will be null.

Select...	To use this rule for documents...
Any	Regardless of user status.
Specified	That have specified user statuses. To add a user status to the table, click  and specify the user status value. To remove a user status, click  .

Specifying Recognition Error Criteria

You can specify the occurrence or non-occurrence of errors during document recognition or attribute extraction as a criterion for matching inbound documents to the processing rule.

Select...	To use this rule for documents...
May have errors	Regardless of whether the document contains errors.
Has no errors	That do not contain errors.
Has errors	That contain errors.

Setting Up Extended Criteria for the Rule

You can specify extracted attributes as criteria for matching inbound documents to the processing rule. Click the **Extended Criteria** tab to provide the details.

Attribute Field

Specify the name of a custom attribute to match on.

Operator Field

Indicate how to match the **Value** (below) against the value extracted from the documents. For example, you can match string values that contain specified characters, match number values that are greater than a specified value, or match date values to a specified date. The operator you can specify depends on the data type of the attribute.

Value Field

Value specifies the value to match against the value extracted from the documents. You can also specify null. Trading Networks might have a null value for an attribute for these reasons:

- For XML documents, the XQL query in the document type did not find the specified attribute in the document.
- The document type of the document is unknown, so Trading Networks did not extract attributes.

If you use the Is Null operator, you do not need to specify a value in this field.

You can specify the following:

To specify...	Do this...
A string	Type the string. This match is case-sensitive.
A number	Type the number.
A date and time	Click  to select the date from a calendar or type the date and time using the format <i>yyyy-mm-dd hh:mm:ss</i> and a 24-hour clock for <i>hh</i> .

Defining Pre-Processing Actions for the Rule

For each pre-processing action you specify, you indicate that you want to:

- Perform the corresponding pre-processing action specified in the document type.
- Perform the pre-processing action in the processing rule.
- Not perform the pre-processing action at all.

Click the **Pre-Processing** tab and define the actions.

Verify Digital Signature Pre-Processing Action

To use this action for XML documents, you must extract the SignedBody and Signature system attributes.

Select...	To...
Defer to document type	Use the Verify Digital Signature pre-processing action in the document type.
Verify digital signature	Always verify digital signatures or digitally sign documents that are processed with this rule. The services used for signing the documents are available in the webMethods Trading Networks <i>Built-In Services Reference Guide</i> .
Do not verify digital signature	Never verify digital signatures or digitally sign documents that are processed with this rule.

Validate Structure Pre-Processing Action

Use this action to invoke a flat file parser to validate the document's structure and content against a specified schema.

Select...	To...
Defer to document type	Use the Validate Structure pre-processing action in the document type.
Validate structure	Always validate the structure of documents that are processed with this rule by picking up the schema.
Validate structure	Never validate the structure of documents that are processed with this rule.

Check for Duplicate Document Pre-Processing Action

Select...	To...
Defer to document type	Use the Check for Duplicate Document pre-processing action in the document type.
Use built-in services to check for duplicate document	Check documents using an option provided by Trading Networks. You can check the database for a document with the same document ID; the same document ID and sender; the same document ID, sender, and receiver; or the same document ID, sender, and document type. To use an option, you must extract the corresponding system attributes.
Use custom services to check for duplicate document	Check documents against a custom duplicate checking service. Click  Select Service and browse the IS namespace for the service. To use the service, you must extract the attributes specified in the service.
Do not check for duplicate document	Never check the uniqueness of documents processed with this rule.

Save Document to Database Pre-Processing Action

You must save a document to the database when you want to:

- Deliver the document using reliable delivery with immediate or scheduled delivery.
- Deliver the document using queue for polling.
- Pass a document to a business process.

- Send a document back to the beginning as a “new” document (for example, because the document did not match any defined document type). This is called resubmitting the document.
- Send a document back through processing rules (for example, because the document was processed by the wrong rule). This is called reprocessing the document.

Select...	To...
Defer to document type	Use the Save Document to Database pre-processing action in the document type.
Save	Always save the document content, attributes, and activity log. Indicate whether to save information for all documents or only for documents that are unique. To only save documents that are unique, you must also check for uniqueness using the Check for Duplicate Documents action.
Do not save	Never save the document content, attributes, or activity log.

Defining Processing Actions for the Rule

Click the **Action** tab and then click **Add Action** to add a processing rule action. Complete your selections for each processing rule action in the dialog box displayed.

Execute a Service Action

Note: If you have not already done so, perform the instructions in ["Creating a Service for the Execute a Service Action"](#) on page 158.

Click  **Select Service** and select the service to invoke for the action. Specify how to invoke the service, as follows:

Selection	Description
Synchronous	Synchronously execute the service once.
Asynchronous	Asynchronously execute the service once.
Service execution task	Trading Networks uses reliable execution to re-execute the service one or more times if the service fails.

To set inputs for the service, click **Set Inputs** and type values for the parameters for which to hardcode an input value.

Alert e-Mail Message Action

Note: If you have not already done so, perform the instructions in ["Using Alert Email Messages" on page 159](#).

Email To Field

To send the email to...	Select...
webMethods administrator	webMethods Administrator.
A sender contact	Sender , and then select the contact type from the list. The contact type must be defined in the sender's profile.
A receiver contact	Receiver , and then select the contact type from the list. The contact type must be defined in the receiver's profile.
A specific address	Email address , and then specify the address.

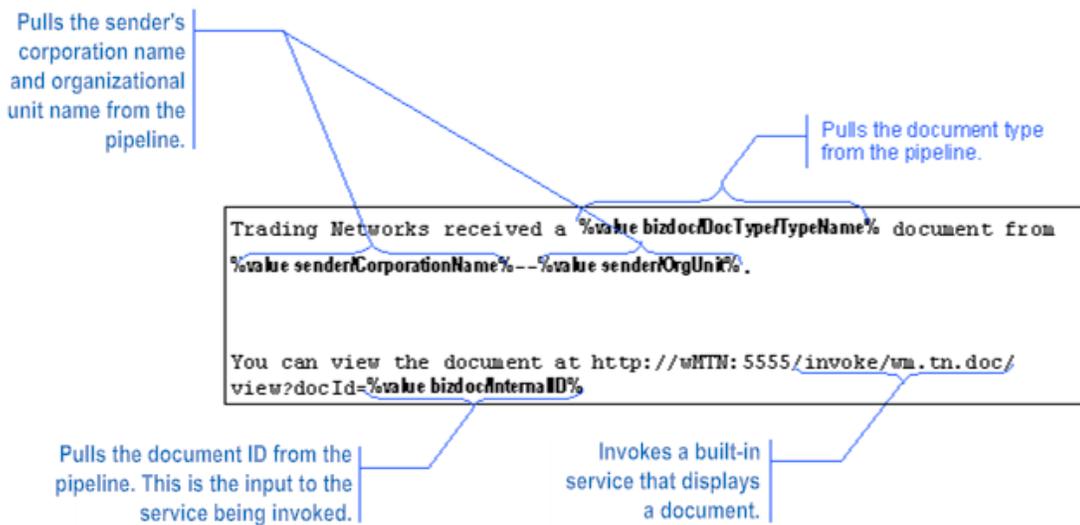
Subject Field

Type a character string for the subject of your email.

Mail Body Field

Type a character string for the body of your email.

At this point, the pipeline contains the information described in ["Creating a Service for the Execute a Service Action" on page 158](#). In addition, if the processing rule specifies the Execute a Service action and the service is executed synchronously, the pipeline also contains any information that the service placed in the pipeline. If you want to make the body of the email message dynamic, you can include information from the pipeline in the body using output template tags. For example, you might want the message to specify the type of document that was received and the sender of the document. You might also want to include a hyperlink that allows the recipient of the email message to view the document. The following shows the text you might specify for the body. (For illustrative purposes, bold text is used for the output template tags.)



If the document type is Purchase Order, the sender's corporation name is XYZ Steel Company, and the organizational unit is Alloys Division, the above sample message body would be rendered as follows:

```
Trading Networks received a Purchase Order document from XYZ Steel Company--
Alloys Division.
You can view the document at
http://wMTN:5555/invoke/wm.tn.doc/view?docId=000233AB907ED1234
```

For information about output template tags, see the *Dynamic Server Pages and Output Templates Developer's Guide*.

Change User Status Action

Use this action to assign a user status to the document. In the **Change To** field, specify a character string (for example, Needs Approval).

Note: You can also change user status using the Java methods `com.wm.app.tn.doc.BizDocEnvelope.setUserStatus(String)` and `com.wm.app.tn.db.BizDocStore.changeStatus(BizDocEnvelope,String,String)` in the Trading Networks Java API. For information, see the *webMethods Trading Networks Built-In Services Reference* or the built-in service `wm.tn.doc:changeStatus`. These APIs are given for creating custom solutions.

Deliver Document By Action

To use this action, you must extract the ReceiverID system attribute from the document.

Note: If you have not already done so, perform the instructions in "[Preparing for Document Delivery](#)" on page 137.

Important: If you use the Execute a Service action and the executed service updates the receiver identified in the document, Trading Networks does not deliver the document to the partner that the service specifies as the receiver unless that service updates the receiver in the bizdoc in the pipeline. If the service does not update the receiver, Trading Networks delivers the document to the receiver identified when it initially received the document for processing.

Select this field...	To...
Immediate delivery	Use immediate delivery. Select the immediate delivery method from the list. If you want to use reliable delivery, specify the Save Document to Database pre-processing action.
Scheduled delivery	Use scheduled delivery. Click the scheduled delivery queue to use in the list. The private queue for the receiving partner is called Receiver's queue in the list. If you want to use reliable delivery, specify the Save Document to Database pre-processing action.
Queue for polling	Place the document in the queue for polling.
Receiver's preferred protocol	Use the immediate delivery method that is identified in the receiving partner's profile as the partner's preferred protocol. The preferred protocol option is primarily for use with webMethods eStandards Modules.

Respond Action

Field	Entry
Content Type	Indicate the content type of the message. For example, if the message is an XML document, specify the content type, <code>text/xml</code> . If the message is in plain text, specify the content type <code>text/plain</code> .
Message	Specify a character string. To use a larger input window to specify the message, click Expand . If you want to make the body dynamic, you can include information from the pipeline. For more information, see " Alert e-Mail Message Action " on page 166 .

Defining a Processing Rule Based on a Saved Document

You can have Trading Networks create a processing rule based on a document that you saved in the Trading Networks database. When Trading Networks creates the processing rule, it sets the criteria, pre-processing actions, and actions as described below. You can edit the rule and add extended criteria.

Setting	Description
Criteria	<ul style="list-style-type: none"> ■ Sender criterion: Set to the sender identified within a document or unknown. ■ Receiver criterion: Set to the receiver identified within a document or unknown. ■ Document type criterion: Set to the document type for the document or unknown. ■ User status criterion: Set to the value of the user status that is associated with the document, if any.
Pre-Processing	<ul style="list-style-type: none"> ■ Verify digital signature: Always verify the document's digital signature. ■ Validate structure: Always validate the document's structure. ■ Check for duplicate document: Use the action in the document type. ■ Save document to database: Always save the document content, attributes, and activity log.
Action	Trading Networks sets the processing action to ignore the document (that is, to not perform any processing actions on the document).

To create a processing rule based on a saved document

1. In My webMethods: **Monitoring > Integration > B2B > Transactions**.
2. Click  **Show Actions** on the row for the document to use to create the processing rule.
3. Select **Create Processing Rule**.
4. You can edit the criteria, pre-processing actions, or processing actions, and you can add extended criteria. For instructions, see "[Creating a Processing Rule](#)" on page 160.

5. Save the processing rule.

Note: Trading Networks saves the new rule at the top of the list of processing rules. If you want to move the rule to another position in the list, see ["Reordering Processing Rules" on page 170](#).

Editing the Default Processing Rule

Trading Networks has a default processing rule named Default Rule. The criteria in this processing rule matches a document that has any sender, any receiver, any document type, and any user status, and might have errors. The default rule does not specify any extended criteria. Any document that Trading Networks receives will meet the criteria of the default processing rule.

The pre-processing actions in the default processing rule indicate that Trading Networks should use the settings in the document type. By default, the Change User Status processing action changes the user status of the document to Ignored. However, you can edit the pre-processing and processing actions in the default processing rule. Do not change the criteria.

The default processing rule must always be the last rule in the full, unfiltered processing list.

To edit the default processing rule

1. In My webMethods: **Administration > Integration > B2B > Processing Rules**.
2. Click **Default rule**. The default processing rule is listed last.
3. To edit the pre-processing actions, click the **Pre-Processing** tab and use the instructions in ["Defining Pre-Processing Actions for the Rule" on page 163](#).
4. To edit the processing actions, click the **Action** tab and see ["Defining Processing Actions for the Rule" on page 165](#).
5. Save the processing rule.

Reordering Processing Rules

The default processing rule must always be the last rule in the full, unfiltered processing list.

To reorder processing rules

1. In My webMethods: **Administration > Integration > B2B > Processing Rules**.
2. Click  for the rule to move.

3. Move the rule up or down, or select a location for the rule. For example, if you want to move rule 4 to position 8, select the rule and specify position 8 in the Select Location dialog box.

Testing the Order of Processing Rules

You can test the order of your processing rules to make sure Trading Networks selects the correct processing rules for your documents.

To test the processing rules, you must have a sample document in your file system. Trading Networks performs these actions on the document:

- Recognizes the document type.
- Extracts attributes as specified in the document type.
- Determines all enabled processing rules the document matches.

Trading Networks does not perform any pre-processing or processing actions on the document.

To test the order of processing rules

1. In My webMethods: **Administration > Integration > B2B > Processing Rules**.
2. Click the **Advanced** tab in the Search panel.
3. Click the **Test** tab.
4. Click **Browse** to select the sample document to test and then click **Test**.

Trading Networks lists the processing rules that match the sample document.

Deleting a Processing Rule

If you no longer need a processing rule, you can delete it in My webMethods. Deleting a processing rule removes it entirely from Trading Networks. You will not be able to recover it.

11 Onboarding New Partners

■ Overview	174
■ Uploading Partner Information	175
■ Managing Partner Onboarding Templates	176
■ Managing Invitation Emails	180
■ Granting Partners Access to the Questionnaire Page	182
■ Approving or Rejecting a Partner	183
■ Viewing and Changing the Status of a Partner	183

Overview

The partner onboarding process can help you automate the creation of partner profiles and trading partner agreements (TPAs) using the partner details provided in a comma-separated values (.csv) file or by way of questionnaire responses received from the partner. For more information about this process, see ["Overview of the Partner Onboarding Process" on page 33](#).

You can also create partner profiles manually by entering partner details on various My webMethods pages. For information about creating profiles in this way, or editing profiles created using either method, see ["Creating Profiles" on page 185](#).

Summary of the Partner Onboarding Process

Following is a high-level overview of the partner onboarding process:

1. The administrator uploads a spreadsheet into My webMethods that contains basic information about one or more partners. For more information on different ways to upload partner information, see ["Uploading Partner Information" on page 175](#).
2. In My webMethods, Trading Networks sets the status of the partner to New (Pending Partners page) and the status of the TPA to Proposed (Trading Partner Agreements page). For more information on partner status, see ["Viewing and Changing the Status of a Partner" on page 183](#).
3. If additional details are needed from the partner (for example, if the spreadsheet contained only the partner name and email address), the administrator does the following:
 - a. Creates a questionnaire template that identifies the additional information needed. For more information on the questionnaire templates, see ["Managing Partner Onboarding Templates" on page 176](#).
 - b. Configures settings for an invitation to be emailed to the partner. For more information on configuring emails, see ["Configuring Invitation Emails" on page 180](#).
 - c. Sends the invitation email to the partner requesting him or her to log in to My webMethods to complete the questionnaire. For more information on sending invitation emails, see ["Sending Invitation Emails" on page 181](#).
 - d. Grants the partner access to the questionnaire page in My webMethods. For more information, see ["Granting Partners Access to the Questionnaire Page" on page 182](#).
4. Trading Networks sets the partner's status to Invited.
5. The partner completes the questionnaire. Trading Networks sets the partner's status to Responded.

Note: If the partner does not complete the questionnaire within a configured time limit, Trading Networks sets the partner's status to Expired. The administrator can restart the onboarding process for that partner by re-inviting the partner.

6. The administrator checks the details that the partner has submitted. If they are complete and in order, the administrator approves the partner. If the details are not in order, the administrator resends the questionnaire for the partner to enter again.
7. If the partner enters the questionnaire correctly this time, the administrator approves the partner. If not, the administrator rejects the partner and Trading Networks sets the partner's status to Rejected. For more information, see "[Approving or Rejecting a Partner](#)" on page 183.
8. If the partner is approved, Trading Networks adds the information from the questionnaire to the partner's profile and TPA. The status of the partner is changed from Pending to Approved and the status of the TPA from Proposed to Agreed.

Uploading Partner Information

The administrator starts the onboarding process for one or more partners by uploading a spreadsheet that contains basic information about the partners. At a minimum, the spreadsheet can contain the partner's name and email address. The administrator can download the sample template provided, enter and save the information, and upload the file.

Alternatively, the administrator can start the onboarding process for a single partner by using the Quick Add feature.

Uploading Partner Information from a Spreadsheet

To upload partner information from a spreadsheet

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Import Information**.
2. On the **Upload From File** tab of the Upload Partners page, click **Browse**.
3. Locate and select the spreadsheet file (in .csv format) containing the partner information, and then click **Open**.

Note: You can also download the sample template provided on the Import Information page, enter and save the information, and upload the file.

4. Click **Upload Partners**.

Note: The spreadsheet must comply with the header information provided in the import template of Import Information page. It must contain, at a minimum, the partner name and email address. If other information

about the partner is supplied in the spreadsheet, the spreadsheet must also contain the partner's external ID.

5. Click the down arrow next to **Upload Errors** and **Successful Uploads** to view the results in the **Upload Results** area. Correct any errors that may have occurred.

Note: You can click the **Status** icon to view messages related to the upload in the **Messages** tab. Messages include upload confirmation and error messages. If an error message appears, you can click the **StackTraces** tab to see the stack trace report for the error.

If successful, Trading Networks sets the status of the partner to **New**.

Adding Partner Information Using the Quick Add Feature

To quickly add a partner to the onboarding process

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Import Information**.
2. On the **Quick Add** tab of the Upload page, enter the partner's name and email address, and then click **Add Partner**.

If successful, Trading Networks sets the status of the partner to **New**.

Managing Partner Onboarding Templates

You can manage templates that are used for the invitation emails. You can use templates to seek information from potential partners by adding them to the questionnaire that you attach to the invitation. Three templates are provided by default:

- Non-EDI Standard Template: used as a standard template for non-EDI questionnaires.
- EDI Standard Template: used as a standard template for EDI questionnaires.
- EDIINT Standard Template: used as a standard template for EDIINT questionnaires.

wm.tn.util:mapStandardFields is the default mapping service that is used to map the standard fields to Trading Networks. If the default template is changed or you create a new template, you will need to create a new mapping service for the edited or new template. For more information about creating a mapping service for new templates, see ["Creating a Mapping Service for New Templates" on page 179](#).

Searching for a Template by Name

To search for the template by name

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Manage Templates**.
2. Enter the template name in the **Keywords** text box, and click **Search**. The search results are displayed.

Adding a New Template

Perform these steps to add a new template that you can attach to an invitation sent to a partner.

To add a new template

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Manage Templates**.
2. Click **Add Template**.
3. Enter the template name and description.
4. To add a new group to the template, click **Manage Groups**. Four default groups (Corporation, HTTP, Contact, and Address) are displayed in the **Groups** section.

Tip: You can reorder the groups in any template by clicking the arrow button and selecting **Move Up**, **Move Down**, or **Select Location**.

5. Add the groups that you want to add to the template by selecting the check box(es) and clicking **OK**.
6. Click **Save & Close**.

Creating a New Group from an Existing Group

To create a new group from an existing group

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Manage Templates**.
2. Click a template link to view the template details.
3. Click **Manage Groups**.
4. Click the **Copy** icon in the **Manage Group(s)** dialog box.
The **Copy Group** dialog box appears, appending the group name with "Copy of".
5. In the **Group Details** section, change the name and description for the group.

6. In the **Group Fields** section, select the appropriate check boxes.
7. Click **OK**.

Adding New Fields to an Existing Group

To add new fields to an existing group

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Manage Templates**.
2. On the Templates page, click **Add Field**.
3. Enter the field name and description.

Note: If you added a group name and description in the **Add New Group** section, the fields are added to this new group.

4. In the **Group** list, select the group to which this new field should belong.
5. If you want the new field to be a mandatory field, select the **Required** check box.
6. Specify whether the data type is **String** or **Binary**.

Note: If you select the data type as **Binary**, you will not be able to set any default or valid values.

7. Specify the **Maximum Length** of the field.
8. Specify the **Default Value** of the field.

Note: Ensure that the default value, if specified, is one of the valid values, if valid values are specified.

9. To add valid values for the field, click the **+** button in **Valid Values** and enter the value in the text box.

Tip: You can add as many valid values as needed by clicking the **+** button and entering the value.

10. Click **OK**.

Managing Fields in a Group

To manage fields within a group

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Manage Templates**.
2. Open the template.

3. Click  in the header of the group whose fields you want to manage.
4. On the Manage Fields dialog box, select whether to display **One field per line** or **Two fields per line**.
5. To add a field to the group, click **Add Field** and specify the field's details. For more information, see "[Adding New Fields to an Existing Group](#)" on page 178.
6. To reorder fields in the group, click  and change the order.
7. Select whether to show or hide a field for the group on the Manage Template page by selecting or clearing the check box next to the field.
8. Click **OK** to save the settings.

Creating a Mapping Service for New Templates

If you make changes to any of the partner onboarding templates, you need to edit the mapping service. If you create a new partner onboarding template, you need to create a new mapping service. If the administrator approves the template that you edited or created, the edited or new mapping service pushes the approved information to the profile/TPA.

Note: If you create a new mapping service, you need to enter the mapping service name in the **Mapping Service** box on the **Manage Templates** page.

To create mapping fields for the edited or new templates

1. In Software AG Designer, create a new flow service. Ensure that the structure of the pipeline input (Pipeline In) to the mapping service is as follows: **Template > Group Name > Field Name**.

Note: **Template** is a fixed variable name, whereas **Group Name** and **Field Name** should reflect the values you defined in the template.

2. Add groups and fields to **Pipeline In** in the flow service.
3. Map the **Pipeline In** fields to the **Pipeline Out** fields. There are three **Pipeline Out** groups, each having a list of fields:
 - **Profile.** For details, see the description for the `wm.tn.rec:Profile` service in *webMethods Trading Networks Built-In Services Reference*.
 - **EDIINT.** For details, see the section on the Extended Fields tab of the profile in *webMethods Module for EDI Installation and User's Guide*.
 - **EDITPA.** For details, see the section on defining EDI trading partner agreements in *webMethods Module for EDI Installation and User's Guide*.

For details about editing or creating a mapping service, see the section on mapping data in flow services in *webMethods Service Development Help*.

Managing Invitation Emails

Managing partner onboarding invitation emails consists of configuring the emails, sending the emails to prospective partners, editing the settings for emails that have not yet been sent, and deleting invitations that have expired.

Configuring Invitation Emails

You can configure the settings that Trading Networks uses to send and track invitation emails.

To configure invitation emails

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Manage Invitations**.
2. In the **SMTP Server** field under **General** in the **Email Preferences** section, enter the SMTP server details that you are using to send email. If you do not enter the SMTP server details, Trading Networks uses the Integration Server SMTP server.
3. In the **Email** field, type the email address of the Trading Networks administrator who will be overseeing partner onboarding activities. This email will be used to send invitation emails to the potential partners.
4. In the Invitations area, do one of the following:
 - a. To create a new invitation from scratch, click **Add**.
 - b. To create an invitation based on an existing one, click the copy icon next to the invitation you want to start with.
5. In the **Name** field, type a name for the invitation email (for example, "Supplier Invitation"). You can use this name to search for specific invitation emails later. If you clicked the copy icon to create the invitation email, a new invitation email is created with "Copy of" prefixed to the name in the **Name** field. You can edit it to enter a new name for the invitation email.
6. In the **Description** field, type a description for the invitation email (for example, "Partner onboarding invitation for XYZ suppliers in Europe").
7. On the Email Invite tab, do the following:
 - a. In the **To** field, click **Select Partners** to select the partners to be invited. The From field is auto-populated with the email address specified in the tn.mail.from property in the Trading Networks Configuration Properties page (**B2B Settings > Edit Properties > TN Configuration Properties**).
 - b. In the **Subject** field, enter the text to appear on the subject line of the email.
 - c. In the **Message Body** field, enter the text to appear in the body of the email.

- d. Click the **Questionnaire** text box and select the questionnaire template you would like to include with the invitation email. When a partner responds to the invitation email by logging in, this template determines the questions that the partner will see when prompted for additional profile details.
- e. Click **Save**.

Sending Invitation Emails

To send the invitation emails

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Manage Invitations**.
2. Click the **Settings** tab and do the following:
 - a. To send the invitation email immediately, select the **Send Invitation Immediately** check box.
 - b. To send the invitation email at another time, specify the send date and time.
 - c. To send a reminder schedule for the invitation email, click **Set Reminder**.
 - d. Select **Daily**, **Weekly**, or **Monthly** from the **Run Schedule** list.
 - e. Select the hours that the invitation should be sent in the **Run schedule at these times** section.
 - f. Click **Save**.
 - g. When you are ready to send the invitation email according to the schedule you just configured, select **Active** from the **Status** list at the top of the page.

Note: Inactive invitations are saved as drafts.

- h. Click **Save & Close**.

Managing Partner Invitation Emails

The information that the partner must provide is defined in the questionnaire template that you specify for that partner. While sending an invitation to a partner, you can associate only one active questionnaire template with the partner. The administrator can edit the settings for invitation emails that have not yet been sent and delete an invitation that has expired.

To manage invitation emails

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Manage Invitations**.
2. In the **Search** area of the **Invitations** page, do one of the following:

- To search for an invitation email by name, click the **Keyword** tab, type the name of the invitation, and click **Search**.
 - To specify detailed search criteria based on invitation expiration, start, or end date, invitation name, partners identified as recipients, invitation email status, or questionnaire template associated with the invitation, click the **Advanced** tab. Add the desired filters and click **Search**.
3. To delete an invitation email that has not already been sent or has not yet expired, select the email in the **Invitations** panel, click **Delete**, and then click **OK** to confirm.
 4. To edit an invitation email, click the **Edit** icon in the **Actions** column next to the email. Edit the email's settings as described in "[Configuring Invitation Emails](#)" on page 180.
 5. To set the status of an invitation email, do one of the following:
 - Double click on the invitation and choose **Active** or **Inactive** from **Status**.
 - Select the checkbox beside the invite in the **Invitations** panel and setting the status to **Active** or **Inactive**. You can do this to set the status of multiple invitations at once.

If the invitation is **Active** and you change it to **Inactive**, the entire schedule will be suspended till the invitation is made active again.

Note: In **Active** mode, you cannot change any of the field values other than reminders and expiration date.

Granting Partners Access to the Questionnaire Page

To grant a partner access to the questionnaire page in My webMethods, you grant permissions to the TN Onboarding Users role to access the questionnaire page for a specific Trading Networks server instance.

Note: Trading Networks automatically adds My webMethods users who receive the questionnaire as a part of the onboarding process to the TN Onboarding Users role.

To grant a partner access to the questionnaire page

1. In My webMethods: **Administration > My webMethods > System Settings > TN Servers**.
2. Click the **Permissions** icon for a Trading Networks Server instance for which you want to grant access.
3. In the Select Roles dialog box, type the name of the role you want to add and click **Search**. In this case, it is the TN Onboarding Users role.
4. Move the role from the **Available** list to the **Selected** list.

