

**Oracle® Retail Extract, Transform, and Load**

Programmer's Guide

Release 13.2.4

**E29068-01**

March 2012

Oracle Retail Extract, Transform, and Load Programmer's Guide, Release 13.2.4

Copyright © 2012, Oracle and/or its affiliates. All rights reserved.

Primary Author: Susan McKibbin

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

## Value-Added Reseller (VAR) Language

### Oracle Retail VAR Applications

The following restrictions and provisions only apply to the programs referred to in this section and licensed to you. You acknowledge that the programs may contain third party software (VAR applications) licensed to Oracle. Depending upon your product and its version number, the VAR applications may include:

(i) the **MicroStrategy** Components developed and licensed by MicroStrategy Services Corporation (MicroStrategy) of McLean, Virginia to Oracle and imbedded in the MicroStrategy for Oracle Retail Data Warehouse and MicroStrategy for Oracle Retail Planning & Optimization applications.

(ii) the **Wavelink** component developed and licensed by Wavelink Corporation (Wavelink) of Kirkland, Washington, to Oracle and imbedded in Oracle Retail Mobile Store Inventory Management.

(iii) the software component known as **Access Via**<sup>™</sup> licensed by Access Via of Seattle, Washington, and imbedded in Oracle Retail Signs and Oracle Retail Labels and Tags.

(iv) the software component known as **Adobe Flex**<sup>™</sup> licensed by Adobe Systems Incorporated of San Jose, California, and imbedded in Oracle Retail Promotion Planning & Optimization application.

You acknowledge and confirm that Oracle grants you use of only the object code of the VAR Applications. Oracle will not deliver source code to the VAR Applications to you. Notwithstanding any other term or condition of the agreement and this ordering document, you shall not cause or permit alteration of any VAR Applications. For purposes of this section, "alteration" refers to all alterations, translations, upgrades, enhancements, customizations or modifications of all or any portion of the VAR Applications including all reconfigurations, reassembly or reverse assembly, re-engineering or reverse engineering and recompilations or reverse compilations of the VAR Applications or any derivatives of the VAR Applications. You acknowledge that it shall be a breach of the agreement to utilize the relationship, and/or confidential information of the VAR Applications for purposes of competitive discovery.

The VAR Applications contain trade secrets of Oracle and Oracle's licensors and Customer shall not attempt, cause, or permit the alteration, decompilation, reverse engineering, disassembly or other reduction of the VAR Applications to a human perceivable form. Oracle reserves the right to replace, with functional equivalent software, any of the VAR Applications in future releases of the applicable program.



---

---

# Contents

<b>Send Us Your Comments</b> .....	xiii
<b>Preface</b> .....	xv
Audience .....	xv
Documentation Accessibility .....	xv
Related Documents .....	xvi
Customer Support .....	xvi
Review Patch Documentation .....	xvi
Oracle Retail Documentation on the Oracle Technology Network .....	xvi
Conventions .....	xvii
<b>1 Introduction</b>	
<b>Technical Specifications</b> .....	1-2
RAC Database Support .....	1-2
jdbcconnectionstring .....	1-2
Check Supported Database Server Requirements .....	1-2
<b>Data Integration</b> .....	1-3
<b>2 RETL Interface</b>	
<b>Terms</b> .....	2-1
<b>RETL XML Interface</b> .....	2-2
<b>Operator Nesting</b> .....	2-3
<b>RETL Java Interface</b> .....	2-3
Initialization .....	2-4
Flows .....	2-4
Operators .....	2-4
Properties .....	2-4
Datasets .....	2-4
<b>3 RETL Program Flow</b>	
<b>General Flow</b> .....	3-1
<b>A Simple Flow</b> .....	3-2
<b>A More Complex Flow</b> .....	3-3
<b>Online Help</b> .....	3-4

Debugging with RETL .....	3-4
Producing Graphical Output of Flows with RETL.....	3-5
Performance Logging with RETL.....	3-6
Sample Log File .....	3-7

## 4 Database Operators

ORAREAD .....	4-1
ORAWRITE .....	4-1
UPDATE .....	4-1
DELETE .....	4-1
INSERT .....	4-1
PREPAREDSTATEMENT .....	4-2
Database Operators XML Specification Table.....	4-3
ORAREAD .....	4-3
ORAWRITE.....	4-4
UPDATE .....	4-9
DELETE .....	4-10
INSERT .....	4-11
PREPAREDSTATEMENT.....	4-12
Database Operator Examples .....	4-13
ORAREAD .....	4-13
ORAWRITE.....	4-13
UPDATE .....	4-14
DELETE .....	4-15
INSERT .....	4-16
PREPAREDSTATEMENT.....	4-17

## 5 RETL Parallel Processing

Pipeline Parallelism .....	5-1
Framework Parallelism .....	5-1
RETL Data Partitioning.....	5-2
Enabling RETL Data Partitioning.....	5-2
Partition Construction .....	5-2
Partitioning Types.....	5-4
Keyed Partitioning.....	5-4
Nonkeyed Partitioning.....	5-4
Partitioners .....	5-5
HASH.....	5-5
IMPORT.....	5-6
SPLITTER .....	5-8
DBREAD, ORAREAD .....	5-9
Operators and Partitioning Types .....	5-11
Parallel Property.....	5-14
Partitioned EXPORT .....	5-14
Partitioned GENERATOR.....	5-15
Partitioned LOOKUP and CHANGECAPTURELOOKUP.....	5-16

Partitioned ORAWRITE .....	5-17
Flows with Multiple Partitioners .....	5-17
HASH Within a Partition .....	5-17
SPLITTER Within a SPLITTER Partition .....	5-20
Two HASH Operators Joined .....	5-21
Funneled Partitions.....	5-22
Data Partitioning Guidelines .....	5-22
<b>Operator Configuration for Data Partitioning .....</b>	<b>5-23</b>
<b>A Final Word on Data Partitioning .....</b>	<b>5-24</b>

## 6 Input and Output Operators

<b>DEBUG .....</b>	<b>6-1</b>
<b>NOOP .....</b>	<b>6-1</b>
<b>EXPORT .....</b>	<b>6-1</b>
<b>IMPORT .....</b>	<b>6-1</b>
<b>Input and Output Operators XML Specification Tables .....</b>	<b>6-2</b>
DEBUG.....	6-2
NOOP.....	6-2
IMPORT.....	6-2
EXPORT.....	6-3
<b>Input and Output Operators Examples.....</b>	<b>6-4</b>
IMPORT.....	6-4
EXPORT.....	6-4