5. Click **Apply** and then click **Save**.

Approving or Rejecting a Partner

After the administrator sends the partner an invitation, the partner logs in to My webMethods using the login URL and credentials available in the invitation. On the Welcome screen, the partner is asked to change the password and enter details in the questionnaire. Once completed, the partner clicks **Submit** to submit the details to the administrator.

After the partner submits the questionnaire, you can approve or reject the partner.

To approve, reject, or ask partner to resend information

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Review Responses**.
2. Click the **Detail** icon next to the partner response to view the details that the partner entered.
3. Click **Approve**, **Reject**, or **Resend** as appropriate.
 - If approved, Trading Networks adds the information from the questionnaire to the partner's profile.
 - If rejected, the partner is rejected. The partner needs to apply again in order to be an approved partner.
 - If the partner is asked to resubmit the information, the partner needs to edit the information and submit again. If approved by the administrator, the partner will be approved.

Viewing and Changing the Status of a Partner

Trading Networks tracks the status of partners who are participating in the partner onboarding process, starting from the time the partner's information is uploaded and ending when the partner either becomes certified or is rejected as a potential partner. You can change the partner's status to reflect the partner's progress through the onboarding process.

To view and change a partner's onboarding status

1. In My webMethods: **Administration > Integration > B2B > Partner Onboarding > Manage Status**.
2. Select the partner whose status you want to change.
3. Click **Change Status**, and then select the appropriate status as follows:

Status	Description
New	The partner's initial information is uploaded.
Invited	An email was sent to the partner asking the partner to complete the partner questionnaire.
Responded	The partner completed the partner questionnaire and is ready to submit test documents.
Expired	The partner did not complete the partner questionnaire within the specified time limit.
Rejected	The partner was rejected and cannot transact on Trading Networks.
ReInvited	The partner was reinvited because he did not respond to the invitation before it expired or the administrator re-sent the invitation to partner.
Approved	The administrator approved the partner and the partner can now start transactions with Trading Networks. The partner is shown in Partner Profiles page.

12 Creating Profiles

■ Overview	186
■ Before You Begin	186
■ Creating a Profile	187
■ Adding Extended Fields to a Profile	209
■ Associating Partners with Partner Groups	210
■ Working with TPAs	211
■ Deleting a Partner Profile	215
■ Finding Partner Certificates that Are Expired or Expiring Soon	215
■ Managing Partner Access to APIs	216

Overview

If you are using Trading Networks, you can add as many partner profiles as you want. If you are using webMethods for Partners, you can add one partner profile.

Profiles might take awhile to create. Software AG recommends that you save your updates frequently.

The procedures described in this chapter create a partner profile from scratch. You can automate some of this process by uploading partner details from a comma-separated values (.csv) file or by way of questionnaire responses received from the partner. For details, see ["Onboarding New Partners" on page 173](#).

Before You Begin

Creating Contact Types

You can specify contacts in profiles. If you need to create a contact type, use the `wm.tn.dictionary:addContactType` built-in service. Contact Types are defined for creating multiple contacts for partners by using this service. It will start showing in the **Partner Profile > Contact Profile > Add Contact > Contact Type** drop-down list.

Creating External ID Types

You specify external IDs in profiles. If you need to create an external ID type, you can use the `wm.tn.dictionary:addIDType` built-in service or you can use My webMethods as described in ["Adding External IDs to a Profile" on page 188](#).

Creating Services for Use with TPAs

You can create services for use with TPAs, as follows:

Service	Description
Initialization	<p>Populates the inputs to the variables in the IS document type with default values.</p> <p>Some webMethods products that take advantage of the TPA feature supply an initialization service. For example, the webMethods Module for EDI supplies the <code>wm.b2b.editn.TPA:initService</code> service, while the webMethods ebXML Module supplies the <code>wm.ip.ebxml.cpa:initTPA</code> initialization service.</p>

Service	Description
	<p>Note: If you create a partner-defined EDITPA from scratch, do not use the <code>wm.b2b.edi.tn.TPA:initService</code> initialization service because it populates the partner-defined EDITPA with hard-coded values other than those in the default EDITPA, which is already tailored to a user's requirements. Instead, manually complete each parameter field in the Input for <code>wm.b2b.edi.tn.TPA:EDITPA</code>.</p>
Validation	<p>Validates the data added to the IS document for the TPA. The validation service must adhere to the service specification <code>wm.tn.rec:TPAValidationService</code>.</p> <p>Note: Specifying a validation service improves processing time because it eliminates the need to validate the TPA every time a process that uses the TPA runs on Integration Server.</p>
Export	<p>Exports a TPA and converts it to an industry-standard format. After creating a TPA, if you need to obtain the industry-standard format of a TPA, you can execute your export service.</p>

Creating a Profile

Creating and Naming a Profile

You must create your own (Enterprise) profile before you can add partner profiles.

Note: While editing a profile, if another user tries to edit at the same time, Trading Networks ensures that data written by you is not over written by the other user and vice versa. While saving the profile, Trading Networks ensures that you are working on the latest version of data, and notify you if another transaction has modified the profile. In that case, you have to cancel the changes you have made, and reopen the profile for edit.

To create and name a profile

1. In My webMethods: **Administration > Integration > B2B > Partner Administration > Partner Profiles**.
2. Do one of the following:
 - To create your own profile, click **Create My Enterprise**.

- To create a partner profile, do one of the following:
 - i. To create the profile by providing new information, click **Add Profile**.
 - ii. To create the profile by copying an existing profile and modifying it, click  in the row for the profile to copy.
3. Fill in these fields:

Field	Entry
Corporation Name	In your profile, the name of your corporation. In a partner profile, the name of the partner's corporation.
Organization Unit	Optional. Name of the organizational unit or division within the corporation (for example, Sales and Marketing).
Status	<p>Whether to enable or disable the profile. When a profile is disabled, the profile owner cannot exchange documents with partners. Do not enable the profile until you are done creating it. You must supply information for all required fields before you can enable the profile.</p> <p>You can also enable or disable a profile on the Partner Profiles page by selecting the checkbox beside the profile and setting the status to Active or Inactive respectively. You can do this to change the status of multiple profiles at once.</p>
Preferred Language	The locale to use.

Adding External IDs to a Profile

To add an external ID to a profile

1. On the Partner Profiles page, select the profile to which you want to add external IDs.
2. Click the **External IDs** tab.
3. If you need to create an external ID type, click **Add ID Type**, provide the name, and click **OK**.
4. Click **Add ID** and fill in these fields:

Field	Entry
ID Type	External ID type to use.
ID Value	Alpha-numeric external ID value. For example, if you selected DUNS as the external ID type and your D-U-N-S number is 987654321, specify 987654321.

Adding Addresses to a Profile

To add addresses to a profile

1. On the Partner Profiles page, select the profile for which you want to add the address.
2. Click the **Addresses** tab.
3. Click **Add Address....**
4. In the Add Address window, you can do either of the following.
 - Manually fill in details for all the fields.
 - Copy or share information from an already existing address.
5. Click **OK**.

Copying or Sharing Addresses

To copy or share information from an already existing address

1. In the Add Address window, click **Select...** next to Select Address.
2. You can either select an address from the list of addresses displayed or search for a particular address using a keyword.

To search for an address, type the keyword and click **Search**. The search is performed across Address Line1, City, State, and Country fields. The addresses that match the criteria is displayed and you can select the required one.

3. Do one of the following:
 - Click **Copy Address**. This copies information from the selected address to the new address. You can modify the fields before saving the new address.
 - Click **Share Address**. This shares information existing in the selected address with the new address. You cannot modify the fields before saving the new address because the same information is shared among various partner profiles.

Note: To modify an address that is shared among different partner profiles, search for the address in the **Addresses** tab and click  to modify. Any

changes made to the address will affect all the partners who share the information.

Adding Contacts to a Profile

To add contacts to a profile

1. On the Partner Profiles page, select the profile for which you want to add the contact.
2. Click the **Contacts** tab.
3. Click **Add Contact...**
4. In the Add Contact window, you can do either of the following.
 - Manually fill in details for all the fields.
 - Copy or share information from an already existing contact.
5. Click **OK**.

Copying or Sharing Contacts

To copy or share information from an already existing contact

1. In the Add Contact window, click **Select...** next to Select Contacts.
2. You can either select a contact from the list of contacts displayed or search for a particular contact using a keyword.

To search for a contact, type the keyword and click **Search**. The search is performed across Role, Last Name, First Name, E-mail, Country, Address, and City fields. The contacts that match the criteria is displayed and you can select the required one.

3. Do one of the following:
 - Click **Copy Contact**. This copies information from the selected contact to the new contact. You can modify the fields before saving the new contact.
 - Click **Share Contact**. This shares information existing in the selected contact with the new contact. You cannot modify the fields before saving the new contact because the same information is shared among various partners.

Note: To modify a contact that is shared among different partner profiles, search for the contact in the **Contacts** tab and click  to modify. Any changes made to the contact will affect all the partners who share the information.

Specifying Document Delivery in a Profile

Trading Networks allows the flexibility of sending documents to partners using immediate delivery, scheduled delivery, or document queues. For more information, see "[Document Delivery](#)" on page 36.

In order to deliver documents in the preferred method, you must add the corresponding delivery method to the partner profiles:

- ["Adding an Immediate Delivery Method to a Profile" on page 191](#)
- ["Adding Scheduled Delivery Method to a Partner Profile" on page 196](#)
- ["Specifying Reliable Delivery Settings in a Partner Profile" on page 198](#)
- ["Specifying Queue for Polling Settings in a Partner Profile" on page 199](#)
- ["Suspending or Resuming Document Delivery for a Partner" on page 200](#)

Adding an Immediate Delivery Method to a Profile

To add an immediate delivery method to a profile

1. On the Partner Profiles page, select the profile to which you want to add an immediate delivery method.
2. Click the **Delivery Settings** tab.
3. Click **Add Delivery Method**.
4. In the **Delivery Method** drop-down list, select the immediate delivery method to add.
5. If you select one of the standard methods such as FTP, FTPS, HTTP, HTTPS, SFTP, SMTP, or Web service in order to create an immediate delivery method, the **Delivery Method Name** field appears, in addition to the fields described in the steps below. Trading Networks automatically prefixes the value you type in this field with the immediate delivery method to form the complete name for the method (for example, HTTP - XYZSteel). Trading Networks also automatically creates the corresponding immediate delivery service with this name.
 - a. For an FTP, FTPS, HTTP, SFTP, and HTTPS immediate delivery method, specify the following delivery details manually:

You can also copy or reuse information from an existing delivery method. For more information, see ["Copying or Sharing Immediate Delivery Methods" on page 195](#)

Field	Description
Host	In your profile, the host name of your system (for example, XYZSteel.com). In a partner profile, the host name of the partner's system.
Port	Port on which the specified host listens for incoming requests.

Field	Description
	<p>Note: If you do not supply a port number, Trading Networks uses the defaults FTP-21, FTPS-21, HTTP-80, or HTTPS-443.</p>
<p>Location</p>	<p>Where to deliver documents using this delivery method.</p> <p>For FTP or FTPS, type the directory path to which to deliver documents on the system identified by Host and Port. The file name of the file is the internal ID Trading Networks generates for the document when the document arrives. For XML documents, the file extension is .xml. For flat file documents, the file extension is dat.</p> <p>For HTTP or HTTPS, type the URL to use (for example, cgi-bin/acceptOrder). Trading Networks automatically fills in the Location with /invoke/wm.tn/receive, the standard URL for document exchange.</p>
<p>User Name</p>	<p>In a partner profile, the user name Trading Networks should supply to connect to the partner's system.</p> <p>The user name can be for an Integration Server or My webMethods user account that has partner authority (that is, is a member of the Trading Networks ACL named TNPartners).</p>
<p>Password</p>	<p>In a partner profile, the password Trading Networks should supply to connect to the partner's system.</p> <p>Note: Passwords used in scheduled delivery queues (public and private) are stored in the Trading Networks database in binary-encoded form, not in clear text. Since trading partners can create scheduled delivery services, Trading Networks cannot determine which user-defined input variable might be a password, and therefore cannot encrypt passwords used in scheduled delivery queues.</p>
<p>Protect Data Channel</p>	<p>For FTPS,</p>

For an E-mail or SMTP immediate delivery method, fill in this field:

Field	Description
E-mail	<p>In your profile, the e-mail address to which you want documents delivered (for example, bizdoc@isteel.com). In a partner profile, the e-mail address to which the partner wants documents delivered.</p> <p>If you want to deliver the documents to multiple e-mail addresses, separate the e-mail addresses using a comma.</p>

For a Web service immediate delivery method, fill in these fields:

Field	Description
Response Handling	<p>Whether the response received after Web service execution should be:</p> <ul style="list-style-type: none"> ■ Added to the bizdoc. Select Additional Content Part. ■ Converted into a new document. Select New Document.
Web Service Connector	<p>Fully qualified name of the Web service connector to use to deliver the document.</p> <p>Click Select. In the Select Web Service Connector dialog box, search for the Web service connector by providing the name of the Integration Server package that contains the Web service connector and the name of the Web service connector namespace. You can provide a string or the entire namespace name. You can use wildcard characters. Then click Search and select the Web service connector from the list.</p>
Mapping Service	<p>Fully qualified name of the mapping service that provides the input to the Web service connector.</p> <p>Click Select. In the Select Mapping IS Service dialog box, follow the instructions for Web Service Connector, above.</p>

Field	Description
Output Parameter	<p>In a partner profile, the output parameter of the Web service connector that contains the required response.</p> <p>Click Select and select the required parameter in the Key field.</p>
Security Options	<p>Certificates to use to secure communication between Trading Networks and the Web service. Select one or more of the options listed below.</p> <ul style="list-style-type: none"> ■ Secure: Use Secure Sockets Layer (SSL) while exchanging documents. ■ Sign: Use the private key required to sign the documents being exchanged. ■ Verify: Verify the digital signatures of the documents being exchanged. ■ Encrypt: Encrypt documents being delivered. ■ Decrypt: Decrypt documents received.

For an SFTP immediate delivery method, fill in these fields:

Field	Description
SFTP User Alias	<p>Specify the SFTP user alias name that you defined on the Integration Server that hosts Trading Networks. You must ensure to use the same SFTP user alias as defined in Integration Server, for Trading Networks to successfully deliver documents to the SFTP server location. You can either make use of an existing SFTP user alias or create a new one on the host Integration Server.</p> <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p>Note: If Trading Networks is in a clustered setup, you must also ensure that the SFTP server alias and SFTP user alias with the same name and configuration exists in every Integration Server node in the cluster to avoid misconfiguration.</p> </div> <p>For instructions to create SFTP server and user aliases, refer to the <i>webMethods Integration Server Online Help</i>.</p>

Field	Description
Location	Type the directory path to which to deliver documents on the system identified by the SFTP server alias.

6. If you select ActiveTransfer as the delivery method:

Note: If ActiveTransfer is installed on a local Integration Server instance, then ActiveTransfer is listed as a delivery method in the **Delivery Method** drop-down list by default. If ActiveTransfer is installed on a remote Integration Server instance, then you need to configure the remote server alias for ActiveTransfer. For more information, see "[Configuring Alias for ActiveTransfer on Remote Server](#)" on page 84.

- a. Click **Select** to select an ActiveTransfer endpoint (Virtual Folder System or VFS) which is the destination of the document delivery.
- b. In the **Select VFS** window, use either of the following methods to find and select the appropriate VFS for a new partner or the VFS associated with an existing partner:
 - Search for the required VFS.
 - Select **Show All Virtual Folders** to view a list of all partners and the virtual folders associated with them.
- c. Select the VFS.
- d. Click **Select** and then **OK**.

Note: If the immediate delivery method is the preferred method to use to deliver documents to the partner, select the **Use as preferred protocol** check box. If you do not designate an immediate delivery method as the preferred protocol, Trading Networks uses the queue for polling by default. The preferred protocol option is primarily for use with **webMethods eStandards Modules**.

Copying or Sharing Immediate Delivery Methods

To copy or share information from an already existing immediate delivery method

1. In the **Add Delivery Method** window, click **Select...**
2. You can either select a delivery method from the **Delivery Method** drop-down list or search for a particular delivery method using a keyword.

To search for a delivery method, type the keyword and click **Search**. The search is performed across the delivery method name, host, and location fields. The delivery methods that match the criteria are displayed and you can select the required one.

3. Do one of the following:

- Click **Copy Delivery Method**. This copies information from the selected delivery method to the new delivery method. You can modify the fields before saving the new delivery method.
- Click **Share Delivery Method**. This shares information existing in the selected delivery method with the new delivery method. You cannot modify the fields before saving the new delivery method because the same information is shared among various partners.

Note: To modify a delivery method that is shared among different partner profiles, search for the delivery method in the **Delivery Settings** tab and click  to modify. Any changes made to the delivery method will affect all the partners who share the information.

Adding Scheduled Delivery Method to a Partner Profile

You add scheduled delivery method to a partner profile by specifying the queue to use for the partner. A queue is associated with a scheduled delivery service and a delivery schedule.

To add a scheduled delivery method to a profile

1. On the Partner Profiles page, select the profile to which you want to add a scheduled delivery method.
2. Click the **Delivery Settings** tab.
3. Click **Add Delivery Method**.
4. In the **Delivery Method** list, click **Queue for Delivery**.
5. Select a public queue or the partner's private queue.
6. To define the partner's private queue, click **Edit** and do the following:
 - a. In the **Delivery Service** field, select the scheduled delivery service to use with the queue, then click the **Inputs** tab and provide the inputs to the scheduled delivery service.

If you want the input parameters to include fields that you leave blank, select **Include empty values for string types?**. If you select this option and you leave a field blank, Trading Networks passes an empty string for the value. If you do not select this option and you leave a field blank, the field is not included in the inputs.

- b. From the **State** list, assign a state to the queue, as follows:

State	Meaning
Enable	Place delivery tasks in the queue and deliver documents.

State	Meaning
Disable	Do not place new delivery tasks in the queue and do not deliver documents. All delivery tasks that are already in the queue remain in the queue. You might disable the queue if you need to fix the scheduled delivery service, or if a processing rule has been placed on the wrong queue and you need to change the rule or reassign tasks to another queue. If the partner profile is disabled, the private queue is disabled.
Drain	Do not place new delivery tasks in the queue but deliver documents in the queue. You might drain a queue if you want to delete it, or perform maintenance changes (for example, change queue settings, update delivery services).
Suspend Delivery	Place new delivery tasks in the queue but do not deliver documents. You might suspend a queue if a trading partner is temporarily unable to accept documents.

You can also change the status on the **Public Queues** page. For this, select the checkbox beside the queue, and select a status from **Change Status**. You can do this to change the status of multiple queues at once.

Note: When the state of a queue is **Disable** or **Drain**, delivery fails because Trading Networks cannot place the delivery task in the queue. Trading Networks sets the delivery task status to FAILED and logs a message to the activity log.

- c. Click the **Schedule** tab. In the **Process Queue** list, select how often to invoke the scheduled delivery service to deliver the documents in the queue, then fill in the fields below.

Important: Only use My webMethods to change the delivery schedule and other queue settings for a scheduled delivery service. Do not try to use Integration Server Administrator to change the schedule, or the data passed to the service will be lost and the service will not be able to deliver documents.

Selection	Action
Run Once	Invokes the scheduled delivery service one time only, at the date and time you specify. This option is typically used in a development environment.

Note: If the scheduled date and time occurs when the host Integration Server is not running, Integration Server

Selection	Action
	throws an exception when you next start Integration Server.
Fixed Interval	<p>Invokes the scheduled delivery service at the fixed interval you specify (for example, every 3 hours). The scheduled delivery service executes for the first time immediately after Trading Networks adds the delivery task.</p> <p>If you want to wait for the scheduled delivery service to finish running before starting the next scheduled invocation of the service, select the Do not overlap task check box.</p>
	<p>Note: When you add or update a fixed interval schedule for a queue whose state is Enabled or Drained, Trading Networks invokes the service immediately. Subsequent invocations are governed by the fixed interval.</p>
Hourly	Invokes the scheduled delivery service at the dates and times you specify.
Daily	
Weekly	Use the Date Range fields to specify the first and last days to invoke the service. Use the Process Queue fields to specify the exact times to invoke the service.
Monthly	
Yearly	If you want to wait for the scheduled delivery service to finish running before starting the next scheduled invocation of the service, select the Do not overlap task check box.

Specifying Reliable Delivery Settings in a Partner Profile

If a partner is using immediate or scheduled delivery, and is using reliable delivery, you must specify reliable delivery settings in the partner's profile.

To specify reliable delivery settings in a partner profile

1. On the Partner Profiles page, select the profile for which you want to specify reliable delivery settings.
2. Click the **Delivery Settings** tab.
3. Fill in these fields:

Field	Entry
Delivery Maximum Retries	Number of times to try to redeliver a document to the partner. Specify an integer greater than 0.
Wait Between Retries	For immediate delivery, the number of milliseconds to wait before trying to redeliver the document the first time. Specify an integer greater than 0.
Retry Factor	<p>For immediate delivery, the factor to use when determining how long to wait before making subsequent attempts to redeliver the document. Trading Networks calculates the time to wait by multiplying the last wait time by Retry Factor. Specify an integer greater than zero.</p> <p>For example, the following shows sample values for the reliable delivery settings and how they affect the wait times between attempts to redeliver a document.</p> <p>Delivery Maximum Retries = 3</p> <p>Wait Between Retries = 10000 ms</p> <p>Retry Factor = 2</p> <p>Wait until next attempt:</p> <ul style="list-style-type: none"> ■ Retry 1: 10000 ms ■ Retry 2: 20000 ms ■ Retry 3: 40000 ms

Note: You do not set the **Wait Between Retries** and **Retry Factor** fields for scheduled delivery because Trading Networks retries the delivery based on the delivery schedule.

Specifying Queue for Polling Settings in a Partner Profile

Trading Networks provides and manages the queue for polling. This queue is always enabled; you cannot disable, suspend, or drain it.

To specify queue for polling settings in a partner profile

1. On the Partner Profiles page, select the profile for which you want to specify the queue for polling settings.

2. Click the **Delivery Settings** tab.
3. In the Settings panel, fill in these fields:

Field	Entry
Polling Method	Method to use to request queued documents on the partner's system.
Polling Frequency	How often to poll the partner's system for documents.

Suspending or Resuming Document Delivery for a Partner

At times you might want to suspend delivery for a partner (for example, when a partner's system is unavailable). The table below describes Trading Networks actions based on the type of delivery.

Type of delivery	Trading Networks actions when delivery is suspended
Immediate delivery method that uses reliable delivery	Delivery tasks for the partner are assigned HELD status and Trading Networks does not try to deliver documents. When delivery is resumed for the partner, Trading Networks changes the status to PENDING and resumes trying to deliver documents.
Immediate delivery method that does not use reliable delivery	Delivering documents fails. Because Trading Networks does not save the document to the database, it cannot hold the delivery.
Scheduled delivery method that uses reliable delivery	Delivery tasks for the partner are assigned HELD status and Trading Networks does not try to deliver documents. When delivery is resumed for the partner, Trading Networks changes the status to QUEUED and resumes trying to deliver documents.
Scheduled delivery method that does not use reliable delivery	Delivering documents fails. Because Trading Networks does not save the document to the database, it cannot hold the delivery.

Delivery tasks are also affected when deliveries are suspended. When delivery is suspended for a partner, Trading Networks assigns the delivery tasks HELD status and does not try to execute the services. When delivery is resumed, Trading Networks updates the task statuses to PENDING and resumes trying to execute the services.

To suspend or resume delivery for a partner

1. On the Partner Profiles page, select the profile for which you want to suspend or resume delivery.
2. Click the **Delivery Settings** tab.
3. You can do either of the following:
 - Suspend the delivery immediately. For this, select the **Suspend Delivery** check box in the Settings panel. To resume delivery, clear the check box.
 - Suspend the delivery using a schedule. For more information, see "[Scheduling a Delivery Suspension](#)" on page 201.

Note: When you select the **Suspend Delivery** check box in the Settings panel for a partner, the suspension schedules set for the partner will be overridden, and all the schedules will become inactive. You cannot add, delete, or edit the delivery schedules for the partner. When you clear the **Suspend Delivery** check box, the status of all the suspension schedules will be restored.

Scheduling a Delivery Suspension

Use the following procedure to schedule a suspension for a time period.

Scheduling a delivery suspension

1. In the **Suspend Delivery Schedule** panel, click **Add** to create a new schedule.
2. In the **Suspend Delivery Schedule** window, do the following.
 - a. Provide a suitable name for the schedule.
 - b. Hover the mouse to set the status to  (enabled) or  (disabled). By default, the status of a schedule is .
 - c. Select a time zone for the schedule. You can schedule a document delivery suspension according to a partner's time zone.
 - d. Choose from **Start** to set the start time, and from **End** or **Duration** to set the end time in the **Suspend Time** panel. If you select the **Duration** check box, select the duration of the suspension schedule in hours or minutes.
 - e. Choose any of the following options from the **Recurrence Pattern** panel to schedule the delivery suspension.
 - **None.** Suspends delivery once for the chosen date and time.
 - **Daily.** Suspends delivery every day for the chosen time.
 - **Weekly.** Suspends delivery every week on the chosen days of the week.
 - **Monthly.** Suspends delivery every month. For example, if you want to schedule a suspension on the 25th of every month, select **Day** and type 25 in

the text box provided. Or, if you want to schedule a suspension on the third Monday of every month, select **The**, and then select Third and Monday from the drop down.

Note: If you want to schedule a delivery suspension on the 31st of every month, the suspension schedule will run only for the months that have 31 days.

- **Yearly.** Suspends delivery every year for the chosen month of the year. For example, if you want to schedule a suspension on January 25th every year, select **On**, followed by January from the drop down and type 25 in the text box provided. Or, if you want to schedule a suspension on the third Monday of January every year, select **On The**, followed by Third, Monday, and January respectively.
- f. Choose from **Start** and **End** in the **Validity Period** panel to set the validity of the suspension schedule. By default, the end date is disabled. To specify the end date, select **Select End Date** checkbox.

If the recurrence pattern of a schedule is **None**, there is no end date for the schedule. The delivery will be suspended once according to the start and end times chosen. For example, if you choose the start time as 10.00 IST and end time as 14.00 IST on January 25th, the delivery will be suspended from 10.00 IST to 14.00 IST. The delivery will resume after 14.00 IST. If you choose a duration of 24 hours, the delivery will be suspended from 10.00 IST on January 25th till 10.00 IST on January 26th.

3. Click **OK**.

Once you create a suspension schedule, you can view it in the **Suspend Delivery Schedule** panel. To change the status of multiple schedules, select the check box beside the schedule and click **Active** or **Inactive**. Similarly, click **Delete** to delete the schedules.

The **Next Suspension Window** column in the **Suspend Delivery Schedule** panel indicates the next delivery suspension planned.

The **Suspension Status** indicates if the delivery was suspended as a result of the schedule. The values are `true` and `false`. If the delivery was suspended because of the schedule, `Suspension Status = true`. If the delivery was suspended, but not because of the schedule (for example, if the user manually selected the **Suspend Delivery** check box in the Settings panel before the schedule could run), or if the delivery was not suspended at all, then `Suspension Status = false`.

Editing an Ongoing Delivery Suspension Schedule

You can increase or decrease the duration or validity of an ongoing delivery suspension schedule. Use the following procedure to edit an ongoing delivery suspension schedule.

1. Locate the ongoing delivery suspension schedule in the **Suspend Delivery Schedule** panel.
2. Click  to edit the end time or end date of the schedule.

Canceling an Ongoing Delivery Suspension Schedule

You can cancel an ongoing delivery suspension schedule, and resume delivery to the partner before the end time of the schedule. Use the following procedure to cancel an ongoing delivery suspension schedule.

1. Locate the ongoing delivery suspension schedule in the **Suspend Delivery Schedule** panel.
2. Click  to cancel the schedule. .

Delivery of documents to the partner will be resumed immediately, and the suspension schedule will remain inactive until you manually change the status to active.

Associate My webMethods User Accounts with a Profile

You can associate one or more My webMethods user accounts with a profile. When you do so, the user account receives Trading Networks partner authority, and the profile owner can use the user account to access Trading Networks.

This procedure explains how to create My webMethods user accounts. An account can specify a user name and password, or it can specify one of the external IDs defined in the profile. The procedure then explains how to associate the profile with an account.

To associate a profile with My webMethods user accounts

1. On the Partner Profiles page, select the profile with which you want to associate a My webMethods user account.
2. Click the **Users** tab.
3. Create a My webMethods user account using one of the following procedures.
 - Create a user account that specifies an external ID by clicking **Add user from external ID**, selecting the external ID to use, and supplying the password to use.
 - Create a user account that specifies a user name and password as follows:
 - i. Click **Add Users**.
 - ii. Click **Add new user** and fill in the fields below.

Field	Entry
User name	User name to assign to the user account.

Note: My webMethods Server adds the user ID to the internal system directory service using the character casing you specify. My webMethods Server typically regards user IDs as case-insensitive; however, My webMethods

Field	Entry
	Server uses the case you specify for actions that are case-sensitive (for example, HTTP authentication). If it is important for Trading Networks to consider case when evaluating user names, particularly when using an Oracle or SQL Server database, consider setting the tn.checkUser.ignoreCase property.
Password and Confirm Password	Password for the user account.

4. Associate the user accounts with the profile as follows:
 - a. Select **Add existing user**.
 - b. Find the user accounts to associate with the profile. Specify a keyword that is contained in the first name, last name, e-mail address, or user name of the user accounts in the **Search for Available Users** field and click **Go**.
 - c. Move the user accounts to associate with the profile from the **AVAILABLE USERS** to the **SELECTED USERS** list.

Adding Security Information

When you add a certificate to a profile, you must also add the CA certificates and private key that are associated with the certificate; all must reside in your file system. Trading Networks can import files that have a .der, .cer, or .p7b extension. If you import a .der file, Trading Networks appends it to the existing certificate chain list. If you import a .cer or .p7b file, Trading Networks replaces the certificate chain list with the .cer or .p7b file.

Likewise, if you add a private key, you must also add the certificate and CA certificates that are associated with the private key.

Note: If you are adding this information to your Enterprise profile, you would specify the private key alias and the keystore alias associated with the certificate.

The following table lists the types of certificate sets you can add to a profile.

Type	Description
Sign/Verify	Certificate information to use to sign a document or verify the digital signature. In your profile, Trading Networks uses the private key associated with the sending partner to digitally sign documents your corporation sends to partners.

Type	Description
	<p>Trading Networks digitally signs a document when the <code>wm.tn.doc:sign</code> built-in service is invoked.</p> <p>In a partner profile, Trading Networks uses the sender's public key that is associated with this partner to verify the document that was digitally signed by the sender (Verify Digital Signature pre-processing action).</p>
Encrypt/ Decrypt	<p>Certificate information to use to decrypt or encrypt information.</p> <p>In your profile, Trading Networks uses the private key of the corporation to decrypt documents that partners send. When the corporation receives these documents from partners, the documents are encrypted using the public key of the corporation.</p> <p>In a partner profile, Trading Networks uses your public key to encrypt information that is being sent to you.</p> <div data-bbox="513 940 1312 1331" style="background-color: #f0f0f0; padding: 10px;"> <p>Note:</p> <ul style="list-style-type: none"> ■ Trading Networks maintains this information for other webMethods products, such as webMethods RosettaNet Module, that take advantage of this feature. You can also add your own functionality that takes advantage of this certificate information. You can obtain the certificate information by using built-in services. ■ Trading Networks does not check whether the CA that signed the certificate is included in the list of trusted CAs maintained by Integration Server. </div>
SSL	Certificate that enables Trading Networks to act as an SSL client and connect to a remote secure server.

Adding Certificate Sets to a Profile

When Trading Networks needs to verify the digital signature of a partner (Verify Digital Signature pre-processing action), it first looks for the partner's certificate set. If it cannot find that certificate set, Trading Networks uses the default certificate set.

Adding the Default Certificate Set to a Profile

To add the default certificate set to a profile

1. On the Partner Profiles page, select the profile to which you want to add the default certificate set.

2. Click the **Certificates** tab.
 3. Under the **Default** branch of the tree on the **Certificates** tab, click  in the **Add/Remove** column and click  in the **Edit** column for the certificate to add.
- Note:** The certificate set that you add first is the *primary* certificate set. This is indicated by a  in the **Status** column.
4. If you are working with a partner profile, add the private key associated with the certificate as follows:
 - a. Click **Browse** next to **Set Private Key** and select the file containing the private key.
 - b. To add a certificate, click **Add Certificate**, click **Browse**, and then upload the certificate file.
 - c. Click **OK**.
 5. If you are working with your Enterprise profile, add the private key associated with the certificate as follows:
 - a. From the **Keystore Alias** list, select the alias of the keystore containing the certificate.
 - b. From the **Key Alias** list, select the alias of the private key used to access the certificate.

Note: If the desired keystore alias is not listed, you can create one by clicking **Create Keystore Alias**. For details, see the section on creating keystore aliases in *webMethods Integration Server Administrator's Guide*.

- c. Click **OK**.
6. Optionally, add a second certificate set as follows:
 - a. Click  in the **Edit** column for the second certificate.
 - b. Click **Browse** next to **Set Private Key** and select the file containing the private key.
 - c. To add the certificate pair, click **Add Certificate**, click **Browse**, and then upload the certificate files.
 - d. Click **OK**.
7. Save the profile.

Adding a Certificate Set to Use with a Specific Partner

To add a certificate sets to use with a specific partner

1. On the Partner Profiles page, select the profile to which you want to add a specific certificate set.
2. Click the **Certificates** tab.

3. Click **Add Certificate Set**.
4. Search for and select the profiles of partners for which to add certificate sets. For details, see "[Searching for Assets](#)" on page 221).
5. For each partner, add a certificate set as follows:
 - a. Expand the tree for the partner.
 - b. Click  in the **Add/Remove** column and click  in the **Edit** column for the certificate to add.

Note: The certificate set that you add first is the *primary* certificate set. This is indicated by a  in the **Status** column.

- c. Add the private key associated with the certificate by doing one of the following:
 - **For a partner profile:** Browse to and select the private key file, and then upload the certificate file.
 - **For your Enterprise profile:** From the respective lists, select the keystore alias of the keystore containing the certificate and the alias of the private key used to access the certificate.
6. Optionally, add a second certificate set as follows:
 - a. Click  in the **Edit** column for the second certificate.
 - b. Click **Browse** next to **Set Private Key** and select the file containing the private key.
 - c. To add the certificate pair, click **Add Certificate**, click **Browse**, and then upload the certificate files.
 - d. Click **OK**.
7. Save the profile.

Updating Certificate Sets

You can add new certificates or private keys to, or remove certificates or private keys from, existing certificate sets associated with your Enterprise profile or a partner profile.

Updating Certificate Sets for Your Enterprise Profile

To update certificate sets for your Enterprise profile

1. Click the **Certificates** tab.
2. Expand the branch for the certificate set to update.
3. Click  in the **Edit** column for the certificate to update, and then do one of the following on the Edit Certificate Set dialog box:

- To add a private key, select the keystore alias and key alias from the respective lists, and then click **OK**.
 - To change the private key, select a new keystore alias and key alias, and then click **OK**.
 - To remove the mapping between a keystore and the profile, select a blank entry from the **Keystore Alias** list, and then click **OK**.
4. To remove a certificate set from the profile, select the check box to the left of the certificate set on the Certificates tab, and then click **Delete**.

Updating Certificate Sets for a Partner Profile

To update certificate sets for a partner profile

1. Click the **Certificates** tab.
2. Expand the branch for the certificate set to update.
3. Click  in the **Edit** column for the certificate to update, and then do the following on the Edit Certificate Set dialog box:
 - To remove a private key, select the **Clear Private Key** check box.
 - To add a private key, click **Browse** next to **Set Private Key** and select the private key file.
 - To add a certificate to the chain, click **Add Certificate** and browse to the certificate to add.
 - To remove a certificate from the chain, select the check box to the left of the certificate to remove, and then click **Delete**.
4. To remove a certificate set from the profile, select the check box to the left of the certificate set on the Certificates tab, and then click **Delete**.

Making Standard Fields Required

To make standard fields required

1. Open the profile:
 - To open your profile, in My webMethods: navigate to **Administration > Integration > B2B > Partner Administration > Partner Profiles**, then click **Enterprise**.
 - To open a partner profile, search for the profile, then click  in the **ACTION** column for the profile.
2. Click the **Extended Fields** tab.
3. Click **Manage Extended Fields**.
4. Click the **Standard** tab.

5. Click the **Actions** icon for any standard field.
6. Select the **Required** check box for the standard fields to make required.

Adding Extended Fields to a Profile

Opening a Profile's Extended Fields

To open a profile's extended fields

1. Open the profile:
 - To open your profile, in My webMethods: navigate to **Administration > Integration > B2B > Partner Administration > Partner Profiles**, then click **My Enterprise**.
 - To open a partner profile, search for the profile, then click  in the **ACTION** column for the profile.
2. Click the **Extended Fields** tab.
3. Click **Manage Extended Fields**.

Adding Extended Fields to a Profile

To add an extended field to a profile

1. Do one of the following:
 - To create an extended field by providing new information, click **Add Extended Field**.
 - To create an extended field by copying an existing extended field and modifying it, click  in the row for the extended field.
2. Fill in the fields below.

Field	Entry
Name	Name of the extended field.
Enabled	Whether to enable or disable the extended field. When you enable a field, it appears in all profiles that contain it. When you disable a field, it does not appear in any profiles. If you create profiles while a field is disabled, the profiles will not contain the field. If you later re-enable the field, if the field is required, you will have to update the new profiles manually.

Field	Entry
Required	Whether the extended field is required or optional. When you define an extended field, Trading Networks automatically adds it to all profiles you create after you add the field, but does not add the field to existing profiles. If you make a field required, therefore, you must add it to existing profiles manually.
Group	Profile field group to which to assign the extended field. If you need to add a new group, type the name of the new group in the Add New Group section of the dialog box and click Add Group .
Description	Optional. Description of the extended field.
Data Type	Data type of the extended field. For String , you can specify the following: <ul style="list-style-type: none"> ■ Maximum Length: Maximum number of characters to allow in the field. The default is no limit. ■ Default: Default value for the field. ■ Valid Values: Acceptable values for the extended field. Click , type a valid value, and repeat as necessary. To remove a value, click .

Associating Partners with Partner Groups

To associate partners with a partner group

1. In My webMethods: navigate to **Administration > Integration > B2B > Partner Administration > Partner Groups**.
2. Click **Add Partner Group**.
3. Type a name and description for the group.
4. Click **Add Members**.
5. Search for profiles of the partners you want associate with the group (see "[Searching for Assets](#)" on page 221).
6. Move the profiles of the partners to associate with the group from the **AVAILABLE PARTNERS** list to the **SELECTED PARTNERS** list.
7. Click **Save & Close**.

Working with TPAs

Defining a TPA

You can define a TPA between any two partners that have partner profiles.

To define a TPA

1. In My webMethods: navigate to **Administration > Integration > B2B > Trading Partner Agreements**.
2. Do one of the following:
 - To create a TPA by providing new information, click **Add TPA**.
 - To create a TPA by copying an existing TPA and modifying it, click  in the row for the TPA to copy.
3. Complete these fields:

Note: The **Sender**, **Receiver**, and **Agreement ID** fields together uniquely identify a TPA. After you create a TPA, you cannot change or update these fields.

Field	Entry
Agreement ID	Unique identifier that indicates the type of agreement between the two partners. The value you specify depends on the application that is going to use the TPA.
Sender	Name of the trading partner that has the sender role in the transaction the TPA will govern. Type the name of the sender or click Edit to select the sender from the Select Partner dialog box. For EDI, to create a template that you will duplicate to create other TPAs, you can use the default value Unknown .
Receiver	Name of the trading partner that has the receiver role in the transaction the TPA will govern. Type the name of the receiver or click Edit to select the receiver from the Select Partner dialog box. For EDI, to create a template that you will duplicate to create other TPAs, you can use the default value Unknown .
Control Number	Value that the application using the TPA expects. For example, the webMethods Module for EDI uses the

Field	Entry
Description	Optional. Description for the TPA.
IS Document Type	<p data-bbox="591 653 1351 909">IS document type that defines the application-specific TPA data. Type the name of the IS document type or click Edit to browse the IS document types and select the one to use.</p> <p data-bbox="591 772 1351 909">Trading Networks displays the data tree with all the high-level parameters of the selected IS document type in the bottom panel of the screen. Expand the high-level parameters to view/add data.</p> <p data-bbox="591 936 1351 1392">For parameters that contain an array of records (📄), click 📄 to view/add values for the elements for the array. When you click 📄, Trading Networks displays the Edit Record List pop-up window where you can create rows and add values for the child elements in the form of a table. Here, each row corresponds to an element of the array. For example, if DocumentSettings is a parameter of type array, Edit Record List helps you create and add values for DocumentSettings[0], DocumentSettings[1], DocumentSettings[2], and so on. The table format also makes it easy to compare the values specified for DocumentSettings[0], DocumentSettings[1], DocumentSettings[2], and so on.</p> <p data-bbox="591 1419 1351 1514">If you already have a file that contains data for the selected IS document type, browse to locate the file and click Load Inputs.</p> <p data-bbox="591 1541 1351 1841">When using the TPA feature for your own application, these are the parameters that your applications need access to during document exchange. webMethods products that use the TPA feature come with the IS document type to use for the TPA data. For example, the webMethods Module for EDI supplies the wm.b2b.editn.TPA:EDITPA IS document type, while the webMethods ebXML Module supplies the wm.ip.ebxml.cpa.rec:tpaDataSchema IS document type.</p>

Field	Entry
Data Status	When the TPA's agreement status is Agreed (see "Changing a TPA's Agreement Status" on page 214), whether the values for the TPA data defined in the IS document type can be modified.
Initialization Service	Optional. Service that populates the inputs to the variables in the IS document type with default values. Type the name of the service or click Edit to browse the services and select the one to use.
Validation Service	Optional. Service to validate the data added to the IS document for the TPA. Type the name of the service or click Edit to browse the services and select the one to use.
	Note: Specifying a validation service improves processing time because it eliminates the need to validate the TPA every time a process that uses the TPA runs on Integration Server.
Export Service	Optional. Service that exports a TPA and converts it to an industry-standard format. Type the name of the service or click Edit to browse the services and select the one to use. After creating a TPA, if you need to obtain the industry-standard format of a TPA, you can execute your export service.

Opening a TPA for Editing

To open a TPA for editing

1. In My webMethods: navigate to **Administration > Integration > B2B > Trading Partner Agreements**.
2. Search for the TPA (see ["Searching for Assets" on page 221](#)), then double click the row containing the TPA.

Editing a TPA

Editing a TPA's Fields

While editing a TPA, if another user tries to edit at the same time, Trading Networks ensures that data written by you is not over written by the other user and vice versa. While saving the TPA, Trading Networks ensures that you are working on the latest

version of data, and notify you if another transaction has modified the TPA. In that case, you have to cancel the changes you have made, and reopen the TPA for edit.

To edit a TPA's field

1. Select **Agreements > Edit**.
2. Update the TPA fields. You cannot modify or update the **Sender, Receiver, or Agreement ID** field.
3. Update the TPA data input in the IS document type fields using one of these methods:
 - To change the values in the IS document type, browse to locate the file that has data for the IS document type and click **Load Inputs**.
 - If an initialization service is specified and you want to invoke it to reset the TPA data inputs of the IS document type to their default values, click **Reset Inputs**.
4. If a validation service is specified and you want to validate the TPA data, click **Validate Data**.
5. Click **Save & Close**.

Changing a TPA's Agreement Status

To change a TPA's agreement status

1. In My webMethods: navigate to **Administration > Integration > B2B > Trading Partner Agreements**.
2. Select the rows for the TPAs whose agreement status to change and click the appropriate status change button, as follows:

Agreement Status	Meaning
 Proposed	TPA is in draft status. Use this status when you are creating the TPA so you can modify the TPA fields, data values in the IS document type, and agreement status.
 Agreed	<p>TPA is in final status. In production, use this status to prevent the accidental deletion of the TPA and to speed up processing. When a TPA is in this status, you can change the Control Number. If the Data Status is Modifiable, you can also change inputs to TPA data variables in the IS document type.</p> <p>When the agreement status is Agreed, the data statuses take effect. Additionally, after the agreement status is Agreed, you <i>cannot</i> delete the TPA agreement.</p>
 Disabled	TPA should not be used. For example, if Partner A is delinquent on payments to Partner B, Partner B might change

Agreement Status	Meaning
	Partner A's agreement status to Disabled so that no further transactions can occur. After Partner A pays his payments, however, Partner B would change the agreement status to Agreed . When a TPA is in the status Disabled , you can change the Control Number .

Deleting a TPA

You can delete any TPA whose agreement status is **Proposed** or **Disabled**. When you delete a TPA, Trading Networks deletes all of the TPA data from the its database.

Note: If you want to delete a non-modifiable agreement, use the `wm.tn.tpa:deleteTPA` built-in service.

Deleting a Partner Profile

When you delete a partner profile, you can no longer view the profile or any related TPAs. Trading Networks removes most of the profile data from its database but keeps all documents, entries, and TPAs associated with the partner.

Search for the profile to delete (see "[Searching for Assets](#)" on page 221). Then Select the check box beside the profile and click **Delete**.

Note: You cannot delete your (the Enterprise) profile.

Finding Partner Certificates that Are Expired or Expiring Soon

Finding Expired Partner Certificates

To find expired partner certificates

1. In My webMethods: navigate to **Administration > Integration > B2B > Partner Administration > Partner Certificates**.
2. Select **Expired** from the **EXPIRING RANGE** list.
3. Click **Go**.

Finding Partner Certificates that Expire Soon

To find partner certificates that expire soon

1. In My webMethods: navigate to **Administration > Integration > B2B > Partner Administration > Partner Certificates**.
2. Specify a time period, as follows:
 - Select a predefined time period from the **EXPIRING RANGE** list.
 - Specify a custom time period. Select **Custom** from the **EXPIRING RANGE** list, use the calendar pickers to specify a start and end dates, and then select the hours and minutes.
3. Click **Go**.

Managing Partner Access to APIs

Using APIs, you can enable business-to-business communication between trading partners by integrating Trading Networks with a API Gateway instance.

You can associate an API Gateway instance with Trading Networks to expose the APIs available in API Gateway and manage API access for each partner. When you add an API to a partner for the first time, a new application with the partner name is created in API Gateway. Any addition or removal of API access to that partner in Trading Networks is updated in the corresponding application on API Gateway. All the supported API types such as REST, ODATA, or SOAP can be exposed to your partners using partner profiles. As an administrator, you can either create a new user or provide access to an existing user in Trading Networks to add or remove API access from the **Partner Profile** page.

The access details such as the API key and authentication mechanism of the associated API Gateway instance appears on the **API** tab of the **Partner Profiles** page. For instructions on how to invoke an API, see the *webMethods API Portal Consumer's Guide*.

Integrating Trading Networks with API Gateway

To integrate Trading Networks with API Gateway, ensure that the following setup is available:

- You must have the permissions to manage applications and APIs in the API Gateway instance.
- APIs that you want to consume must be available in the API Gateway instance when you establish a connection.

- Every API that Trading Networks consumes must have its Identity and Authorize Access Policy values set in API Gateway so that API Gateway can identify the trading partners.
- The API Gateway instance to which you want to connect must have a remote server alias configured in the Integration Server hosting Trading Networks. For instructions on how to create a remote server alias, see *webMethods Integration Server Administrator's Guide*.
- The API Gateway instance to which you want to connect must have the `tn.api.gateway.server.alias` property configured in Trading Networks.

Configuring the API Gateway Property

You need to set a property in Server Configuration Properties for Trading Networks to communicate with an API Gateway instance. Ensure that you create a remote server alias in Integration Server for the API Gateway instance to which you want to establish a connection.

To configure the API Gateway property

1. In My webMethods, go to **Applications > Administration > Integration > B2B Settings > Configure Properties > Trading Networks Configuration Properties**.
2. In Trading Networks Configuration Properties, Type `tn.api.gateway.server.alias=API Gateway remote server alias`.
3. Click **Save**.

This ensures that Trading Networks has the information about the availability of an API Gateway instance to which it can connect.

Viewing the APIs Accessible to a Partner

You can view the list of APIs that are accessible to a partner from an associated API Gateway instance in Trading Networks.

To view the APIs accessible to a partner

- In My webMethods, go to **Applications > Administration > Integration > B2B > Partner Administration > Partner Profiles > API**, click **Add API** to view the list of APIs available for the partner.

Viewing the List of APIs of a Partner Group

A partner group displays all the APIs that belong to all the partners in that group. A single API in a group may be accessible to one or more partners. The **API Associations** column of the **API** panel displays the partners to which the API is accessible. For

example, if P1 and P2 have access to an API, both P1 and P2 appear in the **API Associations** list.

To view APIs in a Partner Group

- In My webMethods, go to: **Applications > Administration > Integration > B2B > Partner Administration > Partner Groups > APIs** to view the list of APIs in a partner group.

Providing API Access for a Partner using Partner Profiles

A partner can consume the APIs exposed through an API Gateway instance after you add the APIs to the partner.

To add APIs to a Partner Profile

1. In My webMethods, go to **Applications > Administration > Integration > B2B > Partner Administration > Partner Profiles**, click on a partner to provide API access.
2. In the **APIs** tab and click **Add API**. A list of APIs that are exposed to the Trading Networks instance appear in the **Add API** dialog box.
3. Select the APIs that you want to add to the **Partner Profile** and click **Apply**.
4. Click **Save & Close** to return to the **Partner Profiles** page.

Revoking API Access for a Partner from Partner Profile

To revoke API access for partners, you must delete the APIs associated with the partners from the Partner Profile tab.

To revoke API access for a Partner

1. In My webMethods, go to **Applications > Administration > Integration > B2B > Partner Administration > Partner Profiles**, click on a partner.
2. In the **APIs** tab, select the APIs and click **Delete**.
3. Click **OK** to confirm the deletion.
4. Click **Save & Close** to return to the **Partner Profiles** page.

Granting Permissions to a User to Access APIs

A user must have permissions to view or edit the APIs. See "[Granting Data Permissions to a Data Set](#)" on page 98 for instructions to provide data permissions.

Providing API Access to a Partner Group

When you add an API to a partner group, all the partners in that group gain access to the API.

To add API access to a Partner Group

1. In My webMethods, go to **Applications > Administration > Integration > B2B > Partner Administration > Partner Groups**.
2. Click the partner group to which you want to add APIs.
3. In the APIs panel of the **Edit Partner Group** page, click **Add APIs** to add APIs to all the members in the partner group.
4. Click **Save & Close**.

Revoking API Access for a Partner Group

When you remove an API to a partner group, all the partners in that group lose access to the API.

To remove API access for a Partner Group

1. In My webMethods, go to **Applications > Administration > Integration > B2B > Partner Administration > Partner Groups**.
2. Click the partner group from which you want to delete the APIs.
3. In the APIs panel of the **Edit Partner Group** page, select the APIs that you want to delete from all the partners in the partner group and click **Delete**.
4. Click **Save & Close**.

13 Searching for Assets

■ Setting Search Options	222
■ Searching Using Keywords (Simple Search)	222
■ Searching Using Keywords and Filters (Advanced Search)	223
■ Improving Trading Networks Performance During Searches	224

Setting Search Options

You can specify options for the pages on which you search. You can do the following:

- Specify the search tab to display and the search to execute, if any, when you first display the page.
- Specify the maximum number of search results to display, the number of rows to display on each results page, the columns to display, and how to sort the results.

For instructions on setting these options, see *Working with My webMethods*.

You can use the wildcard * in searches to match one or more characters, or the wildcard ? to match one character. Searches are case insensitive. You can save searches so you can re-execute them later.

Searching Using Keywords (Simple Search)

To search using keywords

1. In My webMethods: **Administration > Integration > B2B > Partner Administration > asset to search for**.
2. Click the **Keyword** tab and do the following in the text box:

To search for all assets of the specified type that...	Do this in the text box...
Exist in Trading Networks that you are authorized to view	Leave it blank.
Contain one or two keywords	Type one keyword or two keywords separated by a space.
Contain a specified string	Type a string enclosed in double quotes.

3. Click **Go**. Trading Networks searches as follows:

For these assets...	Trading Networks searches...
Profiles	The profile external IDs and corporation name. <ul style="list-style-type: none"> ■ If you type one keyword or a string, it must be either an external ID or the corporation name.

For these assets...	Trading Networks searches...
	<ul style="list-style-type: none"> ■ If you type two keywords, one must be an external ID and the other must be the corporation name.
TPAs	The TPA agreement IDs, sender name, and receiver name. The keyword or a string you type must be either an agreement ID, sender name, or the receiver name.
Other assets	The asset names.

Searching Using Keywords and Filters (Advanced Search)

To search using keywords, filters, and search conditions

1. In My webMethods: **Administration > Integration > B2B > Partner Administration > asset to search for**
2. Click the **Advanced** tab.
3. In the text box, you can type a keyword as described in "[Searching Using Keywords \(Simple Search\)](#)" on page 222
4. Click the **Filters** tab. You can specify one or more search filters. For each filter, specify the field, operator, and value to use in the search. To add another filter, click  . Specify values as described below.
 - For **Name**, **Agreement ID**, and **Description**, type a string.
 - For most other fields, click **Edit** to search for and specify the value.
 - You can search for EDI document types only if the webMethods Module for EDI is installed on the Integration Server that hosts Trading Networks. If you want to search for document types other than XML, flat file, or EDI, select **Other**.
 - In a processing rules search, for **Sender Group** and **Receiver Group**, you can expand the search results to reflect the group membership of the groups you specify. For example, if you select a sender group, you can expand the search to return all rules where the sender is that sender group, the sender is a member of that sender group, or the **Sender** criterion for the rule is set to **Any**. To expand the search, select **Perform Implicit Matching**.
 - You can search for processing rules based on rule ordinals. Ordinals correspond to the positions of rules in the rules list.
5. In the **Search Condition** list, select **AND** to search for assets of the specified type that match all search criteria, or **OR** to search for assets of the specified type that match any search criterion.

Note: You can specify multiple criteria for a single field, and select **AND** or **OR** to search for assets.

6. In a processing rule search, you can specify the criteria order for Trading Networks to follow when searching. Click the **Options** tab. In the Search Order panel, click  and  to move search criteria up or down, respectively, and click **Save**.
7. In a custom attribute, document type, processing rules, or TPA search, you can narrow the results to only those assets that were modified within a specified date range. Use the fields in the **Last Modified** or **Date Range** panel to identify the date range.
8. Click **Go**.

Improving Trading Networks Performance During Searches

Setting a Limit for Search Results

1. Shut down the Integration Server that hosts Trading Networks.
2. Open the `Integration Server_directory \instances\instance_name \config \Cache\webm-cache-config.xml` file and edit these XML tags in the XML tag `<localScheme\>`:

Tag	Value
<code><highUnits></code>	<p>Maximum number of search result rows to cache. Use the format:</p> <pre><highUnits>number_of_rows</highUnits></pre>
<code><lowUnits></code>	<p>Number of search result rows in the cache that will trigger deletion. Use the format:</p> <pre><lowUnits>number_of_rows</lowUnits></pre> <p>Trading Networks uses the <i>Hybrid eviction policy</i> to delete search result rows from the cache. This policy determines which rows to delete based on how often and how recently the rows were accessed; Trading Networks deletes the rows that were least frequently accessed and were not accessed for the longest period.</p> <p>The recommended value for <code><lowUnits></code> is 75 percent of <code><highUnits></code>.</p>

3. Restart Integration Server.
4. In My webMethods: **Administration > Integration > B2B Settings > Configure Properties**.