## 7 Join Operators

<b>INNERJOIN.....</b>	<b>7-1</b>
<b>LEFTOUTERJOIN .....</b>	<b>7-1</b>
<b>RIGHTOUTERJOIN .....</b>	<b>7-1</b>
<b>FULLOUTERJOIN .....</b>	<b>7-1</b>
<b>LOOKUP .....</b>	<b>7-2</b>
<b>DBLOOKUP .....</b>	<b>7-2</b>
<b>Special Notes about Join Operators.....</b>	<b>7-2</b>
<b>Join Operators XML Specification Tables .....</b>	<b>7-3</b>
INNERJOIN .....	7-3
LEFTOUTERJOIN, RIGHTOUTERJOIN, FULLOUTERJOIN .....	7-3
LOOKUP.....	7-3
DBLOOKUP.....	7-4
<b>Join Operators Examples.....</b>	<b>7-5</b>
INNERJOIN .....	7-5
LEFTOUTERJOIN .....	7-5
RIGHTOUTERJOIN.....	7-5
FULLOUTERJOIN .....	7-6
LOOKUP.....	7-6
DBLOOKUP .....	7-6

## 8 Mathematical Operators

<b>BINOP</b> .....	8-1
<b>GROUPBY</b> .....	8-1
<b>GROUPBY on Multiple Partitions</b> .....	8-1
<b>Mathematical Operators XML Specification Tables</b> .....	8-2
<b>BINOP</b> .....	8-2
<b>GROUPBY</b> .....	8-2
<b>Mathematical Operators Examples</b> .....	8-3
<b>BINOP</b> .....	8-3
<b>GROUPBY</b> .....	8-4

## 9 Other Operators

<b>COMPARE</b> .....	9-1
<b>SWITCH</b> .....	9-2
<b>CHANGECAPTURE and CHANGECAPTURELOOKUP</b> .....	9-2
<b>COPY</b> .....	9-3
<b>DIFF</b> .....	9-3
<b>CLIPROWS</b> .....	9-4
<b>PARSER</b> .....	9-4
<b>EXIT</b> .....	9-5
<b>Other Operators XML Specifications</b> .....	9-6
<b>COPY</b> .....	9-6
<b>COMPARE</b> .....	9-6
<b>CLIPROWS</b> .....	9-6
<b>DIFF</b> .....	9-7
<b>CHANGECAPTURE and CHANGECAPTURELOOKUP</b> .....	9-8
<b>SWITCH</b> .....	9-9
<b>PARSER</b> .....	9-10
<b>EXIT</b> .....	9-10
<b>Other Operators Examples</b> .....	9-10
<b>COPY</b> .....	9-10
<b>SWITCH</b> .....	9-11
<b>COMPARE</b> .....	9-11
<b>CHANGECAPTURE</b> .....	9-11
<b>CHANGECAPTURELOOKUP</b> .....	9-11
<b>CLIPROWS</b> .....	9-12
<b>DIFF</b> .....	9-12
<b>PARSER</b> .....	9-12
<b>EXIT</b> .....	9-12
<b>PARSER</b> .....	9-12

## 10 Structures and Data Manipulation Operators

<b>CONVERT</b> .....	10-1
<b>Conversion Functions</b> .....	10-2
<b>FIELDMOD</b> .....	10-3
<b>FILTER</b> .....	10-3

GENERATOR.....	10-3
REMOVEDUP.....	10-4
<b>Structures and Data Manipulation Operators XML Specification Tables .....</b>	<b>10-5</b>
CONVERT.....	10-5
FIELDMOD.....	10-5
FILTER.....	10-6
GENERATOR.....	10-6
REMOVEDUP.....	10-7
<b>Filter Expressions .....</b>	<b>10-7</b>
<b>Structures and Data Manipulation Operators Examples.....</b>	<b>10-8</b>
CONVERT.....	10-8
FIELDMOD.....	10-9
FILTER.....	10-10
GENERATOR.....	10-10
REMOVEDUP.....	10-10
<b>11 RETL Schema Files</b>	
<b>Schema File Requirements.....</b>	<b>11-1</b>
<b>Schema File XML Specification.....</b>	<b>11-2</b>
.....	11-3
<b>Delimited Record Schema.....</b>	<b>11-3</b>
Example Delimited Schema File.....	11-4
<b>Fixed-Length Record Schema.....</b>	<b>11-5</b>
Example Fixed-Length Schema File.....	11-5
nullvalue Considerations.....	11-6
<b>Configure RETL to Print Schema Output in Schema File Format.....</b>	<b>11-7</b>
<b>12 Common Operator Properties</b>	
<b>Common Operator XML Specification .....</b>	<b>12-1</b>
<b>13 Best Practices</b>	
Prerequisites.....	13-2
Interface Design Skills.....	13-2
ETL Technical Skills.....	13-2
<b>Project Initiation/Design/Functional Specification Best Practices .....</b>	<b>13-2</b>
Ask Discovery Questions .....	13-2
General Integration Questions.....	13-2
Application Domain Questions .....	13-3
Data-related and Performance Questions .....	13-3
Map Out The Movement of Data Visually.....	13-3
Define Concrete Functional Requirements for Each Module.....	13-3
Define Concrete Functional Designs for Each Module.....	13-4
Design a Test Plan Early in the Process .....	13-4
Design for Future Usage and Minimize Impact of Potential Changes.....	13-4
Agree on Acceptance Criteria.....	13-4
Document Design Assumptions, Issues, and Risks .....	13-5

<b>Code/Implementation/Test Best Practices</b> .....	13-5
Korn Shell Best Practices .....	13-5
Execute Commands Using \$(command) and Not command.....	13-5
Ensure set -f is set in a Configuration File.....	13-5
Write Flow to an Intermediate File and then Call RETL on that File.....	13-5
Secure/Protect Files and Directories that may Contain Sensitive Information.....	13-6
Make Function Calls Only a Few Layers Deep .....	13-6
Separate Environment Data from the Flow .....	13-6
Enclose Function Parameters in Double Quotes .....	13-6
Set Environment Variable Literals in Double Quotes .....	13-6
Use Environment Variables as \${VARIABLE} Rather than \$VARIABLE.....	13-7
Follow Module Naming Conventions .....	13-7
Log Relevant Events in Module Processing.....	13-8
Place Relevant Log Files in Well-known Directories.....	13-8
Use .ksh Templates .....	13-8
Document Each Flow's Behavior .....	13-8
RETL Flow Best Practices.....	13-8
<b>Review/Product Handoff</b> .....	13-12
Involve Support Personnel Early in the Project.....	13-12
Assign a Long-term Owner to the Project / Product / Interface.....	13-12

## 14 Sort, Merge, and Partitioning Operators

<b>COLLECT and FUNNEL</b> .....	14-1
<b>SORTCOLLECT and SORTFUNNEL</b> .....	14-1
<b>HASH</b> .....	14-1
<b>SPLITTER</b> .....	14-2
<b>SORT</b> .....	14-2
<b>MERGE</b> .....	14-2
<b>Sort and Merge Operators XML Specification Tables</b> .....	14-2
COLLECT/FUNNEL.....	14-2
HASH.....	14-2
SPLITTER .....	14-3
SORT .....	14-3
SORTCOLLECT/SORTFUNNEL .....	14-4
MERGE .....	14-4
<b>Sort, Merge, and Partitioning Operators Tag Usage Examples</b> .....	14-4
COLLECT .....	14-4
HASH.....	14-4
SORTCOLLECT.....	14-5
MERGE .....	14-5

## A Appendix: RETL Data Types

<b>RETL Data Type Properties</b> .....	A-1
<b>RETL Data Type/Database Data Type Mapping</b> .....	A-2
RETL Data Type to Oracle Data Type (ORAWRITE) .....	A-3
Oracle Data Type to RETL Data Type (ORAREAD).....	A-3

## **B Appendix: Database Configuration and Troubleshooting Guide**

RETL Database Configuration and Maintenance Notes .....	B-1
Debugging Database ETL Utilities .....	B-1
Database Semaphore Problems .....	B-1
Runaway Loader Processes .....	B-1
Troubleshooting RETL with Your Database .....	B-2

## **C Appendix: Troubleshooting Guide**

## **D Appendix: Default Conversions**

Default Conversions from UINT8 .....	D-1
Default Conversions from INT8.....	D-1
Default Conversions from UINT16 .....	D-2
Default Conversions from INT16.....	D-2
Default Conversions from UINT32 .....	D-2
Default Conversions from INT32.....	D-3
Default Conversions from UINT64 .....	D-3
Default Conversions from INT64.....	D-3
Default Conversions from SFLOAT .....	D-4
Default Conversions from DEFLOAT .....	D-4

## **E Appendix: Data Partitioning Quick Reference**

Partitioning the Data .....	E-1
Continuing the Partition .....	E-2
Hash Partitioning .....	E-2
Ending the Data Partition .....	E-2

## **F Appendix: Frequently Asked Questions**

## **G Appendix: Database Connections Quick Reference**

Setting the Environment Variables .....	G-1
ORAREAD .....	G-1
ORAWRITE .....	G-2



---

---

# Send Us Your Comments

Oracle Retail Extract, Transform, and Load Programmer's Guide, Release 13.2.3

Oracle welcomes customers' comments and suggestions on the quality and usefulness of this document.

Your feedback is important, and helps us to best meet your needs as a user of our products. For example:

- Are the implementation steps correct and complete?
- Did you understand the context of the procedures?
- Did you find any errors in the information?
- Does the structure of the information help you with your tasks?
- Do you need different information or graphics? If so, where, and in what format?
- Are the examples correct? Do you need more examples?

If you find any errors or have any other suggestions for improvement, then please tell us your name, the name of the company who has licensed our products, the title and part number of the documentation and the chapter, section, and page number (if available).

---

---

**Note:** Before sending us your comments, you might like to check that you have the latest version of the document and if any concerns are already addressed. To do this, access the new Applications Release Online Documentation CD available on My Oracle Support and [www.oracle.com](http://www.oracle.com). It contains the most current Documentation Library plus all documents revised or released recently.

---

---

Send your comments to us using the electronic mail address: [retail-doc\\_us@oracle.com](mailto:retail-doc_us@oracle.com)

Please give your name, address, electronic mail address, and telephone number (optional).

If you need assistance with Oracle software, then please contact your support representative or Oracle Support Services.

If you require training or instruction in using Oracle software, then please contact your Oracle local office and inquire about our Oracle University offerings. A list of Oracle offices is available on our Web site at [www.oracle.com](http://www.oracle.com).



---

---

# Preface

The Oracle Retail Extract Transform and Load (RETL) Programmer's Guide contains a complete description of the RETL programming language, as well as installation, configuration, and maintenance instructions.

## Audience

The RETL Programmer's Guide is intended for all RETL flow developers.

## Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at <http://www.oracle.com/accessibility/>.

### Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

### Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

### Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/support/contact.html> or visit <http://www.oracle.com/accessibility/support.html> if you are hearing impaired.

## Related Documents

For more information, see the following documents in the Oracle Retail Extract, Transform, and Load Release 13.2.4 documentation set:

- *Oracle Retail Extract, Transform, and Load Installation Guide*
- *Oracle Retail Extract, Transform, and Load Release Notes*

## Customer Support

To contact Oracle Customer Support, access My Oracle Support at the following URL:

<https://support.oracle.com>

When contacting Customer Support, please provide the following:

- Product version and program/module name
- Functional and technical description of the problem (include business impact)
- Detailed step-by-step instructions to re-create
- Exact error message received
- Screen shots of each step you take

## Review Patch Documentation

When you install the application for the first time, you install either a base release (for example, 13.1) or a later patch release (for example, 13.1.2). If you are installing the base release, additional patch, and bundled hot fix releases, read the documentation for all releases that have occurred since the base release before you begin installation. Documentation for patch and bundled hot fix releases can contain critical information related to the base release, as well as information about code changes since the base release.

## Oracle Retail Documentation on the Oracle Technology Network

Documentation is packaged with each Oracle Retail product release. Oracle Retail product documentation is also available on the following Web site:

[http://www.oracle.com/technology/documentation/oracle\\_retail.html](http://www.oracle.com/technology/documentation/oracle_retail.html)

(Data Model documents are not available through Oracle Technology Network. These documents are packaged with released code, or you can obtain them through My Oracle Support.)

Documentation should be available on this Web site within a month after a product release.

# Conventions

The following text conventions are used in this document:

<b>Convention</b>	<b>Meaning</b>
<b>boldface</b>	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.



---

---

# Introduction

Oracle Retail Extract Transform and Load (RETL) is a high-performance, scalable, platform-independent, parallel processing data movement tool. RETL addresses several primary needs:

- Database-independent applications
- Platform-independent applications
- Developing applications more quickly than possible with conventional coding methods (such as custom-crafted C or C++ code)
- High-performance data processing

To provide for these needs, RETL defines an interface in XML that applications can call to define ETL functions. This interface is in a well-defined XML form that allows access to any database that RETL supports.

RETL is a cross-platform development tool supported on Oracle Linux. The XML definitions do not change on a per-platform basis, so moving between hardware platforms is simplified.

Development of the XML instructions for RETL is simpler, faster, and less error-prone than writing C or C++ code. Applications can be completed much faster than possible with previous methods.