5. Add or edit these properties:

Property	Value
<code>tn.db.fetchMaxRows</code>	Maximum number of search result rows to retrieve from the database for each search. The default is 10000 rows.
<code>tn.db.sortDocTimestamp</code>	Whether to sort the search result rows based on timestamp in descending order. Values are <code>true</code> (the default) or <code>false</code> .

6. Click **Save**.

Deleting Search Results Periodically

To delete search result rows periodically

1. Shut down the Integration Server that hosts Trading Networks.
2. Open the `Integration Server_directory \instances \instance_name \config \Cache \webm-cache-config.xml` file and edit these XML tags in the XML tag `<localScheme>`:

Tag	Value
<code><expiryDelay></code>	Time in minutes after which a cached search result row is considered expired. Trading Networks can delete a cached row only after it has expired. Use the format: <code><expiryDelay>time</expiryDelay></code>
<code><flushDelay></code>	Time in minutes between periodic cache flushes of the expired search result rows. Use the format: <code><flushDelay>time</flushDelay></code> The first flush starts when the number of rows specified in <code><lowUnits></code> have been cached.

3. Restart Integration Server.

14 Migrating Assets

■ Overview	228
■ Trading Networks Assets You Can Export	228
■ Migration Methods	229
■ Trading Networks Asset Dependencies	229
■ How Trading Networks Detects Assets in the Target System Before Importing Them	232
■ Migrating Assets Using webMethods Deployer	233
■ Migrating Using My webMethods	235
■ Migrating Assets Using the tlexport and timport Utilities	238

Overview

You can migrate Trading Networks assets such as document attributes, document types, and partner profiles from one Trading Networks database to another. Migrate assets when:

- You want to deploy assets from a development environment to a production environment.
- You have multiple Trading Networks instances, each with its own database, and you want each database to have identical assets. You can create the assets on one Trading Networks instance and then migrate the assets to the other instances.
- You want to change the type of database you use for Trading Networks. For example, you were using an Oracle database and now want to use a SQL Server database.

Important: Only use the procedures in this chapter to migrate Trading Networks assets between Trading Networks instances of the same release. If you need to migrate assets from one release of Trading Networks to another, follow the instructions in the *Upgrading Software AG Products*. You cannot share your Enterprise profile with a partner.

Trading Networks Assets You Can Export

You can migrate the Trading Networks assets listed below.

- Data permissions and general functional permissions.
- Document-related assets, as follows:
 - Custom document attributes.
 - Document types.
 - Binary types. Binary types contain the information related to document type categories (flat file, XML, and EDI) and Trading Networks query definitions.
- Processing rules.
- Profile-related assets, as follows:
 - Partner profiles. For each partner, Trading Networks migrates the partner's standard profile fields; private queue; certificate information, if any; and extended profile fields, if any.
 - Custom external ID types.
 - Profile field definitions for standard and extended profile fields.
 - Custom field groups.

- Custom contact types.
- Public delivery queues.
- Partner groups.
- TPAs.
- Archive schedules and criteria for archiving and deleting documents.

Trading Networks assets are available for migration regardless of whether they are disabled, and the state of the assets on the source system is maintained on the target system. Deleted assets are not included in the migration process.

You cannot migrate data for documents; that is, you cannot migrate document content, tasks, or activity log entries.

Migration Methods

You can migrate all assets, all assets of a certain type, or selected assets within an asset type. Below are the methods you can use to migrate assets. Unless otherwise noted, procedures for each method are provided in this chapter.

- Export assets to file using the export feature in My webMethods, then deploy assets using webMethods Deployer. Software AG recommends using this method.

Note: You can use this method to deploy the assets of *all* webMethods products on one Integration Server to another Integration Server using a single script. For complete information, see *webMethods Deployer User's Guide*.

- Migrate Trading Networks assets using the export and import features in My webMethods.
- Migrate Trading Networks assets using the `tnexport` and `tnimport` command line utilities.
- Migrate Trading Networks assets using the `wm.tn.admin:exportData`, `wm.tn.admin:extendedExportData`, `wm.tn.admin:importData`, and `wm.tn.admin:extendedImportData` built-in services.

Trading Networks Asset Dependencies

Some assets require other assets. For example, profiles use assets such as partner groups and contact types, and document types use assets such as document attributes. For assets you migrate to work properly, these required assets must also exist on the target system.

In My webMethods, when you select an asset for export, My webMethods shows dependencies on other assets and you can choose to include those required assets in the

export. When you select an asset for import, however, My webMethods shows some, but not all, of the dependencies.

If you use the command line utilities or the built-in services to migrate, you must manage asset dependencies yourself.

The table below lists all possible dependencies an asset might have. When the table instructs you to migrate an asset, see below for instructions. When the table instructs you to create an asset, see the appropriate chapter in this guide for instructions. You must use the exact same name on the target system as the source system, with the same upper and lower casing.

If you migrate...	Do this to the target system...
Partner profile	<ul style="list-style-type: none"> ■ Migrate extended profile field definitions ■ Migrate any custom external ID types, custom contact types, and partner groups. ■ Migrate any public queues. ■ Create any custom immediate delivery methods.
Queues	Move any custom immediate and scheduled delivery services.
Extended field definition	Migrate any custom profile field groups. The standard groups are Corporation, Address, Contact, Delivery, External ID, Custom, and EDINT (if this is installed). All groups other than these are custom groups. To know the custom groups, execute the getFieldGroups service in the wm.tn.dictionary package. Standard groups need not be migrated because they would be available in the target system.
Document type	<ul style="list-style-type: none"> ■ Migrate any custom document attributes. ■ Move any custom services for document attribute transformation. ■ Move any custom service for the Check for Duplicate Document pre-processing action.
Processing rule	<ul style="list-style-type: none"> ■ Move any custom services for the Check for Duplicate Document pre-processing action and the Execute a Service processing action. ■ For the Deliver Document By action: <ul style="list-style-type: none"> ■ Create any custom immediate delivery method. ■ Migrate any services for the custom immediate delivery method and then use the wm.tn.delivery:registerService service to register the corresponding immediate delivery service using the same immediate delivery method name.

If you migrate...	Do this to the target system...
	<ul style="list-style-type: none"> ■ Manually create queues or use the tlexport and timport utilities to export all public queues. ■ Move any scheduled delivery services. ■ Migrate partner profiles for senders and receivers used as matching criteria. ■ Migrate document types specified as matching criteria. ■ Migrate any custom document attributes used as extended matching criteria.
TPA	<ul style="list-style-type: none"> ■ Migrate partner profiles for the sender and receiver associated with the TPA. ■ Move any initialization, validation, and export services. ■ Migrate the IS document type that defines the application-specific TPA data.
Data permissions	<ul style="list-style-type: none"> ■ Migrate partner profiles, document types, and processing rules that are part of data set definitions. ■ Migrate My webMethods roles that have permission to act on data sets. <div data-bbox="493 1178 1365 1377" style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p>Note: These roles are defined in the My webMethods Server database, not the Trading Networks database, so you cannot migrate them using the tlexport and timport command line utilities. Make sure that the target system uses the same My webMethods Server.</p> </div> <ul style="list-style-type: none"> ■ Grant data permissions to act on a data set to at least one My webMethods role.
General functional permissions	<p>Migrate My webMethods roles that have permission to functions. See the note for Data permissions.</p>
Archive schedule	<ul style="list-style-type: none"> ■ Migrate any partner profiles for the sender and receiver used as selection criteria for archive or deletion. ■ Migrate any document type used as selection criteria for archive or deletion.

If you do not export the required partner profiles and document types for data permissions, Trading Networks throws a warning when you import that says the required assets do not exist in the target system. It then continues to import the other assets.

If the export file contains the required/dependent assets for any asset, timport will ensure that the required assets get migrated first so that no error is thrown.

When importing document types, Trading Networks determines whether the required document attributes exist on the target system with the same name. If they do not, Trading Networks throws an error and does not import the document type.

When importing processing rules, Trading Networks determines whether the required senders, receivers, and document types exist on the target system with the same name. If they do not, Trading Networks throws an error and does not import the processing rule.

When importing partner profiles, if any of the required assets are not defined, Trading Networks throws an error and does not import the profile. For example, if the contact type is not already defined in the target system before you import a contact of that type, Trading Networks throws an error and does not import that profile.

When importing TPAs, if the associated senders and receivers do not exist in the target system with the same name, Trading Networks throws an error and does not import the TPA.

How Trading Networks Detects Assets in the Target System Before Importing Them

When you import an asset, Trading Networks checks whether an asset that has the same internal ID already exists in the target system. If not, in general Trading Networks checks for an asset that has the same name; for profiles, Trading Networks checks for a profile with the same corporation name and unit name, and for TPAs, Trading Networks checks for a TPA with the same sender name, receiver name, and agreement ID.

If Trading Networks finds a matching asset using either check, it asks whether it should overwrite that asset in the target system. If Force Import is true and Trading Networks finds a matching asset using either check, it will overwrite that asset in the target system.

When you import processing rules, Trading Networks replaces matching processing rules and appends non-matching processing rules on the target system according to the rule order that already exists on the target system. If you use My webMethods to import processing rules, you can choose to replace all the processing rules on the target system with the imported processing rules.

Migrating Assets Using webMethods Deployer

Selecting Trading Networks Assets For Export

To select the assets for export

1. In My webMethods: **Administration > Integration > B2B > Asset Management > Export**.
2. From the **TN Server** list, select the Trading Networks Server instance from which to export assets.
3. If you want to export all Trading Networks assets, select the **Select all assets** check box.
4. To select assets by asset type, click the tab whose name corresponds to the asset type.

Note: The Partner Profile tab lists each partner, including your corporation, by corporation name.

5. Do the following in each tab:
 - a. If you want to export all assets of this type, select the **Select all** check box.
 - b. If you want to export only certain assets, type a string to find those assets in the **Keywords** field and click **Search**. Trading Networks searches the same fields it searches when you search for assets (see "[Searching Using Keywords \(Simple Search\)](#)" on page 222). In the search results, do the following to create the list of assets to export:

<u>To do this...</u>	<u>Take this action...</u>
Add an asset	Select the asset's check box and click Add to Export .
Remove all assets on the tab	Click Remove from Export .
Remove all assets	Click Clear . Clear is a global command, which clears all the assets selected for export. Remove From Import removes the selected assets from only the selected asset type. It does not remove all the selected assets from the other asset types.

If you want to also export the dependent assets of the selected assets because they do not exist on the target machine, click **Add Dependencies**.

If you want to navigate to another page after selecting assets for export, the state of the selected assets is maintained when you return to the **Export** page.

6. Click **Continue**. In the Export Summary panel, you can click  for each asset type to see the list of assets on the export list. If you want to remove an asset, click .
7. Click **Asset Extraction**, type the name to use for the export file, and click **OK**. Trading Networks generates the file and names it *SolutionName*.zip. The file contains the XML files for the assets. Click **Save** to save the file.

Building the Exported Assets

For detailed information about the topics in this section, including the properties in the *build.properties* file, see the *webMethods Deployer User's Guide*.

To build the exported assets

1. Go to *Software AG_directory* /common/AssetBuildEnvironment/master_build and open the *build.properties* file.
2. Identify the *SolutionName*.zip export file as follows:

To build the assets from...	Take this action...
A single project	Assign the location of the <i>SolutionName</i> .zip file to the <i>build.source.project.dir</i> property.
Multiple projects	Assign the parent directory of all the project directories to the <i>build.source.dir</i> property.

The *build* Ant target builds the .acdl files using all the source files available in the subdirectories of the *build.source.dir* directory. If the *build.source.dir* property is not specified, the *build* Ant target uses the source file specified by the *build.source.project.dir* property to build the .acdl files.

3. Set the *enable.build.TN* property to `true`.
4. Make sure you have assigned the correct values to the *sag.install.dir*, *build.output.dir*, and *build.version* properties.
5. Save the *build.properties* file.
6. Store the export file in the directory specified by the *build.output.dir* property. To do so, go to the *Software AG_directory* /common/AssetBuildEnvironment/master_build directory and run this command:

```
ant build
```

The *build* Ant target uses the properties in the *build.properties* file, validates the exported assets and the dependencies, and generates these files:

File	Contents
<i>SolutionName</i> .acdl	Metadata for the exported assets; that is, a list of the exported assets and the dependencies. This metadata schema is defined by Deployer and is common across the webMethods product suite.
<i>SolutionName</i> .bin	Actual data for the exported assets

Deploy the Assets to Other Servers

Deploy Trading Networks assets to other webMethods servers. For instructions, see the *webMethods Deployer User's Guide*.

Trading Networks logs the deployment activities as activity logs in My webMethods for easy tracking and monitoring of deployment activities.

Migrating Using My webMethods

Exporting Assets Using My webMethods

To export assets using My webMethods

1. In My webMethods: **Administration > Integration > B2B > Asset Management > Export**.
2. From the **TN Server** list, select the Trading Networks Server instance from which to export assets.
3. If you want to export all Trading Networks assets, select the **Select all assets** check box.
4. To select assets by asset type, click the tab whose name corresponds to the asset type.

Note: The Partner Profile tab lists each partner, including your corporation, by corporation name.

5. Do the following in each tab:
 - a. If you want to export all assets of this type, select the **Select all** check box.
 - b. If you want to export only certain assets, type a string to find those assets in the **Keywords** field and click **Search**. Trading Networks searches the same fields it searches when you search for assets (see "[Searching Using Keywords \(Simple Search\)](#)" on page 222). In the search results, do the following to create the list of assets to export:

To do this...	Take this action...
Add an asset	Select the asset's check box and click Add to Export .
Remove all assets on the tab	Click Remove from Export .
Remove all assets	Click Clear . Clear is a global command, which clears all the assets selected for export. Remove From Import removes the selected assets from only the selected asset type. It does not remove all the selected assets from the other asset types.

If you want to also export the dependent assets of the selected assets because they do not exist on the target machine, click **Add Dependencies**.

If you want to navigate to another page after selecting assets for export, the state of the selected assets is maintained when you return to the **Export** page.

- Click **Continue**. In the Export Summary panel, you can click  for each asset type to see the list of assets on the export list. If you want to remove an asset, click .
- Click **Export** and type the name to use for the export file. For **Type**, choose whether to create an XML or binary-encoded file type. You can read XML files and edit assets if necessary before import (for example, the name of a property), but they are a little slower to import than binary files. Binary files cannot be edited, but they offer better performance during the import. Click **OK**. Trading Networks generates the file with the name you specified and the extension .zip. Click **Save** to save the file to the file system.

Importing Assets Using My webMethods

My webMethods imposes a default size limit of 20 MB for files imported to, or published to, the Integration Server that hosts Trading Networks Server. If the size of the XML or binary file you are importing exceeds 20 MB, contact your system administrator to change the Max File Size content value as described in *Administering My webMethods Server*.

To import assets using My webMethods

- In My webMethods: **Administration > Integration > B2B > Asset Management > Import**.
- From the **TN Server** list, select the Trading Networks Server instance from which to export assets.
- In the Upload File area, click **Browse** to locate and upload the .zip file, XML file, or the binary file you exported and then click **Extract Assets**.

4. If you want to import all Trading Networks assets, select the **Select all assets** check box.
5. To select assets by asset type, click the tab whose name corresponds to the asset type and do the following in each tab:
 - a. If you want to import all assets of this type, select the **Select all** check box.
 - b. If you want to import only certain assets, type a string in the **Keywords** field to find those assets and click **Search**. Trading Networks searches the same fields it searches when you search for assets (see "[Searching Using Keywords \(Simple Search\)](#)" on page 222). In the search results, do the following to create the list of assets to import:

<u>To do this...</u>	<u>Take this action...</u>
Add an asset	Select the asset's check box and click Add to Import .
Remove all assets on the tab	Click Remove from Import .
Remove all assets	Click Clear .

If you want to also import the dependent assets of the selected assets because they do not exist on the target machine, click **Add Dependencies**.

If you want to navigate to another page after selecting assets for import, the state of the selected assets is maintained when you return to the **Import** page.

6. Click **Continue**. In the Import Summary panel, you can click  for each asset type to see the list of assets on the import list. If you want to remove an asset, click  .
7. Click **Import**. In the Import Assets dialog box, do the following:
 - a. An asset with the same internal ID or name as an asset in the export file might exist on the target system. If you are importing document types, document attributes, partner profiles, field groups, contact types, or queues, and you want to overwrite such assets, select the **Force Import** check box.
 - b. If you are importing processing rules and you want the rules you are importing to replace all existing processing rules in the target system, select the **Replace Rule List** check box. If you want to replace only matching processing rules and append non-matching processing rules without changing the rule order on the target system, clear the check box. In the case of appending non-matching processing rules, they are added to the end of the existing list on the target.
 - c. Click **OK**.

- Click the **Results** tab to see the assets that were imported successfully and any warnings or errors that occurred during the import.

Migrating Assets Using the `tnexport` and `tnimport` Utilities

Run these utilities from the `Integration Server_directory\instances\instance_name\packages\WmTN\bin` directory.

`tnexport` Utility

Use the `tnexport` utility to export all or specific assets of the selected types. For example, you can export all Trading Networks processing rules or only specific processing rules.

To export...	Specify...
All assets of a specified type	Export options on the command line.
Specified assets	Export options in an XML file.

You can export the assets to the following:

Export to...	Trading Networks does this...
Binary file	<p>Creates a zip file containing the binary export file. When extracted, you can use the contents of this zip file to import using My webMethods or a custom service. You can import using My webMethods, custom service, or the <code>tnimport</code> utility. Notes about binary files:</p> <ul style="list-style-type: none"> Binary output is much smaller than XML output, and can be imported more quickly than XML output. You can export certificates to a binary file. Binary files cannot be read by humans. If you have special characters in asset names, you might see <code>corrupted file</code> errors when the asset is imported. You can prevent these errors by using a binary file.
XML file	<p>Trading Networks creates a zip file containing an XML export file. When extracted, you can use the contents of this zip file to import using My webMethods or a custom service, or the <code>tnimport</code> utility. Notes about XML files:</p> <ul style="list-style-type: none"> XML output is a little slower to import than binary output.

Export to...	Trading Networks does this...
	<ul style="list-style-type: none"> You cannot export certificates to an XML file. You can read XML files and edit assets if necessary before import (for example, the name of a property).
Solutions package	Go to Asset mgmt > export > select asset > Asset Extraction, which has the solution package details.

Specifying Export Options on the Command Line

To specify export options on the command line, go to the *Integration Server_directory*\instances*instance_name* \packages\WmTN\bin directory and run this command:

```
tnexport outfile [-bin binfile|-xml xmlfile] -zipdir [directory] [options]
```

- outfile* is one of the following:

Parameter	Does this...
-bin [binfile]	Exports assets to a binary file. By default, the name of the file is <code>export</code> . You can specify a different name on <i>binfile</i> .
-xml [xmlfile]	Exports assets to an XML file. By default, the name of the file is <code>export</code> . You can specify a different name on <i>xmlfile</i> .

Trading Networks creates a zip file that contains the binary or XML file in the user temp directory. For solution deployment, you need to extract assets to create the ACDL file. This ACDL file is used to deploy the assets by using webMethods Deployer. You can use *zipdir* to specify a different directory. The directory must already exist. If you want to use a directory relative to the one from which you are running the `tnexport` utility, specify a relative path.

- options* is one or more of the following:

Option	Exports...
-acdl	Generate ACDL file.
	Note: If you want to use the <code>-acdl</code> option, list this option first, followed by the other options.
-attribs	All document attributes.
-types	All document types.

Option	Exports...
-rules	All processing rules.
-profiles	All partner profiles.
-extflds	Values for all extended fields in the profiles.
-flddefs	All definitions for field groups and all profile fields in those groups.
-lkup	Definitions for all custom contact types, external ID types, and partner groups.
-queues	All public and private queues.
-tpas	All TPAs.
-securityData	All certificate information for all partner profiles.
	Note: Certificates can only be exported to a binary file (see the <code>-bin</code> option, above).
-dls	All data permission rules defined in My webMethods.
-fp	All general functional permissions data defined in My webMethods.
-archiveSvcs	All archive schedules and deletion schedules.
-All	All Trading Networks assets.

Specifying Export Options in a File

To specify export options in an XML file, create the file, go to the *Integration Server_directory*\instances*instance_name* \packages\WmTN\bin directory, and run this command:

```
tnexport -file full_path_to_file_name
```

You can specify the options below in the XML file. To see a sample file, go to *Integration Server_directory*\instances*instance_name* \packages\WmTN\bin and open the TNExport.xml file.

Option	Value
<code>exportFileName</code>	Name of the binary file, XML file, or Deployer solution package to which to export assets. By default, the name is <code>export</code> . You can specify a different file name on this option.
<code>zipFileDir</code>	By default, the binary file, XML file, or Deployer solution package is created in the user temp directory. You can use this option to specify a different directory. The directory must already exist. If you want to use a directory relative to the one from which you are running the <code>tnexport</code> utility, specify a relative path.
<code>acdl</code>	Whether to export the assets as a Deployer solution package (<code>true</code>) or an XML or binary file (<code>false</code>). If you specify <code>false</code> , you must specify the <code>exportType</code> option.
<code>exportType {bin xml}</code>	If you specified <code>false</code> on the <code>acdl</code> option, whether to save the export data in binary or XML format.
<code>all{true false}</code>	Whether to export all assets, (<code>true</code>) or specific assets (<code>false</code>). If you specify <code>false</code> , you must specify the <code>asset</code> option.
<code>assettypeall ids</code>	<p>If you specified <code>false</code> on the <code>all</code> option, the assets to export. This element contains these attributes and child elements:</p> <ul style="list-style-type: none"> ■ <code>type</code>. This attribute specifies the asset types to export. Valid values are <code>documenttype</code>, <code>documentattribute</code>, <code>processingrule</code>, <code>partner</code>, <code>fieldgroup</code>, <code>fielddefinition</code>, <code>externalidtype</code>, <code>contacttype</code>, <code>binarytype</code>, <code>profilegroup</code>, <code>queue</code>, <code>tpa</code>, <code>dls</code>, <code>functionalpermission</code>, <code>extendedfield</code>, and <code>archiveschedule</code>. ■ <code>all</code>. This attribute indicates whether to export all assets (<code>true</code>) or only assets that match the internal IDs specified in the <code>ids</code> child element (<code>false</code>). ■ <code>ids</code>. This child element specifies the internal IDs associated with the assets to export. For the <code>queue</code> and <code>archive schedule</code> asset types, the internal ID is considered the name of the asset; for all other asset types, the internal ID is the individual asset ID.

Examples

- You want to export all partner profiles, all extended field values in the profiles, and all the extended field definitions to the XML file `ExportedData.xml`. You want the utility to create the zip file containing the XML data file in `\TradingNetworks\`. Use this command:

```
tnexport -xml ExportedData.xml -zipdir \TradingNetworks\ -profiles
-extflds -flddefs
```

- You want to export all Trading Networks assets to the binary file `ExportedData.bin`, and you want the utility to create the zip file containing the `ExportedData.bin` file in the user temp folder. Use this command:

```
tnexport -bin ExportedData.bin -all
```

- You want to export Trading Networks assets based on the options specified in the XML file `ExportOptions.xml`. Use this command:

```
tnexport -file ExportOptions.xml
```

tnimport Utility

Use the `tnimport` utility to import all assets in a specified file. The file you specify must contain assets you previously exported using either the `tnexport` utility or `My webMethods`.

Note: After using the `tnimport` utility, you must reload the `WmTN` package to be able to use the imported data. You must also update each partner's delivery settings password.

To import...	Specify...
All assets of a specified type	Import options on the command line.
Specified assets	Import options in an XML file.

Specifying Import Options on the Command Line

To specify import options on the command line, go to the `Integration Server_directory\instances\instance_name\packages\WmTN\bin` directory and run this command:

```
tnimport infile [options]
```

- `infile` is one of the following:

Parameter	Value
<code>-bin <i>binfile</i></code>	Full path to the binary file that contains the assets to import.
<code>-xml <i>xmlfile</i></code>	Full path to the XML file that contains the assets to import.
<code>-zip <i>zipfile</i></code>	Full path to the zip file that contains the assets to import.

If you specify a relative path, Trading Networks looks for the directory relative to the directory from which you are running the `tnimport` utility.

- `options` is one or more of the following:

Option	Does this...
<code>-db</code>	By default, the <code>tnimport</code> utility only reads and writes information about the data in your file, which is useful for testing purposes. This parameter tells the utility to actually import the assets into the target system.
<code>-force</code>	If document types, document attributes, partner profiles, field groups, field definitions, and TPAs already exist in the target system under the same ID or name as assets you are importing, this parameter tells the <code>tnimport</code> utility to overwrite the target system assets. If you do not want to overwrite such assets, do not specify this option.
<code>-replaceRuleList</code>	If you are importing processing rules, this parameter tells the <code>tnimport</code> utility to replace all existing rules in the target system with the rules you are importing. If you want to replace only matching rules and append non-matching rules without changing the rule order on the target system, do not specify this option.

Specifying Import Options in a File

To specify import options in an XML file, create the file, go to the `Integration Server_directory\instances\instance_name\packages\WmTN\bin` directory, and run this command:

```
tnimport -file full_path_to_file_name
```

You can specify the options below in the XML file. To see a sample file, go to *Integration Server_directory*\instances*instance_name* \packages\WmTN\bin and open the TNImport.xml file.

Option	Value
<code>importDataFileName</code>	Full path to the export file.
<code>importType</code>	Type of file to import. Valid values are bin, xml, or zip.
<code>dbop</code>	By default, the <code>tnimport</code> utility only reads and writes information about the data in your file, which is useful for testing purposes. Specify this option and set it to <code>true</code> to tell the utility to actually import the assets into the target system.
<code>force</code>	If assets already exist in the target system under the same ID or name as assets you are importing, this option tells the <code>tnimport</code> utility to overwrite (<code>true</code>) or not overwrite (<code>false</code>) the target system assets.
<code>replaceRuleList</code>	If you are importing processing rules, this option tells the <code>tnimport</code> utility to do one of the following: <ul style="list-style-type: none"> ■ Replace all existing rules in the target system with the rules you are importing (<code>true</code>). ■ Replace only matching rules and append non-matching rules to the end of the existing list on the target without changing the rule order on the target system (<code>false</code>).
<code>all</code>	Whether to export all assets, (<code>true</code>) or specific assets (<code>false</code>). If you specify <code>false</code> , you must specify the <code>asset</code> option.
<code>assettypeall ids</code>	If you specified <code>false</code> on the <code>all</code> option, the assets to export. This element contains these attributes and child elements: <ul style="list-style-type: none"> ■ <code>type</code>. This attribute specifies the asset types to import. Valid values are <code>documenttype</code>, <code>documentattribute</code>, <code>processingrule</code>, <code>partner</code>, <code>fieldgroup</code>, <code>fielddefinition</code>, <code>externalidtype</code>, <code>contacttype</code>, <code>binarytype</code>, <code>profilegroup</code>, <code>queue</code>, <code>tpa</code>, <code>dls</code>, <code>functionalpermission</code>, <code>extendedfield</code>, and <code>archiveschedule</code>.

Option	Value
	<ul style="list-style-type: none"> <li data-bbox="607 323 1357 428">■ <code>all</code>. This attribute indicates whether to import all assets (<code>true</code>) or only assets that match the internal IDs specified in the <code>ids</code> child element (<code>false</code>). <li data-bbox="607 449 1312 617">■ <code>ids</code>. This child element specifies the internal IDs associated with the assets to export. For the queue and archive schedule asset types, the internal ID is considered the name of the asset; for all other asset types, the internal ID is the individual asset ID.