This guide shows the application developer how to rapidly deploy RETL in an application, and how to use it to manage your parallel processing system. This guide provides information on how to:

- Install and set up RETL
- Configure RETL
- Administer RETL
- Develop applications using RETL

Select the RETL operators you need to perform your application's data processing tasks. Manage your system resources to scale your application by partitioning data across multiple processing nodes.

RETL resides as an executable on your application platform, and on any other system that serves as a source or target of data.

## Technical Specifications

RETL is certified on the platform listed in this section. For historical reasons, the executable for RETL is called rfx. That executable name is used in this document.

The current configuration included with the RETL install package is documented in the Release Notes. If you have a configuration that is not included in the Release Notes, verify with Oracle Retail Customer Support whether your configuration is supported.

### RAC Database Support

RETL does not support OCI Database driver connection. RETL supports only THIN Database driver connection to connect to RAC and non RAC databases.

RETL database operators have been modified to support RAC databases. (For non RAC databases, no modifications are required to RETL database operators and RETL scripts .)

#### jdbcconnectionstring

For the RAC database, the jdbcconnectionstring property must be added to all database operators (ORAREAD, ORAWRITE, UPDATE, DELETE, INSERT, and PREPAREDSTATEMENT). The value for this property is RAC URL.

The following is an example of the value used for the jdbcconnectionstring property.

```
<PROPERTY name="jdbcconnectionstring" value="jdbc:oracle:thin:@(DESCRIPTION
=(ADDRESS_LIST =(ADDRESS = (PROTOCOL = TCP)(HOST = hostname1)(PORT =
1521))(ADDRESS = (PROTOCOL = TCP)(HOST = hostname2)(PORT = 1521))(LOAD_BALANCE =
yes ))(CONNECT_DATA =(SERVICE_NAME = serviceName))" />
```

---



---

**Note:** For RAC database examples, see Chapter 6, "[Database Operators](#)."

---



---



---



---

**Note:** For answers to frequently asked questions about the jdbcconnectionstring property, see "[Appendix: Frequently Asked Questions](#)."

---



---

### Check Supported Database Server Requirements

Supported On	Versions Supported
Database Server OS	<p>OS certified with Oracle Database 11gR2 Enterprise Edition. Options are:</p> <ul style="list-style-type: none"> <li>■ Oracle Linux 5 Update 5 for x86-64 (actual hardware or Oracle virtual machine)</li> <li>■ Red Hat Enterprise Linux 5 Update 5 (RHEL 5.5) for x86-64 (actual hardware or Oracle virtual machine)</li> <li>■ IBM AIX 6.1 (actual hardware or LPARs)</li> <li>■ Solaris 10 Sparc (actual hardware or logical domains)</li> <li>■ HP-UX 11.31 Integrity (actual hardware, HPVM, or vPars)</li> </ul>

Supported On	Versions Supported
Database Server 11gR2	<p>Oracle Database Enterprise Edition 11gR2 (11.2.0.2) with the following specifications:</p> <p><b>Components:</b></p> <ul style="list-style-type: none"> <li>■ Oracle Partitioning</li> <li>■ Examples CD (Formerly the companion CD)</li> </ul> <p><b>One-off Patches:</b></p> <ul style="list-style-type: none"> <li>■ 10170431—CTWR consumes a lot of CPU cycles.</li> </ul> <p>If ASM is used, apply the following patch to database home.</p> <ul style="list-style-type: none"> <li>■ 11808931—Merge request on top of 11.2.0.2.0 for defects 10410054 and 10422126</li> </ul> <p><b>Other components:</b></p> <ul style="list-style-type: none"> <li>■ Perl compiler 5.0 or later</li> <li>■ X-Windows interface</li> </ul>

## Data Integration

Integrate external data of these forms:

- UNIX flat files
- Oracle database tables



---



---

## RETL Interface

There are two interfaces to RETL. The traditional interface is through a flow file, which is a set of processing instructions in XML format. An additional interface that allows direct access to RETL functionality through a Java class library was introduced in RETL 11.3.

---



---

**Note:** Neither interface is preferred over the other. All functionality in one interface is available in the other. Developers should choose the interface to RETL depending on their specific integration requirements.

---



---

### Terms

Term	Definition
Flow	The instructions that RETL executes. A flow is a collection of operators.
Record	A RETL record is like a database record. It is a collection of fields of different datatypes.
Dataset	A dataset is a collection of records. It is similar to a database table. A dataset that is input to an operator is called an input dataset. A dataset that is output from a record is called an output dataset.
Operator	A RETL operator creates, transforms, or writes records in a dataset. When speaking of an operator independent of an interface type in this document, all uppercase letters are used.
Property	Operators have properties that tell the operator more information on how it should perform its job. For example, the IMPORT operator has an inputfile property that tells the IMPORT operator which file to import.

## RETL XML Interface

When using the RETL XML interface, the flow is specified in a file in XML format. (Consult a reference manual or Web site if you are unfamiliar with XML.) The XML format follows these rules:

- The root node is named FLOW. That is, the entire XML file contents are contained within <FLOW> and </FLOW> tags.
- The FLOW element has an optional attribute named "name." The value of the name attribute is used when logging information about the flow.
- The FLOW element requires two or more OPERATOR children elements that specify the operators.
- The OPERATOR element has a mandatory attribute named "type." The value of the type attribute specifies what type of operator is being specified. The operator type is not case-sensitive, although all lowercase characters are recommended. The different types of operators are described later in this document.
- The properties of an operator are specified using a PROPERTY element of the OPERATOR element. The PROPERTY element has two attributes named "name" and "value." The name attribute specifies the name of the property, and the value attribute specifies the value of the property. The valid properties for each operator are detailed later in this document.
- The input datasets to an operator are specified using the INPUT child element of the OPERATOR element. The INPUT element requires a single attribute named "name." Its value is the name of the dataset. Each input dataset must be an output dataset of another operator.
- The output datasets of an operator are specified using the OUTPUT child element of the OPERATOR element. The OUTPUT element requires a single attribute named "name." Its value is the name of the dataset. Each output dataset must be an input dataset to some other operator.
- XML comments are supported and encouraged.

Here is an example of a flow that reads in data from a file using the IMPORT operator and writes the records to an Oracle database using the ORAWRITE operator:

```
<FLOW name="dataload_flow">
  <!-- Import the data.txt file. -->
  <OPERATOR type="import">
    <PROPERTY name="inputfile" value="data.txt"/>
    <PROPERTY name="schemafilename" value="data.schema.xml"/>
    <OUTPUT name="data.v"/>
  </OPERATOR>

  <!-- Write to the MYDATA table on MYDATABASE -->
  <OPERATOR type="orawrite">
    <INPUT name="data.v"/>
    <PROPERTY name="dbuserid" value="username/password"/>
    <PROPERTY name="dbname" value="MYDATABASE"/>
    <PROPERTY name="tablename" value="MYDATA"/>
  </OPERATOR>
</FLOW>
```

## Operator Nesting

Operator nesting is a way of tying together logically related operators, thus reducing the number of lines in a flow. An OUTPUT for an operator can be replaced with the operator that would receive the corresponding INPUT.