Examples

- You want to import the assets in the XML file `\TradingNetworks\ExportedData.xml`, and you want to overwrite assets on the target system that have the same ID or name as assets you are importing. Use this command:

```
tnimport -xml \TradingNetworks\ExportedData.xml -db -force
```

- You want to view information about the assets in the `ExportedData.bin` file that resides in the same directory as the `tnimport` utility. You do not want to import the assets into your database. Use this command:

```
tnimport -bin ExportedData.bin
```

- You want to import all assets in the `ExportedData.bin` file that resides in the same directory as the `tnimport` utility, and you want to replace the existing list of processing rules in the target system with the processing rules you are importing. Use this command:

```
tnimport -bin ExportedData.bin -db -replaceRuleList
```

- You want to import assets into the target system based on the options specified in the XML file `ExportOptions.xml` that resides in the same directory as the `tnimport` utility. Use this command:

```
tnimport -file ExportOptions.xml
```

- You want to import the assets in the zip file `\TradingNetworks\ExportedData.zip` into your database. Use this command:

```
tnimport -zip \TradingNetworks\ExportedData.zip -db
```

Note that this command does not overwrite the assets on the target system that have the same ID or name as the assets you are importing. The remaining assets are imported.

15 Archiving and Deleting Documents

■ About Archiving and Deleting Documents	248
■ Setting Up Basic Archive and Deletion	249
■ Setting Up Extended Archive and Deletion	251
■ Setting Up Stored Procedure Archive and Deletion	255
■ Using a Partitioned Database for Archival and Deletion	257

About Archiving and Deleting Documents

You can archive documents in the Trading Networks production database. Trading Networks archives documents by moving them from production tables to archive tables. You can also delete documents from the production tables, the archive tables, or both. You can use either of these types of archive and deletion:

Type	Description
Basic	Trading Networks archives or deletes documents after a specified number of days. When Trading Networks archives a document, it copies the document from the production tables to the archive tables and then removes the document and references to the document from the production tables. When Trading Networks deletes a document, it removes all associated document information from both the production and archive tables.
Extended	Trading Networks archives or deletes documents based on criteria you specify, such as the document ID, sender, receiver, and system status (for example, DONE, RESUBMITTED, or ABORTED). You can choose to archive only selected data in documents. You can choose to delete documents from the production tables, the archive tables, or both. You can also set the <i>batchSize</i> , which is the maximum number of documents that can be archived or deleted in a batch.
Stored Procedure	Trading Networks archives or deletes documents using the stored procedure archive script. Using this mechanism of archiving, the archive table uses the same schemas that are used in the production table. When Trading Networks deletes a document, it removes all the associated document information from both, the production and the archive tables.

The archive tables mirror the production tables, as shown below. Data in the archive tables is identical to the data in the production tables. You can access archived documents from My webMethods; for example, you can view archived documents on the Transactions, Tasks, or Activity Log pages. You can work with archived data just like you do with production data (for example, you can resubmit or reprocess transactions, and stop and restart tasks).

Production Table	Archive Table
BizDoc	ARCHIVE_BizDoc

Production Table	Archive Table
BizDocContent	ARCHIVE_BizDocContent
BizDocAttribute	ARCHIVE_BizDocAttribute
BizDocArrayAttribute	ARCHIVE_BizDocArrayAttribute
BizDocRelationship	ARCHIVE_BizDocRelationship
DeliveryJob	ARCHIVE_DeliveryJob
ActivityLog	ARCHIVE_ActivityLog
BizDocUniqueKeys	ARCHIVE_BizDocUniqueKeys
EDITracking	ARCHIVE_EDITracking

Note: If Trading Networks repeatedly deletes data from an Oracle database, the database tables might become skewed. To resolve this issue, regenerate the statistics of the table.

Setting Up Basic Archive and Deletion

Setting the Number of Days to Retain Documents

You specify the number of days to retain documents in the production tables. Trading Networks archives or deletes documents that are older than the specified number of days.

To specify the number of days to retain documents

1. In My webMethods: **Administration > Integration > B2B Settings > Configure Properties**.
2. Add or edit these properties:

Property	Value
tn.archive.archiveAfterDays	Number of days to retain documents before archive. Specify a value from 0 through 730365. If you do not want to archive documents, specify 0 or do not specify this property.

Property	Value
<code>tn.archive.deleteAfterDays</code>	Number of days to retain documents before deletion. Specify a value from 0 through 730365. If you do not want to delete documents, specify 0 or do not specify this property.

3. Click **Save**.

Scheduling Archive and Deletion

The archive and deletion service is the `wm.tn.archive:archive` built-in service. You schedule when to run this service in Integration Server Administrator, by setting up a user task that executes the service. For complete information about scheduling services for execution, see *webMethods Integration Server Administrator's Guide*.

To schedule archive and deletion

1. Open the Integration Server Administrator that hosts Trading Networks.
2. In the **Server** menu of the navigation area, click **Scheduler**.
3. Click **Create a scheduled task**.
4. Set the **Service Information** parameters as follows:

Parameter	Action
folder:subfolder:service	Specify <code>wm.tn.archive:archive</code> .
Run As User	Specify <code>Administrator</code> .
Persistence	Select the Persist after restart check box. This selection tells Integration Server to maintain this user task if Integration Server is restarted.
Clustering	Specify whether you want the task to run on any available Integration Server in your cluster of Integration Servers. If you do not specify this option, the task runs on the Integration Server that hosts Trading Networks.

5. In the **Schedule Type and Details** section of the screen, do one of the following:
 - To run the service at specified intervals, click **Repeating**. In the **Interval** field, specify the number of seconds to wait between executions. For example, if you want the service to run once every 24 hours, specify 86400. In the **Repeating** field, select the **Repeat from end of invocation** check box. This setting tells Integration

Server to wait for the scheduled task to complete before starting the next, regardless of the **Interval** setting.

Note: The service executes for the first time immediately after you schedule the user task.

- To run the service on specified dates and times, click **Complex Repeating**. In the **Start Date** and **Start Time** fields, specify the date and time for the first execution. In the **Run Mask** parameters, specify the schedule for subsequent executions. If you want the service to run indefinitely, do not specify **End Date** or **End Time**.

6. Click **Save Tasks**.

Setting Up Extended Archive and Deletion

You can set up the extended archive and deletion of documents using the `wm.tn.archive:extendedArchive` built-in service and the Integration Server Scheduler, or using My webMethods.

You must schedule when to run the `wm.tn.archive:extendedArchive` service. To do so, you use Software AG Designer to create a wrapper service that executes the service, then use Integration Server Administrator to set up a user task that executes the wrapper service. For complete information about scheduling services for execution, see *webMethods Integration Server Administrator's Guide*.

Setting Up Extended Archive and Deletion Using Integration Server

Setting Archive and Deletion Criteria

To set archive and deletion criteria

1. Open Software AG Designer.
2. In the WmTN package, open the folder **wm > tn > archive** and double-click the `wm.tn.archive:extendedArchiveservice`.
3. Specify input parameters for the service.

Scheduling Archive and Deletion

To schedule archive and deletion

1. Open Software AG Designer.
2. Create a wrapper service that invokes the `wm.tn.archive:extendedArchive` service.
3. Start Integration Server Administrator.
4. Follow the instructions in "[Scheduling Archive and Deletion](#)" on page 250, but specify the wrapper service for `folder:subfolder:service`.

Setting Up Extended Archive and Deletion Using My webMethods

To set up archive or deletion using My webMethods, you set criteria and a schedule for the archival or deletion. Trading Networks archives or deletes documents in batches.

Setting the Archive and Deletion Criteria

To set the archive and deletion criteria

1. In My webMethods: **Administration > Integration > B2B > Archive Schedules**.
2. Do one of the following:
 - a. To create the schedule by providing new information, click **Add Schedule**.
 - b. To create the schedule by copying an existing schedule and modifying it, click  in the row for the schedule to copy.
3. On the Archive Schedule Details page, provide a name and description for the schedule, and indicate whether to enable or disable the schedule. Do not enable the schedule until you are done creating it.
4. Click the Schedule Options tab and complete these fields:

In this field...	Specify...
Action Type	Whether to schedule archive or deletion.
Datasource	If you chose to schedule deletion, whether to delete from the production tables, the archive tables, or both.
Days	Number of days to retain documents. Trading Networks archives or deletes documents that are older than the number of days you specify.
Batch Size	Maximum number of documents to archive or delete in each batch. The default is 100 documents. If you want to change the default, update the tn.archive.batchSize property.
Wait Time	Number of seconds to wait between two batches of archive or deletion during one scheduled execution. The default time is 15 seconds. If you want to change the default, update the tn.archive.batchBackoffTime property.

In this field...	Specify...
Maximum Transactions Per Schedule	<p>Maximum number of documents to archive or delete during one scheduled execution. This value must be greater than or equal to Batch Size.</p> <p>For example, if you set the Batch Size to 50 and the Maximum Transactions Per Schedule to 1000, Trading Networks performs the archive or deletion task 20 times during that scheduled execution.</p> <p>When the total number of documents for archive or deletion is greater than the maximum number of documents per scheduled execution, Trading Networks tries to archive or delete the additional documents during the next schedule execution.</p> <p>The default value is specified on the tn.archive.maxRows property.</p>

- In the Criteria panel, specify the archive or deletion criteria in the fields below. All of the fields are option; if you do not want to limit archive or deletion based on a criterion, leave that field blank.

In this field...	Specify this criterion...
Sender	Name of the sender whose documents to archive or delete.
Receiver	Name of the receiver whose documents to archive or delete.
Document Type	Name of the document type for documents to archive or delete.
Processing Status	<p>Processing status (for example, Accepted or Done) of documents to archive or delete.</p> <p>For information about processing statuses, see <i>webMethods Trading Networks User's Guide</i>.</p>
User Status	User status of the documents to archive or delete.

- In the Data to Archive or Delete panel, select the document data to archive or delete.
 - If you want to archive or delete all document data (that is, the transactions data), select **Transaction Details (BizDoc)**.
 - If you do not want to archive or delete all document data, select one or more of these options:

Select...	To archive or delete...
Transaction Payload Content (BizDocContent)	Document content for the documents.
Transaction Custom Attribute (BizDocAttribute)	All custom document attributes except the custom array attributes for the documents.
Transaction Custom Array Attribute (BizDocArrayAttribute)	All custom array attributes for the documents.
Related Documents (BizDocRelationship)	All documents related to the documents being archived or deleted.
Delivery Tasks	Delivery tasks associated with the documents.
Activity Logs	Activity logs associated with the documents.
Transaction Unique Key Representation (BizDocUniqueKeys)	Unique keys stored for documents for which duplicates were checked.
EDI Tracking Details	Details available in the EDITracking table (see the <i>webMethods Module for EDI Installation and User's Guide</i>).

7. Click **Save**.

Scheduling Archive and Deletion

To schedule archive and deletion

1. In My webMethods: **Administration > Integration > B2B > Archive Schedules**.
2. Select a schedule.
3. Click the **Schedule Settings** tab.
4. In the **Run Schedule** field, specify when to execute the archive or deletion.

Select...	To...
Run Once	Run the archive or deletion only once, on the date and time you specify.

Select...	To...
Fixed Interval	Run the archive or deletion using a fixed interval (for example, every 5 hours).
Hourly	Run the archive or deletion hourly, at the time you specify.
Daily	Run the archive or deletion daily, at the time you specify.
Weekly	Run the archive or deletion weekly, on the day and time you specify.
Monthly	Run the archive or deletion monthly, on the date, day and time you specify.
Yearly	Run the archive or deletion yearly, on the month, date, day, and time you specify.
Do not overlap task	Whether to wait for one archive or deletion to complete before starting the next, regardless of the schedule.

- Click **Save**.
- Set the status of the schedule by selecting **Active** or **Suspend** from **Status** at the top of the page.

You can also set the status on the **Archive Schedules** page by selecting the checkbox beside the schedule and setting the status to **Active** or **Suspend**. You can do this to change the status of multiple schedules at once.

- Click **Save & Close**.

Setting Up Stored Procedure Archive and Deletion

Scheduling Archive and Deletion

You can use the `wm.tn.archive:archiveByStoredProc` service to archive and delete documents from the Trading Networks database. You can schedule when to run this service in Integration Server Administrator, by setting up a user task that executes the service. For complete information about scheduling services for execution, see *webMethods Integration Server Administrator's Guide*.

To schedule archive and deletion

- Open the Integration Server Administrator that hosts Trading Networks.

2. In the **Server** menu of the navigation area, click **Scheduler**.
3. Click **Create a scheduled task**.
4. Set the **Service Information** parameters as follows:

Parameter	Action
folder:subfolder:service	<p>Specify <code>wm.tn.archive:archiveByStoredProc.</code></p> <p>Click Assign Inputs to specify the following inputs parameters:</p> <ul style="list-style-type: none"> ■ <i>operation</i> ■ <i>afterDays</i> ■ <i>batchCount</i> <p>For more information about the parameters, see <i>webMethods Trading Networks Built-In Services Reference</i>.</p>
Run As User	Specify <code>Administrator</code> .
Persistence	Select the Persist after restart check box. This selection tells Integration Server to maintain this user task if Integration Server is restarted.
Clustering	Specify whether you want the task to run on any available Integration Server in your cluster of Integration Servers. If you do not specify this option, the task runs on the Integration Server that hosts Trading Networks.

5. In the **Schedule Type and Details** section of the screen, do one of the following:
 - To run the service at specified intervals, click **Repeating**. In the **Interval** field, specify the number of seconds to wait between executions. For example, if you want the service to run once every 24 hours, specify 86400. In the **Repeating** field, select the **Repeat from end of invocation** check box. This setting tells Integration Server to wait for the scheduled task to complete before starting the next, regardless of the **Interval** setting.

Note: The service executes for the first time immediately after you schedule the user task.

- To run the service on specified dates and times, click **Complex Repeating**. In the **Start Date** and **Start Time** fields, specify the date and time for the first execution. In

the **Run Mask** parameters, specify the schedule for subsequent executions. If you want the service to run indefinitely, do not specify **End Date** or **End Time**.

6. Click **Save Tasks**.

Setting Up Stored Procedure Archive and Deletion Using Integration Server

Setting the Archive and Deletion Criteria

To set archive and deletion criteria

1. Open Software AG Designer.
2. In the WmTN package, open the folder **wm > tn > archive** and double-click the `wm.tn.archive:archiveByStoredProc` service.
3. Specify input parameters for the service.

Scheduling Archive and Deletion

To schedule archive and deletion

1. Open Software AG Designer.
2. Invoke the `wm.tn.archive:archiveByStoredProc` service.
3. Start Integration Server Administrator.
4. Follow the instructions in "[Scheduling Archive and Deletion](#)" on page 255.

Using a Partitioned Database for Archival and Deletion

Apart from using the Basic and Extended methods for archival and deletion, you can partition your database for archiving and deleting data using the scripts that are located at `Software AG_directory\common\db\scripts\oracle\tradingnetworks\65\partition_support`. Partitioning logically separates one set of data from another, thereby enhancing the process of archiving data from production tables to archive tables. For more information on partitioning your database for archival and deletion, see "[Partitioning Trading Networks Database](#)" on page 260.

Note: For `BizDocUniqueKeys`, `ARCHIVE_BizDocUniqueKeys`, `EDITracking`, and `ARCHIVE_EDITracking` tables, Trading Networks supports only the Basic and Extended methods of archival. For deletion, use `wm.tn.archive:purgeBizdocUniqueKeysData` and `wm.tn.archive:purgeEDITrackingData` services respectively. For more information on these services, see *webMethods Trading Networks Built-In Services Reference*.

16 Database Partitioning

■ Overview	260
■ Partitioning Trading Networks Database	260
■ Deleting Run-time Data from a Partitioned Database	261
■ Dropping a table in a Partitioned Database	261
■ Using a Partitioned Database for Archiving Run-time Data	261
■ Deleting Archived Data from a Partitioned Database	262

Overview

You can partition your Trading Networks database to logically separate one set of data from another. Database partitioning is useful especially when you archive data from production tables to archive tables because it provides the flexibility to perform operations at partition level.

Partitioning Trading Networks Database

You can partition your Trading Networks database using the scripts that are located at *Software AG_directory*\common\db\scripts\oracle\tradingnetworks\65\partition_support. Use the following procedure to partition your database.

Note: Trading Networks supports partitioning only for Oracle database.

To partition database

1. To create partitions, you have to first drop the following tables using `ora_tn_drop_ddl_partition_tables.sql`. If you want to retain any data existing in these tables, you have to manually migrate data to the partitioned tables using Oracle export and import utilities.
 - BIZDOC
 - BIZDOCATTRIBUTE
 - BIZDOCARRAYATTRIBUTE
 - BIZDOCCONTENT
 - BIZDOCRELATIONSHIP
 - DELIVERYJOB
 - ACTIVITYLOG
 - ARCHIVE_BIZDOC
 - ARCHIVE_BIZDOCATTRIBUTE
 - ARCHIVE_BIZDOCARRAYATTRIBUTE
 - ARCHIVE_BIZDOCCONTENT
 - ARCHIVE_BIZDOCRELATIONSHIP
 - ARCHIVE_DELIVERYJOB
 - ARCHIVE_ACTIVITYLOG
2. Edit `ora_tn_create_ddl_partition_tables.sql` and modify the partition names and the partition scheme in the sample script to suit your requirements. The sample script

provided has a monthly partition scheme wherein partitions are created for every month, and data is stored in respective partitions.

3. Run `ora_tn_create_ddl_partition_tables.sql` to create tables with partitions.

Deleting Run-time Data from a Partitioned Database

Use the following procedure to delete run-time data from a partitioned database.

To delete run-time data

1. Go to *Software AG_directory* \common\db\scripts\oracle\tradingnetworks\65\partition_support.
2. Edit `ora_tn_purge_runtimedata.sql` and modify the partition names in the sample script according to the partitions you have created.
3. Run `ora_tn_purge_runtimedata.sql` to delete run-time data from the partitions.

Dropping a table in a Partitioned Database

Use the following procedure to drop tables in a partitioned database.

Note: Do not use Database Component Configurator to drop partitioned tables.

To drop a table

1. Go to *Software AG_directory* \common\db\scripts\oracle\tradingnetworks\65\partition_support.
2. Run `ora_tn_drop_ddl_temp_tables.sql` to drop all the temporary tables.
3. Run `ora_tn_drop_ddl_partition_tables.sql` to drop all the partitioned tables.

Using a Partitioned Database for Archiving Run-time Data

Database partitioning enhances the process of archiving data from production tables to archive tables. Use the following procedure to archive run-time data.

To partition database for archiving data

1. Go to *Software AG_directory* \common\db\scripts\oracle\tradingnetworks\65\partition_support.
2. Run `ora_tn_create_ddl_temp_tables.sql` to create temporary tables.
3. Edit `ora_tn_archive_runtime_data.sql` and modify the partition names in the sample script according to the partitions you have created.

4. Run `ora_tn_archive_runtime_data.sql` to archive run-time data.

Deleting Archived Data from a Partitioned Database

Use the following procedure to delete run-time data from a partitioned database.

To delete archived data

1. Go to *Software AG_directory* \common\db\scripts\oracle\tradingnetworks\65\partition_support.
2. Edit `ora_tn_purge_archivedata.sql` and modify the partition names in the sample script according to the partitions you have created.
3. Run `ora_tn_purge_archivedata.sql` to delete archived data from your partitioned database.

17 Caching Assets and Query Results

- Overview 264
- Asset Caching 264
- Query Results Caching 264
- Viewing or Modifying System Cache Settings 265
- Profile Cache Settings 266
- Profile Summary Cache Settings 267
- Profile ID Cache Settings 269
- Document Type Cache Settings 270
- Document Attribute Cache Settings 272
- TPA Cache Settings 274
- TPA ID Cache Settings 275
- Query Results Cache Settings 277

Overview

Integration Server and other Software AG products use Ehcache in many of their own internal processes. The caches that they use are called *system caches* and they belong to *system cache managers*. Trading Networks has its own system cache manager called `SoftwareAG.IS.TN`. The asset caches and query results cache belong to `SoftwareAG.IS.TN`.

With Ehcache, a cache holds elements that are represented as key-value pairs. The key and its value are both Java objects. An element that is placed in cache is identified by its key. For example, when Trading Networks caches a profile, it uses the `profileId` as the key. You use Integration Server Administrator to display the Trading Networks system cache manager. You are not permitted to delete, reload, or shutdown a system cache manager. You are permitted to modify only certain configuration parameters for a system cache manager and its caches.

Asset Caching

Trading Networks uses the following asset caches:

- **Profile cache.** The Profile cache holds elements for profiles of all corporations who are trading partners in your network.
- **Profile Summary cache.** A profile summary is a subset of profiles. The Profile Summary cache holds elements for all profile summaries.
- **Profile ID cache.** The Profile ID cache holds elements for the profile IDs of all profiles present in your network.
- **Document Type cache.** The Document Type cache holds elements for all document types required for XML and flat files.
- **Document Attribute cache.** The Document Attribute cache holds elements for all document attributes required to process documents.
- **Trading Partner Agreement (TPA) cache.** The TPA cache holds elements for trading partner agreements between various partners in your network.
- **TPA ID cache.** The TPA ID cache holds elements for the TPA IDs of all the trading partner agreements between various partners in your network.

Query Results Caching

Trading Networks stores the results of a database query in a session object to optimize query execution. Trading Networks uses the `tn.query.threshold` property for all queries you perform from My webMethods. The number of rows stored in the session object depends on the value you set for this property. The remaining rows are stored in the

TNQueryResults cache. For best performance, set this property to be the page size Trading Networks uses for query results (that is, 50). The remaining rows are stored in the cache.

Viewing or Modifying System Cache Settings

Use the following procedure to view cache settings or to modify the cache-wide settings for an existing Trading Networks cache.

Note: Be aware that many of the parameters associated with a system cache may be system-defined and cannot be edited in Integration Server Administrator. When you edit a system cache, you can see all of its parameters, but you can edit only certain settings. The parameters that can be modified through the Integration Server UI for TN system are Maximum Elements in Memory, Eternal, Time to Live, Time to Idle, and Maximum Entries Local Disk. The parameters Time to Live and Time to Idle are enabled when the Eternal parameter is not selected.

To view or modify cache settings

1. In Integration Server Administrator: **Settings > Caching > SoftwareAG.IS.TN.**
2. Under **Cache List**, click the cache you want to view or modify.
3. View or modify the settings for the selected cache, as listed below.

<u>To view or modify...</u>	<u>Go to...</u>
Profile cache	Profile Cache Settings
Profile Summary cache	Profile Summary Cache Settings
Profile ID cache	Profile ID Cache Settings
Document Type cache	Document Type Cache Settings
Document Attribute cache	Document Attribute Cache Settings
TPA cache	TPA Cache Settings
TPA ID cache	TPA ID Cache Settings
Query Results cache	Query Results Cache Settings

4. Save your changes.

Note: Trading Networks supports the off-heap storage where the underlying cache is Ehcache in BigMemory and does not support Terracotta Server Array(TSA) which is a distributed cache configuration provided by BigMemory Max.

Profile Cache Settings

The Profile cache holds elements for profiles of all corporations who are trading partners in your network. Set the cache configuration values for the specified fields as follows.

Field	Description	Recommended Value
Maximum Elements in Memory	Total number of entries that this cache can keep in on-heap memory.	50
Eternal	Specifies how long entries are to remain in the cache. When selected, the entries in this cache will never expire once they are put into the cache. Note: When selected, entries in cache do not undergo time based eviction, and Time to Live and Time to Idle values are ignored. When <i>Eternal</i> is not selected, the entries remain in cache depending on the values of Time to Live and Time to Idle .	Selected
Overflow to Disk	Specifies where entries are written when the memory-based portion of on-heap and off-heap cache is filled. When selected, the entries are written to disk when the cache is full. Note: This check box is not selected by default. Note that selecting this option can affect performance of Trading Networks especially where database entries for a particular transaction are distributed across cache and disk. Therefore, Software AG recommends	Not selected

Field	Description	Recommended Value
	<p>disabling this option for use with Trading Networks.</p> <p>The following behavior is noticed when cache is full and when the Overflow to Disk property is not selected:</p> <ul style="list-style-type: none"> ■ For a <i>bounded</i> cache (when Maximum Elements in Memory is limited), the entries are evicted from the cache according to the eviction policy. These entries are later fetched from the Trading Networks database when required. ■ For an <i>unbounded</i> cache (when Maximum Elements in Memory is 0), an <code>OutOfMemoryException</code> is thrown. Increasing the Java heap size can accommodate more entries in cache. 	
Memory Store Eviction Policy	Policy used to remove entries from the cache when the value in Maximum Elements in Memory is reached.	LRU (Least Recently Used)

Profile Summary Cache Settings

A profile summary is a subset of profiles. The Profile Summary cache holds elements for all profile summaries. Set the cache configuration values for the specified fields as follows.

Field	Description	Recommended Value
Maximum Elements in Memory	<p>Total number of entries that this cache can keep in on-heap memory.</p> <p>Note: A value of 0 sets no limit. But Software AG recommends setting this field to 0, to have entries of all profile summaries in cache.</p>	0

Field	Description	Recommended Value
Eternal	<p>Specifies how long entries are to remain in the cache. When selected, the entries in this cache will never expire once they are put into the cache.</p> <p>NoteWhen selected, entries in cache do not undergo time based eviction, and Time to Live and Time to Idle values are ignored. When <i>Eternal</i> is not selected, the entries remain in cache depending on the values of Time to Live and Time to Idle.</p>	Selected
Overflow to Disk	<p>Specifies where entries are written when the memory-based portion of on-heap and off-heap cache is filled. When selected, the entries are written to disk when the cache is full.</p> <p>NoteThis check box is not selected by default. Note that selecting this option can affect performance of Trading Networks especially where database entries for a particular transaction are distributed across cache and disk. Therefore, Software AG recommends disabling this option for use with Trading Networks.</p> <p>The following behavior is noticed when cache is full and when the Overflow to Disk property is not selected:</p> <ul style="list-style-type: none"> ■ For a <i>bounded</i> cache (when Maximum Elements in Memory is limited), the entries are evicted from the cache according to the eviction policy. These entries are later fetched from the Trading Networks database when required. ■ For an <i>unbounded</i> cache (when Maximum Elements in Memory is 0), an <code>OutOfMemoryException</code> is thrown. Increasing the Java heap 	Not selected

Field	Description	Recommended Value
	size can accommodate more entries in cache.	
Memory Store Eviction Policy	Policy used to remove entries from the cache when the value in Maximum Elements in Memory is reached.	LRU (Least Recently Used)

Profile ID Cache Settings

The Profile ID cache holds elements for the profile IDs of all profiles present in your network. Set the cache configuration values for the specified fields as follows.

Field	Description	Recommended Value
Maximum Elements in Memory	Total number of entries that this cache can keep in on-heap memory. Note A value of 0 sets no limit. But Software AG recommends setting this field to 0, to have entries of all profile IDs in cache.	0
Eternal	Specifies how long entries are to remain in the cache. When selected, the entries in this cache will never expire once they are put into the cache. Note When selected, entries in cache do not undergo time based eviction, and Time to Live and Time to Idle values are ignored. When <i>Eternal</i> is not selected, the entries remain in cache depending on the values of Time to Live and Time to Idle .	Selected
Overflow to Disk	Specifies where entries are written when the memory-based portion of on-heap and off-heap cache is filled. When selected, the entries are written to disk when the cache is full.	Not selected

Field	Description	Recommended Value
	<p>Note: This check box is not selected by default. Note that selecting this option can affect performance of Trading Networks especially where database entries for a particular transaction are distributed across cache and disk. Therefore, Software AG recommends disabling this option for use with Trading Networks.</p> <p>The following behavior is noticed when cache is full and when the Overflow to Disk property is not selected:</p> <ul style="list-style-type: none"> ■ For a <i>bounded</i> cache (when Maximum Elements in Memory is limited), the entries are evicted from the cache according to the eviction policy. These entries are later fetched from the Trading Networks database when required. ■ For an <i>unbounded</i> cache (when Maximum Elements in Memory is 0), an <code>OutOfMemoryException</code> is thrown. Increasing the Java heap size can accommodate more entries in cache. 	
Memory Store Eviction Policy	Policy used to remove entries from the cache when the value in Maximum Elements in Memory is reached.	LRU (Least Recently Used)

Document Type Cache Settings

The Document Type cache holds elements for all document types required for XML and flat files. Set the cache configuration values for the specified fields as follows.