---



---

**Note:** Using operator nesting does not change the performance of a flow.

---



---

For example, the following two flows are equivalent. The second flow replaces OUTPUT of sales.v in the IMPORT with the ORAWRITE, and the ORAWRITE does not have an INPUT element:

```
<FLOW>
  <OPERATOR type="import">
    <PROPERTY name="inputfile" value="sales.txt"/>
    <PROPERTY name="schemafilename" value="sales-schema.xml"/>
    <OUTPUT name="sales.v">
  </OPERATOR>
  <OPERATOR type="orawrite">
    <INPUT name="sales.v">
    <PROPERTY name="tablename" value="SALES"/>
    ...
  </OPERATOR>
</FLOW>

<FLOW>
  <OPERATOR type="import">
    <PROPERTY name="inputfile" value="sales.txt"/>
    <PROPERTY name="schemafilename" value="sales-schema.xml"/>
    <OPERATOR type="orawrite">
      <PROPERTY name="tablename" value="SALES"/>
      ...
    </OPERATOR>
  </OPERATOR>
</FLOW>
```

## RETL Java Interface

The RETL Java interface exposes the RETL functionality through a Java class library (that is, an API). This section documents how the Java interface compares with the XML interface. Documentation of each class can be found in the Javadocs distributed with the other RETL documentation.

Unless otherwise noted, all classes reside in the `com.retek.retl` package.

---



---

**Note:** When the RETL Java interface is used, RETL runs within the same JVM as the initiating application. This could lead to resource contention issues. The initiating application must be started with appropriate settings for initial heap size, maximum heap size, and thread stack size, as well as any garbage collection parameters.

For more information on parameters to set, consult the "Performance Tuning Guide" and the RETL shell script.

---



---

## Initialization

When the XML interface is used, the RETL run-time environment is set through environment variables and command line options to the RETL shell script. In the Java interface, the RETL runtime environment must be initialized using the RETL class. This class has methods that allow you to set the run-time parameters. See the Javadoc for more information.

## Flows

Just like the XML interface, the Java interface runs flows. The difference is that in the Java interface, a flow is an instance of the Flow class. Operators are added to the flow using the `add()` method. After all of the operators have been added to the flow, the flow is executed using the `run()` method.

## Operators

For each operator type in the XML interface, there is a corresponding class in the Java interface. For example, the IMPORT operator functionality is contained in a class named Import. The operators are documented in Chapters 7 through 13 of this guide.

The class name is always an uppercase first letter followed by lowercase letters, even for operator names that consist of more than one word. For example, the class name for the CHANGECAPTURE operator is `Changecapture`, not `ChangeCapture`.

## Properties

Each operator property is set through a method that has the same name as the property. The valid properties for each operator are documented in chapters 7 through 13.

The method name for setting a property is all lowercase, even when the property name consists of more than one word. For example, the `schemafile` property of the IMPORT operator is set through the `Import.schemafile()` method, not `Import.schemaFile()`. There are only setter methods; no getter methods are implemented.

## Datasets

Operators have `input()` and `output()` methods that set the name of the input and output datasets of an operator. These methods correspond to the INPUT and OUTPUT elements of the XML interface.

The following Java class runs a flow that is identical to the `dataload_flow` example flow from the XML interface:

```
import com.retek.retl.*;
import com.retek.retl.base.RETLException;

...

try
{
    RETLProperties retlProperties = new RETLProperties();
    retlProperties.setRETLHomeDirectory(...);
    RETL.initialize(retlProperties);

    Flow flow = new Flow();

    Import opImport = new Import();
```

```
        opImport.inputfile("data.txt");
        opImport.schemafile("data.schema.xml");
        opImport.output("data.v");
        flow.add(opImport);

        Orawrite opOrawrite = new Orawrite();
        opOrawrite.input("data.v");
        opOrawrite.dbuserid("username/password");
        opOrawrite.dbname("MYDATABASE");
        opOrawrite.tablename("MYDATA");
        flow.add(opOrawrite);

        flow.run();
    }
    catch (RETLEException e)
    {
        RETL.handleException(e);
        System.exit(1);
    }
}
```



---

---

## RETL Program Flow

The following text and diagram provides an overview of the RETL input-process-output model. The data input operators to RETL processing include the following:

- **DBREAD**, where data is read directly from a database
- **IMPORT**, where RETL accepts a data file
- **GENERATOR**, where RETL itself creates data input


Data process operators commonly include those that transform data in the dataset being processed, including **GROUPBY**, which can be used to sum values in a table column.

The data output process can include the use of the following operators:

- **DBWRITE**, where RETL writes data directly to the database
- **DEBUG**, which can be used to print records directly to the screen (stdout)
- **EXPORT**, which can export data to a flat file

### General Flow

The following diagram is a general representation of the form that all flows take. Every flow must have a data input (or source) and a data output (or sink). These flows may optionally perform transformations on the input data before sending it to the data output.




The following are some things to note about all flows:

- When an operator generates an output dataset, that output dataset must be input into another operator. The data "flows" from one operator to another along the datasets (the arrows between the operators above).
- DATA INPUT operators have no dataset inputs - they get data from outside sources such as an Oracle database or flat file.
- Similarly, DATA OUTPUT operators are terminal points within the flow. They do not generate output datasets. They export records to an external repository (for example, Oracle databases or flat files).

The next sections describe specific flows.

## A Simple Flow

The following diagram is one way to represent a simple flow. This flow contains three operators: ORAREAD, FIELDMOD and EXPORT. This flow also contains two datasets; these are the simple arrows that connect the operator boxes above. The arrow from ORAREAD to FIELDMOD means that records flow from the ORAREAD operator to the FIELDMOD operator. Likewise, the arrow from FIELDMOD to EXPORT indicates that the records from the FIELDMOD operator are passed onto the EXPORT operator.



In this flow diagram, several attributes are also shown to give more information about what exactly this flow does. This flow pulls the COMPHEAD table out of the ORA\_SID database (using the ORAREAD operator), renames the column CMPY\_IDNT to COMPANY (using the FIELDMOD operator), and exports the resulting records to the file results.txt (using the EXPORT operator).