Field	Description	Recommended Value
Maximum Elements in Memory	<p>Total number of entries that this cache can keep in on-heap memory.</p> <p>Note: A value of 0 sets no limit. But Software AG recommends setting this field to 0, to have entries of all document types in cache.</p>	0
Eternal	<p>Specifies how long entries are to remain in the cache. When selected, the entries in this cache will never expire once they are put into the cache.</p> <p>Note: When selected, entries in cache do not undergo time based eviction, and Time to Live and Time to Idle values are ignored. When <i>Eternal</i> is not selected, the entries remain in cache depending on the values of Time to Live and Time to Idle.</p>	Selected
Overflow to Disk	<p>Specifies where entries are written when the memory-based portion of on-heap and off-heap cache is filled. When selected, the entries are written to disk when the cache is full.</p> <p>Note: This check box is not selected by default. Note that selecting this option can affect performance of Trading Networks especially where database entries for a particular transaction are distributed across cache and disk. Therefore, Software AG recommends disabling this option for use with Trading Networks.</p> <p>The following behavior is noticed when cache is full and when the Overflow to Disk property is not selected:</p> <ul style="list-style-type: none"> ■ For a <i>bounded</i> cache (when Maximum Elements in Memory is limited), the entries are evicted from the cache according to the 	Not selected

Field	Description	Recommended Value
	<p>eviction policy. These entries are later fetched from the Trading Networks database when required.</p> <ul style="list-style-type: none"> For an <i>unbounded</i> cache (when Maximum Elements in Memory is 0), an <code>OutOfMemoryException</code> is thrown. Increasing the Java heap size can accommodate more entries in cache. 	
Memory Store Eviction Policy	Policy used to remove entries from the cache when the value in Maximum Elements in Memory is reached.	LRU (Least Recently Used)

Document Attribute Cache Settings

The Document Attribute cache holds elements for all document attributes required to process documents. Set the cache configuration values for the specified fields as follows.

Field	Description	Recommended Value
Maximum Elements in Memory	<p>Total number of entries that this cache can keep in on-heap memory.</p> <p>Note: A value of 0 sets no limit. But Software AG recommends setting this field to 0, to have entries of all document attributes in cache.</p>	0
Eternal	<p>Specifies how long entries are to remain in the cache. When selected, the entries in this cache will never expire once they are put into the cache.</p> <p>Note: When selected, entries in cache do not undergo time based eviction, and Time to Live and Time to Idle values are ignored. When <i>Eternal</i> is not selected, the entries remain in cache depending</p>	Selected

Field	Description	Recommended Value
	on the values of Time to Live and Time to Idle .	
Overflow to Disk	<p>Specifies where entries are written when the memory-based portion of on-heap and off-heap cache is filled. When selected, the entries are written to disk when the cache is full.</p> <p>Note: This check box is not selected by default. Note that selecting this option can affect performance of Trading Networks especially where database entries for a particular transaction are distributed across cache and disk. Therefore, Software AG recommends disabling this option for use with Trading Networks.</p> <p>The following behavior is noticed when cache is full and when the Overflow to Disk property is not selected:</p> <ul style="list-style-type: none"> ■ For a <i>bounded</i> cache (when Maximum Elements in Memory is limited), the entries are evicted from the cache according to the eviction policy. These entries are later fetched from the Trading Networks database when required. ■ For an <i>unbounded</i> cache (when Maximum Elements in Memory is 0), an <code>OutOfMemoryException</code> is thrown. Increasing the Java heap size can accommodate more entries in cache. 	Not selected
Memory Store Eviction Policy	Policy used to remove entries from the cache when the value in Maximum Elements in Memory is reached.	LRU (Least Recently Used)

TPA Cache Settings

The TPA cache holds elements for trading partner agreements between various partners in your network. Set the cache configuration values for the specified fields as follows.

Field	Description	Recommended Value
Maximum Elements in Memory	<p>Total number of entries that this cache can keep in on-heap memory.</p> <p>NoteSoftware AG recommends not to set this field to 0 because of the large size of each agreement.</p>	1000
Eternal	<p>Specifies how long entries are to remain in the cache. When selected, the entries in this cache will never expire once they are put into the cache.</p> <p>NoteWhen selected, entries in cache do not undergo time based eviction, and Time to Live and Time to Idle values are ignored. When <i>Eternal</i> is not selected, the entries remain in cache depending on the values of Time to Live and Time to Idle.</p>	Selected
Overflow to Disk	<p>Specifies where entries are written when the memory-based portion of on-heap and off-heap cache is filled. When selected, the entries are written to disk when the cache is full.</p> <p>NoteThis check box is not selected by default. Note that selecting this option can affect performance of Trading Networks especially where database entries for a particular transaction are distributed across cache and disk. Therefore, Software AG recommends disabling this option for use with Trading Networks.</p>	Not selected

Field	Description	Recommended Value
	<p>The following behavior is noticed when cache is full and when the Overflow to Disk property is not selected:</p> <ul style="list-style-type: none"> ■ For a <i>bounded</i> cache (when Maximum Elements in Memory is limited), the entries are evicted from the cache according to the eviction policy. These entries are later fetched from the Trading Networks database when required. ■ For an <i>unbounded</i> cache (when Maximum Elements in Memory is 0), an <code>OutOfMemoryException</code> is thrown. Increasing the Java heap size can accommodate more entries in cache. 	
Memory Store Eviction Policy	Policy used to remove entries from the cache when the value in Maximum Elements in Memory is reached.	FIFO (First In First Out)

TPA ID Cache Settings

The TPA ID cache holds elements for the TPA IDs of all the trading partner agreements between various partners in your network. Set the cache configuration values for the specified fields as follows.

Field	Description	Recommended Value
Maximum Elements in Memory	<p>Total number of entries that this cache can keep in on-heap memory.</p> <p>NoteA value of 0 sets no limit. But Software AG recommends setting this field to 0, to have entries of all TPA IDs in cache.</p>	0
Eternal	Specifies how long entries are to remain in the cache. When selected, the entries in	Selected

Field	Description	Recommended Value
	<p>this cache will never expire once they are put into the cache.</p> <p>Note:When selected, entries in cache do not undergo time based eviction, and Time to Live and Time to Idle values are ignored. When <i>Eternal</i> is not selected, the entries remain in cache depending on the values of Time to Live and Time to Idle.</p>	
Overflow to Disk	<p>Specifies where entries are written when the memory-based portion of on-heap and off-heap cache is filled. When selected, the entries are written to disk when the cache is full.</p> <p>Note:This check box is not selected by default. Note that selecting this option can affect performance of Trading Networks especially where database entries for a particular transaction are distributed across cache and disk. Therefore, Software AG recommends disabling this option for use with Trading Networks.</p> <p>The following behavior is noticed when cache is full and when the Overflow to Disk property is not selected:</p> <ul style="list-style-type: none"> ■ For a <i>bounded</i> cache (when Maximum Elements in Memory is limited), the entries are evicted from the cache according to the eviction policy. These entries are later fetched from the Trading Networks database when required. ■ For an <i>unbounded</i> cache (when Maximum Elements in Memory is 0), an <code>OutOfMemoryException</code> is thrown. Increasing the Java heap size can accommodate more entries in cache. 	Not selected

Field	Description	Recommended Value
Memory Store Eviction Policy	Policy used to remove entries from the cache when the value in Maximum Elements in Memory is reached.	LRU (Least Recently Used)

Query Results Cache Settings

The TNQueryResults cache stores the results of a database query. Set the cache configuration values for the specified fields as follows.

Field	Description	Recommended Value
Maximum Elements in Memory	Total number of entries that this cache can keep in on-heap memory.	500
Eternal	Specifies how long entries are to remain in the cache. Note When selected, entries in cache do not undergo time based eviction, and Time to Live and Time to Idle values are ignored. When <i>Eternal</i> is not selected, the entries remain in cache depending on the values of Time to Live and Time to Idle .	Not Selected
Time to Live	Maximum amount of time (in seconds) that an element can remain in the cache regardless of whether it is accessed or not. Note A value of 0 indicates that elements in this cache do not have a time-to-live expiration. The default value is 0.	1200
Time to Idle	Maximum amount of time (in seconds) that an element can remain in the cache without being accessed.	300

Field	Description	Recommended Value
	<p>Note:A value of 0 indicates that elements in this cache do not have a time-to-idle expiration. The default value is 0.</p>	
<p>Overflow to Disk</p>	<p>Specifies where entries are written when the memory-based portion of on-heap and off-heap cache is filled. When selected, the entries are written to disk when the cache is full.</p> <p>Note:This check box is not selected by default. Note that selecting this option can affect performance of Trading Networks especially where database entries for a particular transaction are distributed across cache and disk. Therefore, Software AG recommends disabling this option for use with Trading Networks.</p> <p>The following behavior is noticed when cache is full and when the Overflow to Disk property is not selected:</p> <ul style="list-style-type: none"> ■ For a <i>bounded</i> cache (when Maximum Elements in Memory is limited), the entries are evicted from the cache according to the eviction policy. These entries are later fetched from the Trading Networks database when required. ■ For an <i>unbounded</i> cache (when Maximum Elements in Memory is 0), an <code>OutOfMemoryException</code> is thrown. Increasing the Java heap size can accommodate more entries in cache. 	<p>Not selected</p>
<p>Memory Store Eviction Policy</p>	<p>Policy used to remove entries from the cache when the value in Maximum Elements in Memory is reached.</p>	<p>LRU (Least Recently Used)</p>

18 Managing Events

■ Overview	280
■ Run-Time Events	280
■ Event Groups	286
■ Configuring Run-time Events	288

Overview

Trading Networks run-time events are organized into several event groups. Grouping of events makes it easy to manage events. To enable Trading Networks to publish run-time events, you have to enable the event properties for the corresponding event groups. This chapter describes the types of run-time events, the event groups, and how to configure the event properties to enable or disable the events.

Run-Time Events

Trading Networks publishes the following run-time events:

- **Alert Email.** Trading Networks can be configured to send alert emails. It publishes this event when it sends an alert email to a contact in the sender's or receiver's profile.

Information published:

- Timestamp of the event
- Internal Document ID
- Email address of the sender or receiver

- **Batch Delivery Complete.** Trading Networks publishes this event when it delivers a batch of documents present in a queue, at a scheduled time.

Information published:

- Timestamp when the batch is delivered
- Queue where the documents are placed

- **Document Attribute Added.** Trading Networks publishes this event when a new attribute is added to a document or an existing attribute in a document is updated.

Information published:

- Timestamp of the event
- Internal Document ID
- Sender ID
- Receiver ID
- Custom attributes (ID and value of the attribute added or updated)

- **Document Persisted.** A document can be processed using a processing rule, or simply saved to the Trading Networks database. Trading Networks publishes this event when it saves a document to its database.

Information published:

-
- Timestamp of the event
 - Internal Document ID
 - Indication if the document is saved to the Trading Networks database
 - **Document Processing Complete.** Trading Networks publishes this event when it completes processing a document.

Information published:

- Internal Document ID
- Timestamp of the event
- Timestamp when the document was last processed
- Type ID for the document type
- Document ID
- Sender ID
- Receiver ID
- Group ID to which the document belongs
- Conversation ID for the document's conversation
- System status of the document. The possible values are:
 - NEW
 - DONE
 - ABORTED
 - DONE W/ERRORS
 - POLLABLE
 - QUEUED
 - ACCEPTED
 - ACCEPTED W/ERRORS
 - RESUBMITTING
 - REPROCESSING
 - RESUBMITTED
 - REPROCESSED
 - RESUBMITTED AND ABORTED
 - REPROCESSED AND ABORTED
 - RESUBMITTED W/ERROR
 - REPROCESSED W/ERROR

- NOT ROUTED
- User status of the document
- Custom attributes
- Indication that the document is saved to the Trading Networks database
- Indication that the document is large
- Original Sender ID
- Original Receiver ID
- Number of times the document has been reprocessed or resubmitted
- Indication that the document has been reprocessed or resubmitted
- **Enqueue Document.** Trading Networks publishes this event when it places a document in a queue, to be delivered later.

Information published:

- Internal Document ID
- Queue where the documents are placed
- Timestamp of the event
- **Immediate Delivery Complete.** Trading Networks publishes this event when it delivers a document directly to the receiving partner. It uses the reliable delivery approach where it makes repeated attempts to deliver the document. This event occurs whether the delivery is a success or a failure.

Information published:

- Task ID
- Internal Document ID
- Server on which the process is running
- Time to wait between retries
- Number of retries
- Retry limit
- Status of the delivery process. The possible values are:
 - DONE
 - FAILED
 - NEW
 - DELIVERING
 - PENDING
 - STOPPED

- QUEUED
- HELD
- Timestamp when the delivery process was created
- Timestamp when the status of the delivery process was last updated
- Indication if the completed delivery process is a success or a failure
- Message indicating the success of a delivery process, or reason for failure
- Time taken to deliver the document in milliseconds
- Fully qualified service name
- **Queue for Polling.** Queue for polling is the default delivery method. Trading Networks publishes this event when it places a document in a queue, saves the document to its database, and sets the document's processing status to POLLABLE. When a partner polls for documents, Trading Networks delivers all documents in the queue for which the partner is the receiver.

Information published:

- Timestamp of the event
- Internal Document ID
- **Relate Documents.** Trading Networks publishes this event when it establishes a one-way relationship between two documents.

Information published:

- Timestamp of the event
- Internal Document IDs of the two documents
- Relationship shared between the documents
- **Reprocess Document.** Trading Networks publishes this event when a document is sent back through processing rules to be reprocessed.

Information published:

- Internal Document ID
- Timestamp of the event
- **Resubmit Document.** Trading Networks publishes this event when a document is resubmitted, to be processed as a new document.

Information published:

- Internal Document ID
- Timestamp of the event
- **Routing Rule Selected.** Trading Networks publishes this event when it matches a routing rule with an incoming document.

Information published:

- Timestamp of the event
- Internal Document ID
- Name of the routing rule
- Post-processing actions. The possible values are:
 - execute.a.service
 - deliver.document.by
 - change.user.status
 - alert.email
 - respond.with
- **Send Document.** Trading Networks publishes this event when it sends a document over a transport protocol.

Information published:

- Timestamp of the event
- Internal Document ID
- Transport Protocol. The possible values are:
 - POLL
 - QUEUE_DELIV
 - RECEIVERS_QUEUE
 - PREFERRED
 - PRIMARY_HTTP
 - PRIMARY_HTTPS
 - PRIMARY_EMAIL
 - PRIMARY_FTP
 - PRIMARY_FTPS
 - SECONDARY_HTTP
 - SECONDARY_HTTPS
 - SECONDARY_EMAIL
 - SECONDARY_FTP
 - SECONDARY_FTPS
 - BATCH_FTP
 - PRIMARY SFTP

- SECONDARY SFTP
- **Service Invoke.** Trading Networks publishes this event when it invokes a service for processing a document.
Information published:
 - Timestamp of the event
 - Internal Document ID
 - Method of service invocation. The possible values are:
 - ASYNC
 - SYNC
 - RELIABLE
 - Fully qualified service name
- **Start Delivery.** Trading Networks publishes this event when it starts to deliver a document or a batch of documents.
Information published:
 - Timestamp of the event
 - Internal Document ID
 - Job ID
 - Indication if the job is held
 - Queue where the documents are placed
- **System Status Changed.** System Status refers to the processing status of a document. Trading Networks publishes this event when it changes the processing status of a document, for example, NEW.
Information published:
 - Timestamp of the event
 - Internal Document ID
 - New system status of the document
- **Task Action.** Trading Networks publishes this event when it performs an action on a delivery task related to a document.
Information published:
 - Job ID
 - Timestamp of the event
 - Action performed. The possible values are:
 - Restart

- Stop
- Reassign
- Remove
- New server to which the job is reassigned
- **User Status Changed.** User Status refers to a user-defined status of a document, for example, Accepted. Trading Networks publishes this event when it changes the user status of a document.

Information published:

- Timestamp of the event
- Internal Document ID
- New user status of the document
- **Validation.** Trading Networks publishes this event when it validates a document against a specified schema.

Information published:

- Timestamp of the event
- Internal Document ID
- Validation status (success or failure)
- Error messages in case of failure
- **Verification.** Trading Networks publishes this event when it verifies the digital signature of a document.

Information published:

- Timestamp of the event
- Internal Document ID
- Verification status (success or failure)
- Error messages in case of failure

Event Groups

The run-time events are grouped into event groups. To control whether Trading Networks publishes a group of events, you set the corresponding event property to `true`.

Event group	Run-time events	Event property
Transaction Events	<ul style="list-style-type: none"> ■ Routing Rule Selected ■ Document Persisted ■ Document Attribute Added ■ System Status Changed ■ User Status Changed ■ Verification ■ Validation ■ Alert Email ■ Queue for Polling 	tn.eda.events.transactions
Transaction Complete Events	Document Processing Complete	tn.eda.events.completedTransactions
Delivery Task Events	<ul style="list-style-type: none"> ■ Enqueued Documents ■ Start Delivery ■ Send Document ■ Immediate Delivery Complete ■ Batch Delivery Complete 	tn.eda.events.deliveryTasks
Service Execution Task Events	Service Invoke	tn.eda.events.serviceExecutionTasks
User Initialized Task Events	<ul style="list-style-type: none"> ■ Reprocess Document ■ Resubmit Document ■ Relate Documents ■ Task Action 	tn.eda.events.userInitiatedTasks

Configuring Run-time Events

To control whether Trading Networks publishes a group of events, you set the corresponding event property to `true`. To disable a group of events, you set the corresponding event property to `false`. By default, all event properties are set to `false`.

To enable run-time events

1. In My webMethods: **Administration > Integration > B2B Settings > Configure Properties**.
2. Enable the event groups by setting the corresponding event properties to `true`.

To enable...	Set...
Transaction Events	<code>tn.eda.events.transactions</code> to <code>true</code>
Transaction Complete Events	<code>tn.eda.events.completedTransactions</code> to <code>true</code>
Delivery Task Events	<code>tn.eda.events.deliveryTasks</code> to <code>true</code>
Service Execution Task Events	<code>tn.eda.events.serviceExecutionTasks</code> to <code>true</code>
User Initialized Task Events	<code>tn.eda.events.userInitiatedTasks</code> to <code>true</code>

3. Save your changes.

19 Working with Dashboards and Charts

■ Overview	290
■ Before You Begin	290
■ Dashboards	291
■ Charts	292

Overview

Trading Networks provides a set of dashboards and charts that give you a graphical view of partners and their transactions on a daily, weekly, and monthly basis. The information you can view includes:

- Transaction volume summary
- Transaction volume and total value trend for a document attribute
- Percentage of successful and failed transactions
- Number of late FA violations

You can also view the top-tier and bottom-tier partners based on transaction volume, and the most common and least common document types exchanged between partners.

The dashboards and the charts display data according to the time zone of My webMethods server. However, for non-standard time zones that have 15 or 45 minutes GMT offset, select the nearest available standard time zone to view the dashboards and charts. For example, for Eucla in Australia (GMT +8:45), Chatham Islands in New Zealand (GMT +12:45), and Kathmandu (GMT +5:45), select the nearest available standard time zone.

Before You Begin

To view data in the form of dashboards and charts, Trading Networks has to subscribe to the run-time events that belong to the Transaction Complete Events and the Transaction Events event groups. For more information on run-time events, see "[Run-Time Event Notifications](#)" on page 61.

To use the dashboards and charts, ensure the following:

- Enable the Transaction Complete Events event group to enable Trading Networks to publish the Document Processing Complete run-time event. For this, set `tn.eda.events.completedTransactions` event property to `true`.

In addition to this, to view the Success and Failure Transaction Summary chart and the Late FA (Functional Acknowledgement) Violations Summary chart, enable the Transaction Events event group to enable Trading Networks to publish all the run-time events that belong to this group. For this, set `tn.eda.events.transactions` event property to `true`.

For more information on the event properties, see "[Event Properties](#)" on page 304.

Note: Late FA Violations Summary chart is available only if you have installed webMethods Module for EDI. For more information on installing webMethods Module for EDI, see *webMethods Module for EDI Installation and User's Guide*.

- Enable Trading Networks to subscribe to the run-time events that belong to the Transaction Complete Events and the Transaction Events event groups. To enable Trading Networks to subscribe to the events, install and configure Software AG Universal Messaging. For more information, see *Installing Software AG Products*.
- Set the `tn.chart.fetchMaxRows` property to a suitable value. By default, `tn.chart.fetchMaxRows=500000`. For more information on the property, see ["Miscellaneous Properties" on page 314](#).

Dashboards

The following list describes the available executive dashboards and the details you can view in each dashboard.

- **Transaction Volume Summary.** A transaction volume summary dashboard that summarizes the following data over a day, week, or a month.
 - Transaction Volume Summary by Sender
 - Transaction Volume Summary by Document Type
 - Transaction Volume Summary by Hour of the Day
- **Cross Partner Metrics.** A cross partner metrics dashboard that summarizes the following data for a selected document attribute of type NUMBER.
 - Volume and Value 7 day and 30 Day Trend
 - Volume and Value by Partners 7 Day Trend
 - Successful and Failed Transactions 7 Day Trend
 - Late FA Violations by Partners 7 Day Trend
- **Partner Metrics.** A dashboard that summarizes the following data for a selected partner and a selected document attribute of type NUMBER.
 - Volume and Value 7 Day and 30 Day Trend
 - Successful and Failed Transactions 7 Day Trend
 - Late FA Violations 7 Day Trend

Viewing Dashboards

Use the following procedure to view dashboards.

Note: By default, My webMethods sets the compatibility mode for Internet Explorer to IE8. To view the dashboards and charts in Internet Explorer, set the compatibility mode to IE10. To change the default Internet Explorer compatibility setting, see *Administering My webMethods Server*.

To view dashboards

1. In My webMethods: **Monitoring > Integration > B2B > Analytics**.
2. Select a dashboard to view data.
3. Depending on the dashboard selected, specify the following criteria.
 - For Transaction Volume Summary dashboard, choose from **Date** to view daily, weekly, or, monthly data.

To view weekly data, select a day that belongs to the week for which you want to view data. For example, if the date selected is December 22, 2014, the data displayed is for the entire week starting from December 21, 2014 to December 27, 2014 (Sunday to Saturday).

To view monthly data, select any day that belongs to the month for which you want to view data. For example, if the date selected is December 22, 2014, the data displayed is for the entire month of December 2014.
 - For Cross Partner Metrics dashboard, choose from **Attribute** and **Date** to view daily, weekly, or, monthly data for the selected document attribute of type NUMBER.
 - For Partner Metrics dashboard, choose from **Date**, **Attribute**, and **Partner** to view daily, weekly, or, monthly data for the selected partner and the selected document attribute of type NUMBER. When you click **Select** to choose a partner, the **Select Partner** pop-up window displays the list of partners. For more information on searching for partners, refer "[Searching for Assets](#)" on page 221.
4. Click  to refresh data for all the individual charts in the dashboard.
5. Click  to export data to a .csv file.

Charts

The following list describes the available charts and the details you can view by drilling down each data point in the chart.

- **Transaction Volume Summary by Sender.** A bar chart that summarizes the transaction volume by sender over a day, week, or month.

Drill down the chart to view:

 - Transaction volume by receivers for the selected sender
 - Transaction volume by document types exchanged between the selected partners (selected sender and receiver)
 - Transaction volume for the selected document type exchanged between the partner pair at various time intervals during the day

- **Transaction Volume Summary by Receiver.** A bar chart that summarizes the transaction volume by receiver over a day, week, or month.

Drill down the chart to view:

- Transaction volume by senders for a receiver
- Transaction volume by document types exchanged between the selected partners (selected sender and receiver)
- Transaction volume for the selected document type exchanged between the partner pair at various time intervals during the day

- **Transaction Volume Summary by Document Type.** A bar chart that summarizes the transaction volume by document types over a day, week, or month.

Drill down the chart to view:

- Transaction volume by senders for a selected document type
- Transaction volume by receivers who receive the document type from the selected sender
- Transaction volume for the document type exchanged between the partner pair at various time intervals during the day

- **Transaction Volume Summary by Hour of the Day.** A bar chart that summarizes the transaction volume by hour of the day over a day, week, or month.

Drill down the chart to view:

- Transaction volume by senders at a specific time interval during the day
- Transaction volume by receivers who receive document types from the selected sender during the chosen time
- Transaction volume by document types exchanged between the selected partner pair during the chosen time

- **Transaction Volume and Total Value Trend Summary.** A chart that summarizes the transaction volume and total value trend for a document attribute of type NUMBER over a day, week, or month.

Drill down the chart to view the transaction volume and total value trend by partners for the selected attribute. The data is sorted in ascending order based on the total value if you select a data point on the total value graph, and transaction volume if you select a data point on the transaction volume graph.

- **Transaction Volume and Total Value Trend Summary by Partners.** A chart that summarizes the transaction volume and total value trend by partners for a document attribute of type NUMBER over a day, week, or month.

Drill down the chart to view the transaction volume and total value trend for the selected attribute and the selected partner.

- **Success and Failure Transaction Summary.** A set of pie charts that summarize the percentage of successful and failed transactions over a day, week, or month.

The successful transactions pie chart in turn shows the percentage of successful transactions whose status is either DONE, REPROCESSED, or RESUBMITTED. Similarly, the failed transactions pie chart in turn shows the percentage of failed transactions whose status is either DONE W/ERRORS, REPROCESSED W/ERRORS, or RESUBMITTED W/ERRORS.

- **Late FA Violations Summary.** A chart that summarizes the late FA violations by partners over a day, week, or, month.

An FA (functional acknowledgement) is a type of transaction set sent by the receiver of an EDI transmission to the sender, acknowledging that the message has been received and its syntax is acceptable. For more information on FAs, refer *webMethods Module for EDI Installation and User's Guide*.

Drill down the chart to view the late FA count for the selected partner.

Note: By default, when the Late FA document attribute is added or updated, Trading Networks publishes the Document Attribute Added runtime event (an event that is a part of the Transaction Events event group). To enable Trading Networks to publish this event for any other document attribute in addition to the Late FA attribute, refer "[tn.eda.events.emittableDocumentAttributeNameFilter](#)" on page 315.

Viewing Charts

Use the following procedure to view charts.

Note: By default, My webMethods sets the compatibility mode for Internet Explorer to IE8. To view the dashboards and charts in Internet Explorer, set the compatibility mode to IE10. To change the default Internet Explorer compatibility setting, see *Administering My webMethods Server*.

To view charts

1. In My webMethods: **Monitoring > Integration > B2B > Analytics**.
2. Select a chart to view data.
3. Depending on chart selected, specify the following criteria.

Note: To view weekly data, select a day that belongs to the week for which you want to view data. For example, if the date selected is December 22, 2014, the data displayed is for the entire week starting from December 21, 2014 to December 27, 2014 (Sunday to Saturday). For monthly data, if the date selected is December 22, 2014, the data displayed is for the entire month of December 2014.

- For Transaction Volume Summary by Sender and Transaction Volume Summary by Receiver charts, choose from **Date** to view daily, weekly, or, monthly data. Click  to choose from **Number of Partners**, and from **Sort By** to sort data

- in ascending or descending order based on total volume. By default, the charts display data for top 10 partners arranged in descending order.
- For Transaction Volume Summary by Document Type chart, choose from **Date** to view daily, weekly, or, monthly data. Click  to choose from **Number of Document Types**, and from **Sort By** to sort data in ascending or descending order based on total volume. By default, the chart displays data for 10 most commonly used document types arranged in descending order.
 - For Transaction Volume Summary by Hour of the Day, choose from **Date** to view daily, weekly, or, monthly data.
 - For Transaction Volume and Total Value Trend Summary chart, choose from **Date** and **Attribute** to view daily, weekly, or, monthly data for the selected document attribute of type NUMBER. The drill down chart, by default, displays top 10 partners for the selected attribute.
 - For Transaction Volume and Total Value Trend Summary by Partners chart, choose from **Date** to view daily, weekly, or, monthly data. Click  to choose from **Attribute**, from **Sort By** to sort data in ascending or descending order based on total value or total volume, and from **Direction** to display data for inbound or outbound transactions, or both.
 - For Success and Failure Transaction Summary chart, choose from **Date** to view daily, weekly, or, monthly data.
 - For Late FA Violations Summary chart, choose from **Date** to view daily, weekly, or, monthly data. Click  to choose from **Number of Partners**, and from **Sort By** to sort data in ascending or descending order based on total number of late FA violations. By default, the chart displays data for top 10 partners with late FA violations arranged in descending order.
4. Click a bar on a bar chart or a data point on a line chart to drill down for more details.
 5. Click  to take you back to the main chart.
 6. Click  to export data to a .csv file.

A Configuration Properties

■ Introduction	298
■ Activity Log Properties	298
■ Cluster Properties	299
■ Database Properties	300
■ Database Query Properties	301
■ Data Level Security Properties	302
■ Debugging Properties	302
■ Document Archiving Properties	302
■ Document Validation Property	303
■ Document Delivery Properties	303
■ Event Properties	304
■ Flat File Property	306
■ Large Document Handling Properties	306
■ Localization Properties	307
■ Performance Properties	308
■ Profile Properties	310
■ Task Properties	311
■ Miscellaneous Properties	314

Introduction

This appendix contains a description of the Trading Networks properties. You can set these properties using one of the following methods:

- Execute the `wm.tn.admin:setProperties` service.
- Specify values for the properties on the **Administration > Integration > B2B Settings > Configure Properties** page in My webMethods.
- Edit the properties file, `properties.cnf`, in a text editor. This file is located in the `Integration Server_directory\instances\instance_name\packages\WmTN\config` directory. If you set properties using this method, you should reload the WmTN package after you save the properties file.

Trading Networks Server uses default values for many of the properties. If a property has a default, it is listed with the description of the property.

Note: Trading Networks maintains a second properties file, called `default_properties.cnf`, in the `Integration Server_directory\instances\instance_name\packages\WmTN\config` directory. This file contains the default values for Trading Networks properties.

Activity Log Properties

`tn.docType.EnableLogDocumentType`

Whether to log auditing information to the activity log when you modify document types from My webMethods (as opposed to modifying document properties in the Trading Networks properties file). Set the property to `true` or `false`. The default is `false`.

`tn.partnerOnboarding.EnableLog`

Whether to log all partner onboarding activity in the Activity Log. If you set the value of `tn.partnerOnboarding.EnableLog` to `true`, it will log all partner onboarding activity in the Activity Log table. If you set the value to `false`, Trading Networks will not log any partner onboarding activities in the Activity Log. Set the property to `true` or `false`. The default is `true`.

`tn.procRule.EnableLogProcessingRule`

Whether to log auditing information to the activity log when you modify processing rules from My webMethods. Set the property to `true` or `false`. The default is `false`.

tn.profile.EnableLogProfileChanges

Whether to log auditing information to the activity log when you modify profiles or data or general functional permissions from My webMethods. Set the property to `true` or `false`. The default is `false`.

tn.queue.EnableLogQueueStartStop

Whether to prevent start and stop messages for queues from being logged into the Activity Log table. If you set the value of this property to `true`, webMethods Trading Networks will log the start and stop messages for all the queues in the Activity Log table. If you set the value to `false`, webMethods Trading Networks will not log the start and stop messages for any queues in the Activity Log table. The default is `true`.

tn.tpa.EnableLogTPA

Whether to log auditing information to the activity log when you modify profiles or data or general functional permissions from My webMethods. Set the property to `true` or `false`. The default is `false`.

Cluster Properties

tn.cluster.notifyCacheChange

Whether a host Integration Server will notify other Integration Servers in the cluster when you modify Trading Networks document types, processing rules, document attributes, profile summaries, or delivery services. Specify `false` or specify `true` and use the [tn.cluster.sync.remoteAliases](#) property to identify the Integration Servers to notify. The default is `true`.

tn.cluster.notifyProfileAddUser

When a user account is created on one Integration Server in a cluster, whether the user account is also created on other servers in the cluster. Specify `false` or specify `true` and use the [tn.cluster.sync.remoteAliases](#) property to identify the Integration Servers to notify. The default is `true`.

tn.cluster.notifyPropertyChange

When you update Trading Networks properties on one Integration Server in a cluster, whether the Trading Networks properties are updated on other Integration Servers in

the cluster. Specify `false` or specify `true` and use the [tn.cluster.sync.remoteAliases](#) property to identify the Integration Servers to notify. The default is `true`.

tn.cluster.sync.remoteAliases

List of remote server aliases for other Integration Servers in the cluster. There is no default; if you do not set this property, Trading Networks cannot notify other Integration Servers when there are changes. Separate server names with a comma.

Database Properties

tn.db.log.level

Controls the activity logs created for every incoming document, thereby reducing the number of messages added to the Activity Log table. Trading Networks logs messages to the Activity Log table according to the value set for this property:

Value	Messages Logged
0	Logs error messages
1	Logs error messages and warnings
2	Logs all messages

The default value is 2.

tn.db.readonly

When Trading Networks points to a read-only database, set this property to `true`. If you set this property to `false`, Integration Server logs an SQL exception in the server log file. The default value is `false`.

tn.store.encoding

The encoding your database is using. When a string is passed to an INSERT or UPDATE statement, if that string is too long to fit into a column, this property indicates where to truncate the string. If you do not set this property, Trading Networks uses the default system encoding.

tn.db.TNLocalDateFormat

Determines the timestamp to use when storing data, to use either the local time zone or GMT. This property applies only to documents that are stored after the property is set. Existing records are not updated.

- `tn.db.TNLocalDateFormat=false` Uses GMT timezone. Default value.
- `tn.db.TNLocalDateFormat=true` Uses the local timezone.

Database Query Properties

tn.query.maxrows

Default number of rows that query services return. The query services are services in the `wm.tn.query` folder. An input variable to the query services allows you to specify a maximum number of rows to return from a query. If you do not specify the number of rows to return when you invoke a query service, the query service uses this property. If a query returns more rows than allowed by this property, the query service silently drops the rows that exceed the number you specify. The default is 0, which indicates no maximum.

tn.query.threshold

Number of rows of query results to store in the session object; storing query results in the session object optimizes query execution. Trading Networks stores the number of rows you specify in the session object on the host Integration Server and uses the caching component to store the remaining rows of query results in Ehcache.

Trading Networks uses this property for all queries you perform from My webMethods (for example, queries for documents, tasks, and activity logs). In addition, this is the default for queries that you perform using the built-in services in the `wm.tn.query` folder.

If you are using Trading Networks in a cluster, do not store query results in the session object; set this property to -1. If you are not using Trading Networks in a cluster, for best performance, set this property to be the page size Trading Networks uses for query results (that is, 50). The default is 50.

Data Level Security Properties

tn.dls.EnableLogDataSetChanges

Whether to log auditing information to the activity log when you modify data permissions or general functional permissions in My webMethods. Set the property to `true` or `false`. The default is `false`.

Debugging Properties

tn.store.logsql

Whether Trading Networks should write each SQL statement to the host Integration Server server log as it executes the statement. Writing the SQL statements to the log will cause the Integration Server to run more slowly. You should only use this feature in a development environment. Set the property to `true` or `false`. The default is `false`.

Note: To view the SQL statements in the Integration Server server log, you must set the detail level for the log to 9 or higher. You specify the level of detail you want to view in the server log using the `-debugIntegration Server` startup option.

Document Archiving Properties

tn.archive.archiveAfterDays

Number of days to retain documents in the production tables. Trading Networks archives documents that are older than the specified number of days. Specify a number from 0 through 730365. If you do not want to archive documents, do not specify this property or set it to 0. There is no default.

tn.archive.batchSize

Number of documents to archive or delete in a batch. If the number of documents available for archive or deletion is less than this property, Trading Networks archives or deletes only the number of documents in the batch. There is no limit for the batch size. The default is 100.

tn.archive.batchBackoffTime

Seconds to pause between two batches of archive or deletion. There is no limit for the pause time. The default is 15.

tn.archive.deleteAfterDays

Number of days to retain documents in the production tables. Trading Networks deletes documents that are older than the specified number of days. Specify a number from 0 through 730365. If you do not want to delete documents, do not specify this property or set it to 0. There is no default.

tn.archive.maxRows

Maximum number of documents to archive or delete in each scheduled archive or deletion. The this property must be greater than or equal to the `tn.archive.batchSize` property. The default is -1, which indicates no limit to the maximum number of documents that can be archived or deleted in a scheduled archive or deletion.

Document Validation Property

tn.doc.validate.max_errs

Maximum number of errors Trading Networks can return when errors occur during XML document validation. The default is 10.

Document Delivery Properties

tn.delivery.ftpTransferType

Transfer type for FTP delivery of documents. Specify `active` or `passive`. The default is `passive`.

tn.mail.from

The From e-mail address to use when you deliver documents using an E-mail immediate delivery method. There is no default.

tn.mail.subject

Subject line to use when you deliver documents using an E-mail immediate delivery method. There is no default.

tn.delivery.active.transfer.server.aliases

Aliases for ActiveTransfer Server on remote Integration Server instances. The aliases are necessary to deliver documents to ActiveTransfer Server that are installed on remote Integration Server instances and also provide notifications when a partner profile is added, updated, or deleted. There is no default.

Event Properties

tn.eda.events.transactions

Controls whether Trading Networks publishes the run-time events that belong to the Transaction Events group. When the `tn.eda.events.transactions` parameter is set to `true`, Trading Networks publishes the events listed below. When the parameter is set to `false`, these events are not published. The default is `false`. The run-time events that fall into this category are:

- Routing Rule Selected
- Document Persisted
- Document Attribute Added
- System Status Changed
- User Status Changed
- Verification
- Validation
- Alert Email
- Queue for Polling

tn.eda.events.completedTransactions

Controls whether Trading Networks publishes the run-time events that belong to the Transaction Complete Events group. When the `tn.eda.events.completedTransactions` parameter is set to `true`, Trading Networks publishes the Document Processing Complete event. When the parameter is set to `false`, the event is not published. The default is `false`.

tn.eda.events.deliveryTasks

Controls whether Trading Networks publishes the run-time events that belong to the Delivery Task Events group. When the `tn.eda.events.deliveryTasks` parameter is set to `true`, Trading Networks publishes the events listed below. When the parameter is set to `false`, these events are not published. The default is `false`. The run-time events that fall into this category are:

- Enqueued Document
- Start Delivery
- Send Document
- Immediate Delivery Complete
- Batch Delivery Complete

tn.eda.events.serviceExecutionTasks

Controls whether Trading Networks publishes the run-time events that belong to the Service Execution Task Events group. When the `tn.eda.events.serviceExecutionTasks` parameter is set to `true`, Trading Networks publishes the Service Invoke event. When the parameter is set to `false`, the event is not published. The default is `false`.

tn.eda.events.userInitiatedTasks

Controls whether Trading Networks publishes the run-time events that belong to the User Initialized Task Events group. When the `tn.eda.events.userInitiatedTasks` parameter is set to `true`, Trading Networks publishes the events listed below. When the parameter is set to `false`, these events are not published. The default is `false`. The run-time events that fall into this category are:

- Reprocess Document
- Resubmit Document
- Relate Documents
- Task Action

Flat File Property

tn.ff.contenttypes

Content types to associate with flat file documents. Separate content types with a comma. If you do not include this property in the properties.cnf file, documents with the content type text/plain are treated as flat files.

You cannot register these content types as flat files:

- application/x-www-form-urlencoded
- text/xml
- application/x-wmrpc2
- application/x-x509v3-bin
- application/x-wmrpc
- application/x-wmrpc-bin
- application/x-wmidatabin
- application/wm-soap
- application/soap
- application/soap+xml

Large Document Handling Properties

tn.BigDocThreshold

Threshold size at which to consider a document large.

Value	Meaning
Any negative integer ($-n$)	Consider no documents large.
0	Consider all documents large.
Positive integer	Number of bytes over which to consider a document large.

The default is -1 .

tn.xml.xqlThreshold

Maximum number of bytes to read into memory from the start of an XML document to perform the XQL queries that are defined in an XML document type. The XQL queries can be identifying queries or queries to extract attributes. In your document types, if you have XQL queries that point to nodes beyond the number of bytes you specify, the queries will fail.

Specify a positive integer greater than 1000. If you specify an integer less than 1000, Trading Networks reads 1000 bytes. If you do not specify this property, Trading Networks uses [tn.BigDocThreshold](#).

Localization Properties

Use the localization properties to change the Integration Server's locale. When you change the locale, Trading Networks displays error messages and screen labels using the language associated with the locale. To display the messages and screen labels in the language, Trading Networks must have the appropriate resource bundle. These resource bundles reside in the `tncore.zip` file, which was copied to your system during installation. If you change the locale but the resource bundle does not exist for that locale, Trading Networks tries to find the best match. For information about how Trading Networks selects the resource bundle that matches best, see the Oracle documentation for the `java.util.ResourceBundle` class. By default, Trading Networks uses the locale of the host machine.

To change the locale you must set, at a minimum, the [tn.locale.language](#) and [tn.locale.country](#) properties. If you do not specify both of these properties, Trading Networks continues to use the default locale (that is, the locale of the host machine).

tn.locale.language

Two-letter ISO 639 language code for the language to use for messages and screen labels.

tn.locale.country

Two-letter ISO 3166 country code for the country associated with the language.

tn.locale.variant

Variant code that further distinguishes the locale, if necessary. Typically you use the variant code if there is a difference based on the platform on which Integration Server is running. If you need to specify more than one variant code, list the most important codes first, and separate them with an underscore. For more information about variant codes, see the Oracle documentation for the `java.util.Locale` class.

Performance Properties

tn.clean.routePipeline

Whether to remove all data from the pipeline except *bizdoc*, *rule*, *TN_parms*, and *\$tnReprocess* when a document is routed. Removing pipeline contents can help prevent the pipeline from growing very large, which could improve system performance. Set the property to `true` or `false`. The default is `false`.

tn.db.fetchMaxRows

Maximum number of result rows to return for each query. To return all rows, set the property to 0. The default is 10000.

tn.db.sortDocTimestamp

Whether to sort result rows based on timestamp in descending order. Set the property to `true` or `false`. The default is `true`.

tn.doc.saveAsDone

When Trading Networks receives a document, it saves information about the resulting BizDocEnvelope for that document to its database and sets the status of the BizDocEnvelope. Based on this property, Trading Networks sets the status of the BizDocEnvelope to one of the following:

Status	Description	Set property to...
NEW	Trading Networks has received the document but has not yet processed it. This is the default.	<code>false</code>
DONE	Typically, the DONE status means that Trading Networks has completed processing the document. However, you can use this property to have Trading Networks immediately set the status to DONE. This saves a database update per document if Trading Networks completes processing of the document normally with no errors.	<code>true</code>

Important! Integration Server ends unexpectedly while processing a document, upon it restarts, because the status is already DONE, processing of the

Status	Description	Set property to...
	document will appear to be complete. There is a risk that the document might not be processed completely.	

tn.receive.clearTNOjects

Whether to drop pipeline variables that represent Trading Networks objects after the `wm.tn:receive` entry point service finishes executing. Dropping these objects reduces the time it takes to process each document and improves overall performance because the content handler does not have to format the objects to return to the client. You specify the pipeline variables to drop using the `tn.receive.clearKeys` property. Set the property to `true` or `false`. The default is `true`.

tn.receive.clearKeys

Pipeline variables that represent Trading Networks objects to drop from the pipeline when the `tn.receive.clearTNOjects` property is set to `true`. Separate the variables with a command. Your list can contain these objects:

- `node`
- `bizdoc`
- `sender`
- `receiver`
- `flags`
- `ffdata`
- `TN_parms`

If you specify a pipeline variable other than the ones listed above, Trading Networks ignores it, and Trading Networks does not drop the variable. The default is null (no variables appear in the pipeline).

tn.store.pool.stmt

Enables caching of JDBC statements. Set the property to `true` or `false`. The default is `true`. When you set this property to `true`, all PreparedStatements are cached in memory to speed up subsequent access to data from database. When you set this property to `false`, the PreparedStatements are not cached in memory. Instead, they are created every time to access data.

tn.resubmit.return.bizdocs

Whether the `wm.tn.doc:resubmit` and `wm.tn.doc:resubmits` services should return BizDocEnvelopes in the pipeline. You use these services to resubmit documents that

are in the Trading Networks database for processing. When you resubmit a document, Trading Networks has a BizDocEnvelope for the original document that was in the database and forms an additional BizDocEnvelope for the new instance of the document that is being resubmitted. This property controls whether the `wm.tn.doc:resubmit` and `wm.tn.doc:resubmits` built-in services leave these BizDocEnvelopes in the pipeline when returning to the caller. Performance is improved if the BizDocEnvelopes are not returned in the pipeline. Set the property to `true` or `false`. The default is `false`.

tn.tpa.cacheSize

Maximum number of trading partner agreements (TPAs) in Agreed status to cache in memory. Caching these TPAs speeds up the run-time retrieval of TPAs; however, it also increases the amount of memory used by Integration Server. The default is 1000.

tn.xml.cacheXQLQueryResults

Whether to save and reuse the results of identifying XQL queries executed against an XML document during document recognition. Several of your XML document types might use the same XQL queries, so setting this property can improve performance. Set the property to `true` or `false`. The default is `true`.

Profile Properties

tn.default.idType

External ID type (for example, D-U-N-S® number) to initially display when creating or editing a profile in My webMethods. The default is `DUNS`.

tn.export.deleted.profiles

Displays deleted profiles in the Export Data dialog box. Set this property to `true` or `false`. The default is `false`.

tn.api.gateway.server.alias

Ensures that Trading Networks has the information about the availability of a webMethods API Gateway instance it can connect with.

Task Properties

tn.task.attempt.notPersisted

For you to use reliable delivery for a document, the document must be saved to the Trading Networks database. If reliable delivery is specified for a document and the document has not been saved, this property specifies whether to attempt a single asynchronous execution of the delivery service. Set this property to `true` or `false`. The default is `true`.

tn.task.dbupdate.enableNotifications

Whether Trading Networks should publish an IS document (that is, raise an Integration Server event) when it attempts to update its database for a task after a database failure occurs. If you want Trading Networks to publish an IS document every time it retries to update the database, whether the attempt is successful or failed, set the property to `true`. If you want Trading Networks to not publish an IS document when it tries to update the database after a database failure, set the property to `false`. The default is `false`.

The IS document that Trading Networks publishes is in the format defined by the `wm.tn.rec:TaskDbUpdate` IS document type. You can use Software AG Designer to define a trigger that subscribes to the IS document and takes appropriate actions (see the *Publish-Subscribe Developer's Guide*).

tn.task.dbupdate.retryEnabled

Whether Trading Networks should try to update its database after it encounters a database failure when it tries to update the information for a task (for example, because it lost connectivity to the database). In this case, the information that Trading Networks is maintaining in memory for the task will not match the information stored in the database. If Integration Server is shut down, the task information will be incorrect upon restart. Set this property to `true` or `false`. If you set the property to `true`, also specify the [tn.task.dbupdate.retrySweepTime](#) property. The default is `true`.

tn.task.dbupdate.retrySweepTime

Number of seconds between attempts to retry updating the database information for a task. Trading Networks uses this property when the [tn.task.dbupdate.retryEnabled](#) property is set to `true`. The default is 300.

tn.task.init.groupSize

Number of tasks to process at one time during Integration Server startup. Trading Networks uses a background thread during startup to read the task data from the database and create the tasks in groups. The number of tasks in a group is the setting you specify for this property. After creating a group of tasks, Trading Networks pauses for the amount of time specified by the [tn.task.init.pauseLength](#) property. Pausing between groups makes CPU cycles available to other Integration Server functions that are initializing during startup. This improves performance when starting Trading Networks when there are many pending tasks. The default is 100.

tn.task.init.pauseLength

Number of seconds to pause between processing groups of tasks during Integration Server startup. Trading Networks uses a background thread during startup to read the task data from its database and create the tasks in groups. The number of tasks in a group is the setting you specify for the [tn.task.init.groupSize](#) property. After creating a group of tasks, Trading Networks pauses for the amount of time specified by this property. Pausing between groups makes CPU cycles available to other Integration Server functions that are initializing during startup. This improves startup performance when starting Trading Networks when there are many pending tasks. The default is 10.

tn.task.maxRetries

Number of times to try to re-execute a failed service specified on the Execute a Service processing action when using reliable execution. Specify a positive integer. The default is 5.

tn.task.ttw

Number of milliseconds to wait before making the first attempt to re-execute a failed service specified on the Execute a Service processing action when using reliable execution. Set this property to a positive integer. The default is 300.

tn.task.retryFactor

Factor to use to determine how long to wait before making the second and subsequent attempts to re-execute a failed service specified on the Execute a Service processing action when using reliable execution. Trading Networks calculates the time to wait by multiplying the last wait time by this property. Set this property to a positive integer. The default is 1.

tn.task.notifyFailure

Whether Trading Networks should publish an IS document (that is, raise an Integration Server event) when a task fails. Set this property to `true` or `false`. The default is `false`.

The IS document that Trading Networks publishes is in the format defined by the `wm.tn.rec:TaskFailure` IS document type. You can use Software AG Designer to define a trigger that subscribes to the IS document and takes appropriate actions (see the *Publish-Subscribe Developer's Guide*).

tn.task.pauseOnMaxThreads

Number of milliseconds to pause when the task manager runs out of threads. You control the number of threads that the task manager is allocated using the `tn.task.pauseOnMaxThreads` property. If the number of threads specified by the `tn.task.threadpool.pct` property are in use and Trading Networks has more tasks to process, Trading Networks pauses the number of milliseconds specified by this property to allow some of the working threads to complete and become available. The default is 500.

tn.task.sweepTime

Number of seconds the task manager thread remains idle before checking for tasks it needs to perform (for example, a task to deliver a document or execute a service). The thread becomes idle when all tasks are completed, failed, or in the wait period between attempts. The thread waits the number of seconds you specify with this property before waking up to check for tasks it needs to process. The tasks that the task manager tries to process when it wakes up are those tasks that failed on their last attempt and have not yet reached their retry limit.

The minimum allowable setting for this property is 1 second. If you try to set this property lower than 1 second, the task manager remains idle for 1 second before checking for tasks it needs to perform. The default is 300.

Note: The task manager wakes up immediately when a new task arrives.

This property can affect how long to wait between attempts to retry a task. Typically, the wait between retries is governed by one of the following:

Task	Governed by...
Delivery	The Wait between retries specified in the delivery settings of the receiving partner's profile.

Task	Governed by...
Service execution	The <code>tn.task.ttw</code> property.

However, Trading Networks could wait as long as this property specifies. If the task manager thread is idle when the wait between retries expires, the task manager thread will not wake up to process the task until the `tn.task.sweepTime` expires.

tn.task.threadpool.pct

Percentage of the Integration Server thread pool the task manager can use to process tasks simultaneously. This is to prevent the task manager from consuming all Integration Server threads when the task manager has a large number of tasks to process. Set this property from `.03` through `.9`. If you set the property below the range, the task manager ignores the setting and uses `.03`. If you set the property greater than the range, the task manager ignores the setting and uses `.9`. No matter how small you set the property, the task manager is guaranteed to always have at least one thread with which to execute tasks. The default is `.5`.

tn.task.orderingMethod

Enables task manager to order the pending jobs according to the task creation time or document creation time. The default value is `default`. When you set this property to `default`, the task manager orders the pending jobs according to the task creation time. When `tn.task.orderingMethod= BizDocCreationTime`, the task manager orders the pending jobs according to the document creation time. `BizDocCreationTime` has the same value as the `DocTimestamp` variable of the `BizDocEnvelope`. For more information on `BizDocEnvelope`, see *webMethods Trading Networks Built-In Services Reference*.

Miscellaneous Properties

tn.chart.fetchMaxRows

Whether to display data in the form of dashboards and charts. Set a suitable value for this property. When the number of records in the dashboard tables are less than or equal to the value set, My webMethods displays the data. When the number of records is greater than the value set, My webMethods throws an error when you access the **Analytics** link. By default, `tn.chart.fetchMaxRows = 500000`.

tn.check.ambiguous.docs

Whether to search through all document types to find a single matching document type for an inbound document, or to use the first match. To find a single document type, set

the property to `true`; if Trading Networks finds multiple document types that match it considers the document an unknown document because it does not know which document type to use. To use the first match, set the property to `false`. The default is `true`.

tn.checkUser.ignoreCase

Whether to consider case when mapping partners to the user that initiated a document transaction. To not consider case when evaluating sender IDs, set the property to `true`. To consider case, set the property to `false`. The default is `false`.

Note: This parameter applies only when Trading Networks uses either the Oracle or the Microsoft SQL Server database. Trading Networks uses the default `false` for other databases.

tn.decimalFormatPattern

To accommodate more than two digits after the decimal point. This property specifies the pattern that Trading Networks should use to format decimal numbers. For example, if you want Trading Networks to retain up to six digits after the decimal point, specify `tn.decimalFormatPattern=##.#####`. The pattern can contain any of the following symbols:

- # Indicates a digit. Leading or trailing zeroes are omitted.
- 0 Indicates a digit. Leading or trailing zeroes are included.
- . Indicates a placeholder for a decimal separator.
- , Indicates a placeholder for a grouping separator.

tn.eda.events.emittableDocumentAttributeNameFilter

Determines the set of document attributes for which Trading Networks will publish the Document Attribute Added run-time event. By default, Trading Networks publishes this event when Late FA document attribute is added or updated. To enable Trading Networks to publish this event for any other document attribute apart from Late FA attribute, set this property to a fully qualified custom java class. This class should include the set of document attributes apart from Late FA attribute, that will enable Trading Networks to publish the event when the specified set of document attributes are added or updated. For example, `tn.eda.events.emittableDocumentAttributeNameFilter = com.softwareag.app.tn.policies.eda.EmittableAttributeNameFilter`. The `EmittableAttributeNameFilter` custom class will filter all the document attributes and enable Trading Networks to publish the event only for the specified document attributes.

The custom class should implement the `com.softwareag.app.tn.policies.IInvoker<Boolean, String>` interface, and

override the `invoke (final String attrname)` method. The `attrname` parameter corresponds to the document attribute, and the output of the method is a Boolean value. If the value is `true`, Trading Networks will publish the event for the document attribute specified. If the value is `false`, Trading Networks will not publish the event for the document attribute specified.

tn.ism.stats.saveFrequency

Frequency (in seconds) to save document processing statistics to the database. Document statistics are viewed in My webMethods. The default is 30.

tn.mail.onuserfail

Whether to send an e-mail message to the Integration Server Administrator to alert to these errors:

- Sender of the document is not the same as the current user.
- There is no profile for current user.
- Trading Networks could not retrieve the profile summary for the sender.
- Sender's profile is Inactive, which does not allow for sending documents.
- Sender's profile has been deleted.
- Internal exception prevented processing from completing.