Here is XML for the above flow:

```

<FLOW name="cmpy.flw">
  <OPERATOR type="oraread">
    <PROPERTY name="dbname" value="ORA_SID" />
    <PROPERTY name="connectstring" value="userid/passwd"/>
    <PROPERTY name="query">
      <![CDATA[
        select * from COMPHEAD
      ]]>
    </PROPERTY>
    <OUTPUT name = "data.v"/>
  </OPERATOR>
  <!-- This is the format for a comment.
  -- They can of course be
  -- multi-line, but cannot be nested.
  
```

```


-->
  <OPERATOR type="fieldmod">
    <PROPERTY name="rename" value="CMPY_IDNT=COMPANY" />
    <INPUT name = "data.v" />
    <OUTPUT name = "comphead.v" />
  </OPERATOR>
  <OPERATOR type="export">
    <INPUT name="comphead.v" />
    <PROPERTY name="outputfile" value="results.txt" />
  </OPERATOR>
</FLOW>

```

## A More Complex Flow

Flows can get much more complex than the simple flow shown previously.

The following flow pulls in data simultaneously from a flat file and an Oracle database. The dataset from the file is filtered based upon the specified filter criterion, then sorted. In parallel, the dataset from the Oracle database is modified using the FIELDMOD operator, and then sorted. The resulting datasets are joined, and the results are output to a flat file using the EXPORT operator.



Note the following:

- RETL can deal with multiple data sources within a single flow.
- RETL can output results to multiple files or database tables.

## Online Help

RETL has an online help system you can use to obtain operator syntax and usage information directly from the rfx binary. The following help is available:

- Specify property names for each operator
- Confirmation of whether a property is required or optional
- Valid values for each property
- A brief description of each property
- Default values for each property (if the property is optional)

For more information about this feature, see information about the "-o" option in the ["rfx Command Line Options"](#) section in Chapter 2, ["Installation and System Configuration."](#)

## Debugging with RETL

It is often useful for RETL flow programmers or system administrators to diagnose what is happening inside of RETL. This can assist in tracking down errors or performance bottlenecks. There are several mechanisms that RETL offers to aid in debugging:

- Verbose messages

Verbose can be turned on using the following variables:

```
export RFX_DEBUG=1
export RFX_SHOW_SQL=1
```

- RFX\_DEBUG turns on additional informational messages about RETL operations to standard output.
- RFX\_SHOW\_SQL turns on additional logging of database operations to standard output.

For more information about printing in schema file format, see ["Configure RETL to Print Schema Output in Schema File Format"](#) in Chapter 5, ["RETL Schema Files."](#)

For more information about printing graphical output of flows, see ["Producing Graphical Output of Flows with RETL."](#)

For more information about performance logging, see ["Logger Configuration"](#) in Chapter 2, ["Installation and System Configuration."](#)

## Producing Graphical Output of Flows with RETL

Beginning with version 11.2, RETL offers an option to print graphs of flows. This allows a flow developer or user to get an actual picture of a RETL flow, and of RETL operators and link names that connect each operator. In addition to a visual aid, this graph is invaluable for tracing problems with a flow. RETL error messages often denote which operator threw a particular error. By noting the link names that connect operators, you can trace problems directly back to the source operator, fix the problem, and rerun the flow. An example of running a flow with the `-g` command line option is as follows:

1. Enter the following:

```
rfx -f <flow name.xml> -g
```

2. Load `rfx<random#>_serial_flow.html` in any Web browser to view the output.


---

### XML Flow

```
<FLOW name="example.flw">
  <OPERATOR type="import">
    <PROPERTY name="inputfile"
      value="input0.dat" />
    <PROPERTY name="schemafile"
      value="schema0.xml" />
    <OUTPUT name="import0.v" />
  </OPERATOR>
  <OPERATOR type="fieldmod">
    <INPUT name="import0.v" />
    <PROPERTY name="keep"
      value="A D" />
    <OUTPUT name="leftside.v" />
  </OPERATOR>
  <OPERATOR type="import">
    <PROPERTY name="inputfile"
      value="input1.dat" />
    <PROPERTY name="schemafile"
      value="schema1.xml" />
  </OPERATOR>
  <OPERATOR type="fieldmod">
    <INPUT name="import1.v" />
    <PROPERTY name="keep"
      value="B D" />
    <OUTPUT name="rightside.v" />
  </OPERATOR>

  <OPERATOR type="lookup">
    <INPUT name="leftside.v" />
    <INPUT name="rightside.v" />
    <PROPERTY name="tablekeys"
      value="D" />
    <OUTPUT name="joined.v" />
  </OPERATOR>
  <OPERATOR type="debug">
    <INPUT name="joined.v" />
  </OPERATOR>
</FLOW>
```

### RETL Visual Graph



This is an invaluable tool when debugging a RETL flow. For example, RETL might return the following error:

```
Exception in operator [lookup:1]
Record (A|B) can't be found in lookup relation!
```

To determine which lookup operator the error came from, follow these steps:

1. Print a RETL flow graph (see RETL syntax above).
2. Locate the operator with the name lookup:1 in the flow graph. (This can also be located directly in the XML, by counting down to the first lookup operator in the flow.)
3. Locate link names that connect to lookup:1, and search the XML flow file for the lookup operator that contains these link names. For example, in the above flow, lookup:1 has links named leftside.v and rightside.v as INPUTs, and joined.v as OUTPUT. You can use any of these to search for the appropriate lookup in the flow, because link names must be unique.

## Performance Logging with RETL

RETL records the following diagnostic information when the `retek.retl.performance` log4j logger is set to level DEBUG:

- Processing time for each operator
- Count of processed records for each operator
- Throughput in records per second for each operator
- Major events taking place within each operator (for example, start/stop times for queries in DBREAD operators, calls to native database load utilities in DBWRITE operators, or calls to external gsort in SORT)

This information is logged in the performance log file. Also, the first three items are used to create an Operator Time Comparison Chart and a Performance Graph.

The file locations of the Operator Time Comparison Chart and the Performance Graph are output when RETL completes processing:

```
All threads complete
```

```
Flow ran successfully
```

```
Open operator-time-comparison-chart-file in your browser to see the performance page.
```

```
1 File written: performance-graph-file
```


```
To view the output, use this command to load the file into dotty:<br>
```

```
dotty performance-graph-file
```


## Sample Log File

```
[2004-12-08 16:04:15,193]PERF: oraread:3 Starting query 1 of 1: select * from
stresstests_1000000
[2004-12-08 16:04:15,256]PERF: oraread:3 Ending query 1
[2004-12-08 16:04:15,384]PERF: oraread:4 Starting query 1 of 1: select * from
stresstests_1000000 ORDER BY SEQFIELD
[2004-12-08 16:04:15,420]PERF: oraread:4 Ending query 1
[2004-12-08 16:04:15,479]User time: sortfunnel:1 started!
[2004-12-08 16:04:15,480]User time: sort:1 started!
[2004-12-08 16:04:15,480]User time: oraread:3 started!
[2004-12-08 16:04:15,481]User time: oraread:4 started!
[2004-12-08 16:05:06,460]User time: Operator "pipeline:2" finished - processed
1,000,000 records in 50.977 seconds ( 19,616.69 records/sec)
[2004-12-08 16:05:06,464]User time: Operator "oraread:3" finished - processed
1,000,000 records in 50.968 seconds ( 19,620.154 records/sec)
[2004-12-08 16:05:06,475]PERF: sort:1 Starting gsort
[2004-12-08 16:05:20,426]PERF: sort:1 Ending gsort
[2004-12-08 16:05:20,529]User time: export:1 started!
[2004-12-08 16:06:14,122]User time: Operator "import:1" finished - processed
1,000,000 records in 53.685 seconds ( 18,627.177 records/sec)
[2004-12-08 16:06:14,125]User time: Operator "pipeline:1" finished - processed
1,000,000 records in 118.643 seconds ( 8,428.647 records/sec)
[2004-12-08 16:06:14,128]User time: Operator "sort:1" finished - processed
1,000,000 records in 118.643 seconds ( 8,428.647 records/sec)
[2004-12-08 16:06:14,191]User time: Operator "pipeline:3" finished - processed
1,000,000 records in 118.708 seconds ( 8,424.032 records/sec)
[2004-12-08 16:06:14,193]User time: Operator "oraread:4" finished - processed
1,000,000 records in 118.708 seconds ( 8,424.032 records/sec)
[2004-12-08 16:06:14,201]User time: Operator "pipeline:0" finished - processed
2,000,000 records in 118.72 seconds ( 16,846.361 records/sec)
[2004-12-08 16:06:14,204]User time: Operator "sortfunnel:1" finished - processed
2,000,000 records in 118.72 seconds ( 16,846.361 records/sec)
[2004-12-08 16:06:14,206]User time: Operator "export:1" finished - processed
2,000,000 records in 53.675 seconds ( 37,261.295 records/sec)
```

### Sample Operator Time Comparison Chart



### Sample Performance Graph





---

---

## Database Operators

This chapter describes database operators.

### ORAREAD

ORAREAD performs a read from Oracle databases. RETL uses JDBC technology to read data out of Oracle databases.

### ORAWRITE

ORAWRITE performs a load to Oracle databases. RETL uses SQL\*Loader (sqlldr) to load the records.

---

---

**Note:** The ORAWRITE operator may implicitly drop fields when writing to the database, if fields in the incoming record schema do not exist in the database table. Also, the incoming record schema determines what is set for nullabilities and maxlengths of fields as well. If there are any inconsistencies between the incoming record and the format of the database table, RETL may not flag this as a problem; instead, it relies on SQL\*Loader to throw errors or reject inconsistent records. This only happens when appending to an existing table and does not affect database writes when specifically creating or recreating a table.

---

---

### UPDATE

The UPDATE operator updates records in an Oracle database. Fields in the records are used to provide updated values and also select records for updating.

### DELETE

The DELETE operator deletes records from an Oracle database table. Fields in the records are used to select records for deleting.

### INSERT

The INSERT operator inserts records into an Oracle database table. This is the same functionality that ORAWRITE performs, only using a different technology (JDBC instead of SQL\*Loader). ORAWRITE is preferred from a performance standpoint.

## PREPAREDSTATEMENT

The PREPAREDSTATEMENT allows you to execute SQL commands using values from the current record as input into the SQL command.

The PREPAREDSTATEMENT can perform inserts, deletes, and updates like the INSERT, UPDATE, and DELETE operators, only there is much more flexibility. For example:

- Record field names can be used more than once.
- Complex WHERE clause logic can be used.
- SQL calculations can be used.

The statement property specifies the SQL to execute. The SQL specified can contain question marks ( ? ) that act as placeholders for values from the current record. The fields property contains the names of the fields that will be used to provide the values. The fields must be specified in the same order as the question mark placeholders, but the same field can be specified multiple times.

---

---

**Note:** One important restriction is that there must be a database column with the same name for any field specified in the fields property.

---

---

All of the following are valid values for the statement property:

- `UPDATE EMPLOYEES SET SALARY = ? WHERE EMPLOYEE_ID = ? AND JOB_TITLE = 'SOFTWARE ENGINEER'`

This statement applies the new salary only if the employee is a software engineer. The JOB\_TITLE field does not need to be in the RETL dataset.

- `UPDATE EMPLOYEES SET SALARY = SALARY * 1.10 WHERE EMPLOYEE_ID = ?`

This statement gives a 10% raise to an employee. The 10% calculation is performed by the database, not RETL.

- `DELETE FROM EMPLOYEES WHERE EMPLOYEE_ID = ? AND LOCKED = "false"`

This statement deletes a record from the EMPLOYEES table, provided its LOCKED column is set to false.

- `INSERT INTO EMPLOYEES (LAST_NAME, FIRST_NAME, SALARY, JOB_TITLE) VALUES (?, ?, 0, 'UNKNOWN')`

This statement inserts a record into the employees table with a salary of 0 and a title of 'UNKNOWN'. The input dataset does not need to have fields named SALARY and JOB\_TITLE.

## Database Operators XML Specification Table

The following table outlines database read/write operators.

---



---

**Note:** Throughout the RETL Programmer's Guide, *schema* can apply to the data structure that the RETL applies to a dataset (such as a schema file), or *schema* can apply to the database owner of a table (the *schemaowner* property of a database write operator).

---



---

### ORAREAD

Element Name	Name	Value
PROPERTY	dbname	Name of the Oracle database.
PROPERTY	connectstring	username/password
PROPERTY	maxdescriptors	Maximum number of columns allowed to be returned from the database. This is an optional property. The default value is 70.
PROPERTY	Query	<p>The SQL query (select statement) used to extract the data from the database. SQL syntax must be executable for the database environment.</p> <p><b>Note:</b> This property must be enclosed in a CDATA tag in order to allow all SQL query types. See examples for more information.</p> <p>If any kind of function is applied to a column in the select statement, the column will be extracted as a DFLOAT even if the column in the database is defined as an int or string type. You may have to perform a CONVERT to change datatypes in order to compensate for this behavior.</p>
PROPERTY	datetotimestamp	This is an optional property. When it is set to true, ORAREAD returns DATE columns with the time included. When false, the time is not included. The default for this property is "false".
PROPERTY	hostname	<p>Hostname or IP address</p> <p>This is an optional property. The fully-specified hostname or IP address where the database resides. This defaults to the localhost if not specified.</p> <p><b>Note:</b> This property should only be specified when it is known that connections are being made to a remote database.</p> <p>This can be specified as a default in rfx.conf for convenience.</p>
PROPERTY	port	<p>Port number</p> <p>This is a required property. The port on which the database resides. This defaults to 1521 if not specified in an operator or in rfx.conf.</p> <p><b>Note:</b> There are preset defaults in rfx.conf for the necessary port numbers. The default for Oracle is 1521. The Oracle utility <i>tnsping</i> can be used to obtain the port number.</p> <p>Verify with your database administrator the proper port for your database installation.</p>