Set the property to `true` or `false`. The default is `false`.

tn.portal.rule.preventValidation

Whether to validate a processing rule name. Set this property to `true` or `false`. If the value is `true`, Trading Networks does not validate the processing rule name (allows special characters). If the value is `false`, Trading Networks validates the processing rule name and prompts you to provide a valid name.

tn.portal.savetodisk.chunksize

Configure the chunk size used to save transaction content on the client side by using this property in the `properties.cnf` file and set the chunk size range from 1024 bytes(1Kb) to 104857600 bytes (100Mb). The default is 100.

tn.recursion.max

Maximum number of successive times to use a processing rule for a specific document. Trading Networks recursively uses a processing rule if the processing rule uses the Execute a Service processing action to synchronously invoke a service that calls the `wm.tn:submit` service to submit the same BizDocEnvelope to Trading Networks, and

Trading Networks selects the same processing rule again. If Trading Networks tries to recursively use a processing rule more times than the number you specify, Trading Networks aborts the document processing. Set this property to a positive integer. The default is 0, which means recursive use of processing rules is not allowed.

tn.server.seq

Trading Networks uses this property when generating internal IDs. To guarantee that the internal IDs are unique across all Integration Servers, IDs consist of the hash code of the host name, the server sequence number, the system timestamp, and the counter. Set this property from 0 through 31. There is no default.

The Integration Server sequence number is the `tn.server.seq` property. Use this property only if you are running more than one Trading Networks instance on the same machine. In this situation, give each instance a unique `tn.server.seq`.

If you run multiple Trading Networks instances on one machine and do not set this property, Trading Networks might generate duplicate internal IDs. If you supply an invalid value, Trading Networks ignores the setting and uses 0 for the server sequence number portion of generated internal IDs, in which case, again, Trading Networks might generate duplicate internal IDs.

tn.transport.user

User account to use to invoke a service that Trading Networks is executing reliably due to a processing action in a processing rule. This affects the user account used to invoke the built-in delivery services (which are in the `wm.tn.transport` folder), custom delivery services that you create, and services invoked by service execution tasks. To have Trading Networks use the currently connected user to invoke services, do not set the `tn.transport.user` property. By default, Trading Networks will use the currently connected user to invoke services. This is the user that submitted the document to Trading Networks, unless you are resubmitting or reprocessing a document. When you resubmit or reprocess a document, the current user is the user issuing the resubmit or reprocess.

If you want to use a user account other than the currently connected user's account, use this property to specify the user name. If you want an Administrator account to execute services from processing rules, specify `Administrator`. Some Integration Server Administrators might delete the built-in Administrator account and create a new one with a different name. If `tn.transport.user` is `Administrator`, Trading Networks will use an account from the Administrators group, though it might not be the built-in Administrator account.

Important: Make sure the ACL you use to protect the service grants Execute access to the user you specify.

B Server Configuration Parameters

■ Introduction	320
■ Server Configuration Parameters	320

Introduction

This appendix contains a description of the Trading Networks properties you can specify in the server configuration file (server.cnf), which is located in the *Integration Server_directory\instances\instance_name\config* directory.

Trading Networks Server uses default values for many of the properties. If a property has a default, it is listed with the description of the property.

Server Configuration Parameters

watt.WmTN.aclMap

Indicates whether Trading Networks maps the ACLs to the Trading Networks public services. During the first startup, Trading Networks maps the ACLs to the Trading Networks public services. After the mapping is done, Trading Networks adds the `watt.WmTN.aclMap` property to the Integration Server server.cnf file, and sets the value of this property to the current build version of Trading Networks. This property does not exist in the server.cnf file by default. During subsequent startups, if you want to force Trading Networks to map the ACLs again, remove the `watt.WmTN.aclMap` property from the server.cnf file and restart Trading Networks.

watt.WmTN.mwsMap

Indicates whether Trading Networks adds My webMethods roles related to Trading Networks to their respective ACLs. During the first startup, Trading Networks adds the roles to their respective ACLs. After adding the roles, Trading Networks adds the `watt.WmTN.mwsMap` property to the Integration Server server.cnf file, and sets the value of this property to the current build version of Trading Networks. This property does not exist in the server.cnf file by default. During subsequent startups, if you want to force Trading Networks to add the roles to the ACLs again, remove the `watt.WmTN.mwsMap` property from the server.cnf file and restart Trading Networks.

C XQL Reference

■ XQL Reference Information	322
■ Sample XML File	322
■ Example XQL Queries	323
■ Extracting Custom Attributes From XML Documents	324

XQL Reference Information

Use the following link for more information about the W3C, World Wide Web Consortium reference information for XQL: <http://www.w3.org/TandS/QL/QL98/pp/xml.html>.

Sample XML File

The example XQL queries in the next section reference the following XML file.

```
<?xml version='1.0'?>
<!-- This file represents a fragment of a book store database. -->
<bookstore specialty='novel'>
  <book style='autobiography'>
    <title>Seven Years in Trenton</title>
    <author>
      <first-name>Joe</first-name>
      <last-name>Bob</last-name>
      <award>Trenton Literary Review Honorable Mention</award>
    </author>
    <price>12</price>
  </book>
  <book style='textbook'>
    <title>History of Trenton</title>
    <author>
      <first-name>Mary</first-name>
      <last-name>Bob</last-name>
      <publication>
        Selected Short Stories of
        <first-name>Mary</first-name> <last-name>Bob</last-name>
      </publication>
    </author>
    <price>55</price>
  </book>
  <magazine style='glossy' frequency='monthly'>
    <title>Tracking Trenton</title>
    <price>2.50</price>
    <subscription price='24' per='year' />
  </magazine>
  <book style='novel' id='myfave'>
    <title>Trenton Today, Trenton Tomorrow</title>
    <author>
      <first-name>Toni</first-name>
      <last-name>Bob</last-name>
      <degree from='Trenton U'>B.A.</degree>
      <degree from='Harvard'>Ph.D.</degree>
      <award>Pulitzer</award>
      <publication>Still in Trenton</publication>
      <publication>Trenton Forever</publication>
    </author>
    <price intl='canada' exchange='0.7'>6.50</price>
    <excerpt>
      <p>It was a dark and stormy night.</p>
      <p>But then all nights in Trenton seem dark and
      stormy to someone who has gone through what
      <emph>I</emph> have.</p>
```

```

    <definition-list>
      <term>Trenton</term>
      <definition>misery</definition>
    </definition-list>
  </excerpt>
</book>
<my:book style='leather' price='29.50' xmlns:my='http://www.placeholder
-name-here.com/schema/'>
  <my:title>Who's Who in Trenton</my:title>
  <my:author>Robert Bob</my:author>
</my:book>
</bookstore>

```

Example XQL Queries

The following lists sample XQL queries.

/bookstore

Find the root element (`bookstore`) of the document

//author

Find all `author` elements anywhere within the current document.

//author/first-name

Find all `first-name` elements within an `author` element anywhere in the current document.

bookstore/*/title

Find all `title` elements that are grandchildren of `bookstore` elements.

//@style

Find all `style` attributes anywhere within the current document.

bookstore/book/price/@exchange

Find the `exchange` attributes of `price` elements that are contained inside `book` elements that are contained in a `bookstore` element.

//book[excerpt]

Anywhere within the current document, find all `book` elements containing at least one child `excerpt` element.

//book[excerpt]/author[degree]

Anywhere within the current document, find all `author` elements containing at least one child `degree` element that are contained inside a `book` element that has at least one `excerpt` element.

//book[excerpt][title]

Anywhere within the current document, find all `book` elements containing at least one `excerpt` element and at least one `title` element.

//author[(degree \$or\$ award) \$and\$ publication]

Anywhere within the current document, find all `author` elements that contain at least one `degree` element or `award` element and at least one `publication` element.

```
//degree[@from $ne$ 'Harvard']
```

Anywhere within the current document, find all `degree` elements where the `from` attribute is not equal to Harvard.

```
//(book/author)[end()]
```

Anywhere within the current document, find the last `author` element from the entire set of `author` elements that are contained within `book` elements.

```
//my:book
```

Anywhere within the current document, find all `book` elements with the namespace `my`.

```
/*my:book/my:title
```

Find the `title` element with a `my` prefix within a `book` element with a `my` prefix that is a child of the root node of the document.

```
//book[index() $le$ 1]
```

Find the first two books anywhere within the current document.

```
//author[publication!count() $gt$ 1]
```

Anywhere within the current document, find all `author` elements that contain more than one `publication` element.

```
//author[$any$ last-name = 'Bob']
```

Find all `author` elements where at least one of the `last-name` elements is "Bob" anywhere within the current document.

```
//book $union$ //book/author
```

Anywhere within the current document, find all `book` elements and all `author` elements that are children of `book` elements.

```
/bookstore/book/author[last-name/regex('Bo. ')]
```

Find all the `author` elements that have `last-name` elements that contain the string "Bo." and another single character.

```
//book[@style/regex("no.*")]
```

Find all the `book` elements that have a `style` attribute that contains the regular expression pattern "no.*"

```
/comment()[0]
```

Find the first `comment` element starting from the root of the current document.

Extracting Custom Attributes From XML Documents

For custom attributes that are extracted from an XML document, Trading Networks returns specific values for different attribute data types:

- **STRINGLIST** and **STRING**
- **NUMBERLIST** and **NUMBER**
- **DATETIMELIST** and **DATETIME**

The following tables demonstrate the values that Trading Networks returns using the example XQL query of: `/node/item`. In some cases, the query returns a null value.

Extracting STRINGLIST or STRING ATTRIBUTES

This table shows the values Trading Networks stores for a **STRINGLIST** or **STRING** attribute based on the value returned by the XQL query `/node/item`

XML Data	XQL query returns (result returned as String List)	Results stored as STRINGLIST	Results stored as STRING
<code><node></node></code>	null	null	null
<code><node> <item></item> </node></code>	<code>{""}</code> (String List with one element, an empty String)	<code>{null}</code> (String List with one element, a null String)	null
<code><node> <item>value</item> </node></code>	<code>{"value"}</code> (String List with one element)	<code>{"value"}</code> (String List with one element)	"value"
<code><node> <item>value</item> <item></item> </node></code>	<code>{"value", ""}</code> (String List with two elements)	<code>{"value", null}</code> (String List with two elements)	"value"
<code><node> <item></item> <item>value</item> </node></code>	<code>{"", "value"}</code> (String List with two elements)	<code>{null, "value"}</code> (String List with two elements)	null
<code><node> <item></item> <item></item> </node></code>	<code>{"", ""}</code> (String List with two elements)	<code>{null, null}</code> (String List with two elements)	null
<code><node> <item>valuea</item> <item> </item> <item></item> <item>valuez</item> </node></code>	<code>{"valuea", "", "", "valueb"}</code> (String List with four elements)	<code>{"valuea", null, null, "valueb"}</code> (String List with four elements)	"valuea"

Extracting NUMBERLIST or NUMBER ATTRIBUTES

This table shows the values Trading Networks stores for a **NUMBERLIST** or **NUMBER** attribute based on the value returned by the XQL query `/node/item`

XML Data	XQL query returns (result returned as a String List)	Results stored as NUMBERLIST	Results stored as NUMBER
<code><node></node></code>	null	null	null

XML Data	XQL query returns (result returned as a String List)	Results stored as NUMBERLIST	Results stored as NUMBER
<node> <item></item> </node>	{""} (String List with one element, an empty String)	{null} (Number List with one element, a null element)	null
<node> <item>123</item> </node>	{"123"} (String List with one element)	{"123"} (Number List with one element)	"123"
<node> <item>123</item> <item></item> </node>	{"123", ""} (String List with two elements)	{"123", null} (Number List with two elements)	"123"
<node> <item></item> <item>123</item> </node>	{"", "123"} (String List with two elements)	{null, "123"} (Number List with two elements)	null
<node> <item></item> <item></item> </node>	{"", ""} (String List with two elements)	{null, null} (Number List with two elements)	null
<node> <item>123</item> <item> </item> <item></item> <item>456</item> </node>	{"123", "", "", "456"} (String List with four elements)	{"123", null, null, "456"} (Number List with four elements)	"123"

Extracting DATETIMELIST or DATETIME ATTRIBUTES

This table shows the values Trading Networks stores for a **NUMBERLIST** or **NUMBER** attribute based on the value returned by the XQL query `/node/item`

XML Data	XQL query returns (result returned as a String List)	Results stored as DATETIMELIST	Results stored as DATETIME
<node></node>	null	null	null
<node> <item></item> </node>	{""} (String List with one element, an empty String)	{null} (Datetime List with one element, a null element)	null
<node> <item>2004/12/31</item> </node>	{"2004/12/31"} (String List with one element)	{"2004/12/31"} (Datetime List with one element)	"2004/12/31"
<node> <item>2004/12/31</item> <item></item> </node>	{"2004/12/31", ""} (String List with two elements)	{"2004/12/31", null} (Datetime List with two elements)	"2004/12/31"

XML Data	XQL query returns (result returned as a String List)	Results stored as DATETIMELIST	Results stored as DATETIME
</node>	with two elements)	with two elements)	
<node> <item></item> <item>2004/12/31</item> </node>	{"", "123"} (String List with two elements)	{null, "2004/12/31"} (Datetime List with two elements)	null
<node> <item></item> <item></item> </node>	{"", ""} (String List with two elements)	{null, null} (Datetime List with two elements)	null
<node> <item>2004/12/31</item> <item> </item> <item></item> <item>2005/01/31</item> </node>	{"2004/12/31", "", "", "2006/01/31 " (String List with four elements)	{"2004/12/31", null,null, "2005/0 1/31"} (Datetime List with four elements)	"2004/12/31"

D Large Document Handling

■ Overview	330
■ Sending Large XML Documents for Processing	330
■ Minimum Hardware Requirements	331
■ Configuring Large Document Handling	331
■ Defining Document Types for Large Documents	333
■ Creating Services that Recognize Large Document Handling	333
■ Increasing the Size of the Largest Document that Can Be Saved When Using DB2	336

Overview

By default, when Trading Networks receives a document, it keeps the document content in memory during processing. If a document is large, Trading Networks might encounter memory constraint problems. These problems can occur when Trading Networks tries to execute an XQL query against an XML document during document recognition or tries to perform pre-processing or processing actions.

You can set a property to specify when a document should be considered large. When Trading Networks receives a document that fits your specification, it processes the document as a *large document*. Trading Networks does not try to read content for large documents into memory; rather, it writes the document content to hard disk drive space (called *tspace*) and stores only a reference to the document content in memory. When Trading Networks needs to access document content during processing, if the content is in *tspace*, Trading Networks either retrieves a certain number of bytes of the document content or uses a Java InputStream to read the document content. The document content remains on disk until the service that processes the document, and all services invoked from that service, complete, and the time to live period you specify expires.

The Trading Networks features listed below are not available for large documents.

Feature	If specified, Trading Networks...
Queue for polling delivery method	Logs a message to the activity log and continues processing. The document is not available for your partner to poll.
Web service delivery method	Logs a message to the activity log and continues processing.
Validate Structure pre-processing action	Does not validate the document, logs a message to the activity log, and continues document processing.
Verify pre-processing action	Does not verify the digital signature, logs a message to the activity log, and continues document processing.

Sending Large XML Documents for Processing

Make sure trading partners do not use the `$xmldata` variable to send large XML documents to Trading Networks. For information about other methods partners can use to submit XML documents, see the topic about passing XML data to services in *Software AG Designer Online Help*.

Minimum Hardware Requirements

The minimum and recommended hardware requirements for handling large documents are shown below.

Hard Disk Space (MB)	RAM	# of CPUs
100 (256)	512 MB (1GB)	1 (2)

The hard disk drive space identifies only the amount of hard disk drive space that Trading Networks requires for normal processing. It does not account for the hard disk drive space that is necessary for temporarily saving the large documents during processing. To use large document handling, you must define hard disk drive space where Trading Networks can temporarily save documents (rather than store them in memory). The size of the hard disk drive space for temporarily saving documents will vary based on the number of documents you process concurrently and the size of the documents you process. For example, if your typical concurrent document load is 10, you would need a hard disk drive space that is 10 to 15 times the combined size of the documents being processed concurrently.

Configuring Large Document Handling

To configure large document handling

1. In My webMethods: **Administration > Integration > B2B Settings > Configure Properties**.
2. Add or update the `tn.BigDocThreshold` and `tn.xml.xqlThreshold` properties.
3. Click **Save**.
4. Shut down the host Integration Server.
5. Go to the `Integration Server_directory \instances\instance_name \config` directory, open the `server.cnf` file, and update the properties listed below.
 - Set the `watt.server.tspace.location` property to the absolute directory path to the hard disk drive space in which to temporarily store large documents. The directory you specify is on the same machine as the host Integration Server. Examples are `watt.server.tspace.location=D:\LargeDocTemp` for Windows, or `watt.server.tspace.location=/opt/webmethods/` for UNIX. If you do not set this property, Trading Networks uses the Java system property `java.io.tmpDir`, which defaults to the value of the environment variable `Temp` on most operating systems.

- Note:**
- In a clustered environment, each Integration Server must have its own `tspace` location. Do not share the `tspace` location.
 - If you are setting up large document handling for multiple `webMethods` products, all products use this property. Integration Server manages the allotment of space to each product.
 - Each file stored in this directory is given the name `DocResxxxxx.dat`, where `xxxxxx` can vary in length and character.

- Set the `watt.server.tspace.max` to the maximum number of bytes to store at any one time in the hard disk drive space defined on the `watt.server.tspace.location` property. If Trading Networks (or another `webMethods` product) tries to write a large document to the hard disk drive space that will cause the specified number of bytes to be exceeded, an error message is displayed on the Integration Server console and the document is not stored. Specify a positive integer in bytes. The default is 52,428,800 bytes (50 MB).

If you are setting up large document handling for multiple `webMethods` products, all products use this property. The Integration Server makes sure the hard disk drive space used for all products does not exceed the value you specify.

- Note:** The size of the hard disk drive space for temporarily saving documents varies based on the number of documents you process concurrently and the size of the documents you process. For example, if your typical concurrent document load is 10, you would need a hard disk drive space that is 10 to 15 times the combined size of the documents being processed concurrently.

- When you submit a document using a Java `InputStream` over a network to Integration Server, Integration Server makes a request to read from a network `InputStream`. Set the `watt.server.keepAliveTimeout` property to the number of milliseconds Integration Server should wait for a response. If Integration Server does not receive a response in the specified amount of time, it times out with the error `Connection reset by peer`. Specify a positive integer in milliseconds. Software AG recommends that you set this property to 180000 (3 minutes). The default value is 15000 (15 seconds).

- Note:** The value you specify does not have to be long enough to accommodate reading the entire document, only to accommodate the length of time it takes to receive a response to a read request.

- Set the `watt.server.tspace.timeToLive` property to the number of milliseconds to temporarily store documents in `tspace`. Setting this property prevents the deletion of a document as soon as it is created, thus preventing any resulting exceptions that occur when you try to read back from the document. For example, if you want to delete the documents from `tspace` three minutes after they are created, set this property to 180000 milliseconds.

The default value 0 tells Integration Server to delete documents from tspace only after creating a document in tspace. For example, if tspace contains 10 documents and this property is set to 0, Integration Server deletes the 10 documents only after it creates a document.

6. Save the file and restart Integration Server.

Defining Document Types for Large Documents

Large XML files that are sent to Trading Networks for processing must contain the XML declaration (for example, `<?xml...`) as the first line. No blank space can exist before the XML declaration, or Trading Networks will throw an exception.

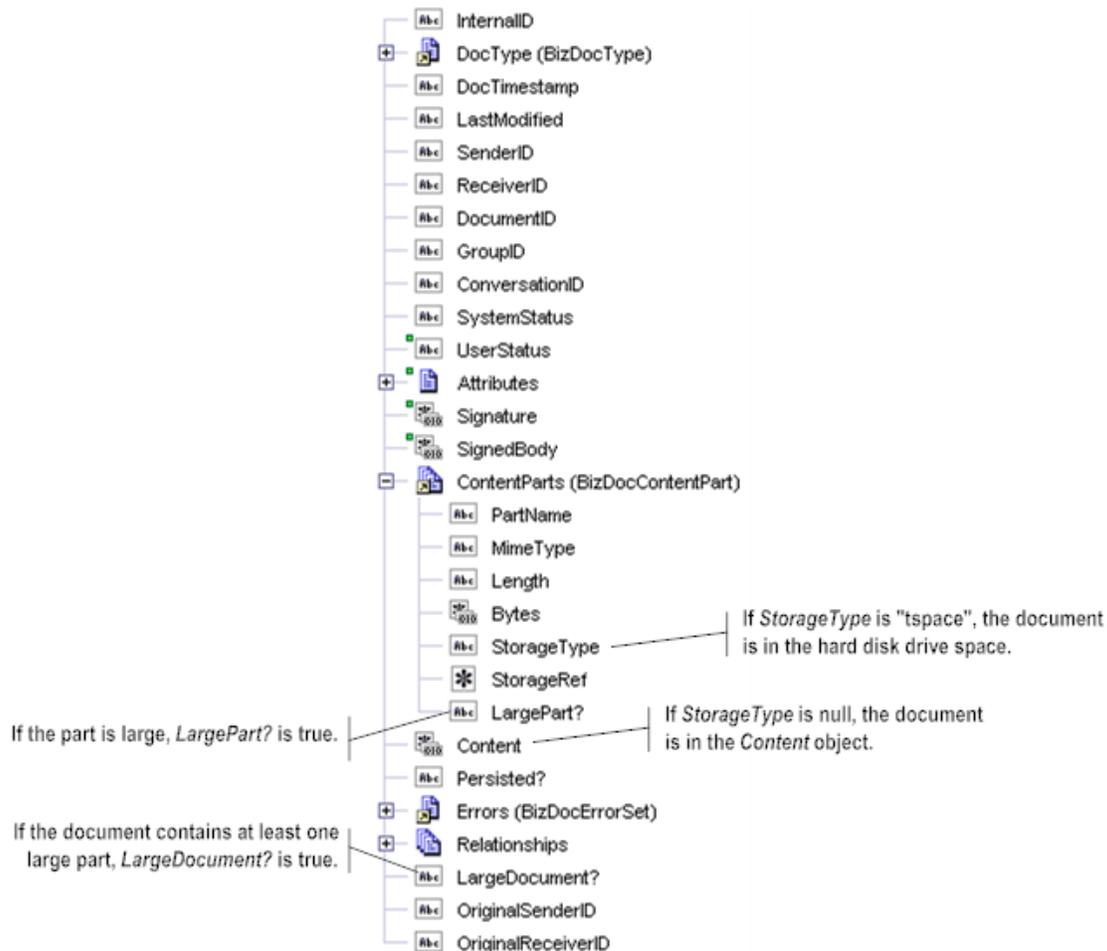
XML document types can include XQL queries for document recognition and XQL queries for extracting document attributes. Trading Networks does not read the entire content of large XML documents into memory to perform XQL queries; instead, it reads only the number of bytes you specify on the `tn.xml.xqlThreshold` into memory. The XQL queries, therefore, must match nodes in the first portion of the document. If the XQL queries identify nodes in the portion that is not in memory, the query fails, and Trading Networks cannot match the document type to the document or extract document attributes.

Creating Services that Recognize Large Document Handling

Code custom immediate and scheduled delivery services and services you use with the Execute a Service processing action to retrieve document content from hard disk drive space for large documents and from memory for regular documents. For example, a processing rule might match both large documents and regular documents. If the rule invokes a service, that service must take appropriate actions to obtain the document content whether the content is in hard disk drive space or memory.

Determining Where the Document Content is Stored

Trading Networks keeps information about all documents in the pipeline in the bizdoc variable. Fields within the bizdoc indicate whether Trading Networks considers the document large, and therefore, indicates whether the document is stored in memory (if the document is not considered large) or on hard disk drive space (if the document is considered large). The figure below shows how to use fields in the bizdoc to determine where the document content is stored.



Use the following fields in *bizdoc* to determine whether the document is stored in memory or on the hard disk drive space:

Field	Indicates...
LargePart?	Whether a particular content part is considered large. If not, the value is <code>false</code> . If so, the value is <code>true</code> ; Trading Networks sets <code>StorageType</code> to <code>tspace</code> and leaves the <code>Content</code> variable null.
LargeDocument?	Whether at least one content part in the document is considered large. If so, the value is <code>true</code> . If not, the value is <code>false</code> ; Trading Networks stores the document content in the <code>Content</code> variable and leaves <code>StorageType</code> null.

Retrieving Document Content from Hard Disk Drive Space

Use the `wm.tn.doc:getContentPartData` service to retrieve document content from the hard disk drive space. When you invoke this service, specify the following as input:

Input Variable	Description
<i>BizDocEnvelope</i>	BizDocEnvelope Trading Networks created for the document.
<i>partName</i>	Part name of the document content to retrieve.
<i>getAs</i>	Whether to: <ul style="list-style-type: none"> ■ Retrieve a specified number of bytes of the document content (<code>bytes</code>). The service will return a byte array that contains the retrieved bytes. ■ Obtain the document content using a Java <code>InputStream</code> object (<code>stream</code>). The service will return a Java <code>InputStream</code> object.
<i>startIndex</i>	If you specified <code>bytes</code> , the index into the document at which to start retrieving content. For example, to retrieve document content from the beginning, specify 0.
<i>byteCount</i>	If you specified <code>bytes</code> , specify the number of bytes of content to retrieve.

You can invoke this service for all documents, regardless of whether a document is considered large or not.

Retrieving the Content for Document Delivery

To deliver a document, you need the original document content (that is., the content Trading Networks received before processing). Use the `wm.tn.doc:getDeliveryContent` service to retrieve the document content. When you invoke this service, specify the following as input:

Input Variable	Description
<i>BizDocEnvelope</i>	BizDocEnvelope Trading Networks created for the document.

You can invoke this service for all documents. The service determines whether the document is considered large or not. If the document is not considered large, the service returns a byte array that contains the retrieved delivery content. If the document is

considered large, the service retrieves the delivery content using a Java InputStream object.

Increasing the Size of the Largest Document that Can Be Saved When Using DB2

When you define columns for a DB2 database, DB2 requires you to specify a size limit. When defining the Trading Networks database, the DB2 SQL script defines the size of the BizDocContent.Content field in the BIZDOCCONTENT table to have a 2M size limit. As a result, the largest document that Trading Networks can save to the DB2 database is 2M.

If you are using a DB2 database for the Trading Networks database and want to save documents that are larger than 2M, update the size limit for the BizDocContent.Content field to accommodate the largest document you expect to process. If you use the Trading Networks Archive tables, also update the size of the ARCHIVE_BizDocContent.Content field.

E Example of User Status in Document Processing

■ An Example of User Status while Processing a Document	338
---	-----

An Example of User Status while Processing a Document

1. Trading Networks receives a document and recognizes that its document type is PO. The document type does not specify to extract the UserStatus system attribute, so the user status is null.
2. The document matches the processing rule criteria that the document type must be PO and the document can have any user status. The rule specifies the Alert e-Mail action, and the e-mail contains a URL for a Web page that displays the PO. The processing rule also changes the user status to Needs Approval.
3. The user clicks the URL to open the Web page to view the PO. The Web page contains an HTML form that lets the user accept or reject the PO.
4. If the user accepts the PO, the code uses the built-in service `wm.tn.doc:changeStatus` to set the user status associated with the document to Accepted. If the user rejects the PO, the code uses the built-in service to set the user status associated with the document to Rejected.
5. The code for the Web form invokes the `wm.tn:submit` built-in service that sends the document back to the Integration Server for processing. Trading Networks recognizes the document type is PO and determines the user status.
6. The document matches either the processing rule that indicates the document type must be PO and the user status is Accepted or the document type is PO and the user status is Rejected. The rule specifies the Execute a Service action, and the service processes the purchase order.