Element Name	Name	Value
PROPERTY	sp_prequery	Special query that allows the flow developer to specify a stored procedure to be run prior to the main query that ORAREAD executes. The current implementation does not allow variables to be specified, and it does not allow data to be returned.
PROPERTY	sp_postquery	Special query that allows the flow developer to specify a stored procedure to be run after the main query that ORAREAD executes. The current implementation does not allow variables to be specified, and it does not allow data to be returned.
PROPERTY	numpartitions	Optional property indicating the number of data partitions to create. A query must be specified for data partition.
PROPERTY	arraysize	Optional property indicating the number of rows to pre-fetch. The default is 100. Increasing this value may improve performance, because fewer trips are made to the database. However, memory usage increases because more records are being held in memory at one time.
PROPERTY	useString	Optional property used for testing the latin type characters.  <b>Note:</b> When this property is added with value as true, the method used will be <code>getString( )</code> . The default method used is <code>getBytes( )</code> . In the flow XML, add the property as follows.  <pre>&lt;PROPERTY name="useString" value="true" /&gt; //</pre> By default it is false.
PROPERTY	jdbccconnectionstring property	Provide RAC URL to connect to RAC database. Use this property only if you want to connect to RAC database.  Example value:  <pre>jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS_ LIST = (ADDRESS = (PROTOCOL = TCP) (HOST = hostname1) (PORT = 1521)) (ADDRESS = (PROTOCOL = TCP) (HOST = hostname2) (PORT = 1521)) (LOAD_ BALANCE = yes )) (CONNECT_DATA = (SERVICE_NAME = serviceName)))</pre>
OUTPUT		The output dataset name.

## ORAWRITE

Element Name	Name	Value
PROPERTY	Dbname	Name of the Oracle database.
PROPERTY	dbuserid	username/password
PROPERTY	maxdescriptors	Maximum number of columns allowed to be written to the database.
PROPERTY	schemaowner	The database owner of the table.
PROPERTY	tablename	Name of the Oracle table to which the data is written.

Element Name	Name	Value
PROPERTY	method	<p>This is an optional property to specify the loading method. The default is conventional.</p> <p>direct—Load the data using the SQL*Loader utility with direct=true.</p> <p>conventional—Load the data using SQL*Loader utility with direct=false.</p>
PROPERTY	mode	<p>This is an optional property.</p> <p>append—Add the new data in addition to existing data on the target table.</p> <p>truncate—Delete the data from the target table before loading the new data.</p> <p>The default value is append.</p>
PROPERTY	createtablemode	<p>This is an optional property that allows the developer to specify table creation mode. If this property is not given, the flow will fail if the specified table does not already exist.</p> <p>recreate—Drops the table and recreates the same table before writing the data.</p> <p>create—Use this option when the target table does not exist. It will create the table before writing the data only if the table does not exist.</p>
PROPERTY	jdbcdriver	<p>Connects to the database through the specified database driver.</p> <p>This is an optional property that defaults to THIN.</p>
PROPERTY	allowedrejects	<p>A value greater than zero that indicates how many errors are allowed before SQL*Loader exits with a critical error. This is an optional property that defaults to 50.</p>
PROPERTY	exit_threshold	<p>Optional property that allows configuration of the level of errors or warnings from SQL*Loader to cause RETL to abort. Valid values are as follows, from most strict to least strict in causing RETL to abort:</p> <p>warning—Any SQL*Loader warnings or fatal errors will cause RETL to abort.</p> <p>exceeded_rejects—SQL*Loader rejects exceeding the allowedreject property, or SQL*Loader fatal errors, will cause RETL to abort.</p> <p>fail_or_fatal—Only SQL*Loader failures or fatal errors will cause RETL to abort.</p> <p>The default value is fail_or_fatal.</p>
PROPERTY	numloaders	<p>Specifies the degree of parallelism by executing a number of SQL*Loaders running in parallel. May be specified in either conventional or direct mode, although there are certain restrictions for each. (See Oracle SQL*Loader syntax for information.)</p> <p><b>Note:</b> Multiple tempdir properties should be used in conjunction with the numloaders property in a one-to-one mapping to maximize performance. For example, if running 8 ways parallel, 8 distinct temporary directories should be specified to maximize performance.</p>

Element Name	Name	Value
PROPERTY	parallel	<p>true or false</p> <p>Optional property that allows RETL to execute one SQL*Loader per RETL partition. This behaves similar to the numloaders property, but it connects SQL*Loader sessions directly to each parallel RETL partition.</p> <p><b>Note:</b> Parallel ORAWRITE automatically inherits parallelism from the numpartitions property in rfx.conf, in that one SQL*Loader session is spawned for each partition (for example, if numpartitions is set to 4, then 4 SQL*Loaders will be run in parallel). See the "Configuration" section in <a href="#">Chapter 2, "Installation and System Configuration"</a> for more on setting numpartitions.</p> <p>To override parallelism in ORAWRITE, set parallel to "false" and set numloaders to the desired number of concurrent SQL*Loaders.</p> <p>When running ORAWRITE in parallel mode, there are restrictions imposed by SQL*Loader. Refer to Oracle SQL*Loader syntax for more information on these limitations.</p>
PROPERTY	tempdir	<p>The temporary directories to use for the SQL*Loader data files. May be specified multiple times by specifying separate properties.</p> <p><b>Note:</b> To maximize performance, temporary directories should reside on separate disk controllers when specifying tempdir with parallel SQL*Loaders through the numloaders property.</p>
PROPERTY	partition	<p>Optional name of the partition of the Oracle table to which data is written.</p>
PROPERTY	outputdelimiter	<p>Optional. Output delimiter to be used when sending data to SQL*Loader.</p> <p><b>Note:</b> By default, SQL*Loader is sent fixed-format data. The outputdelimiter property can significantly speed up the performance of flows that have records with many fields that are not populated with data.</p> <p>The outputdelimiter should never be set to a value that can be part of the data.</p>
PROPERTY	loaderoptions	<p>Optional. Name-value pair options to be passed to SQL*Loader through the SQL*Loader parfile.</p> <p>In particular, the rows setting can be tweaked to improve performance. The setting specifies the number of rows per commit and defaults to 5000. Increasing this value will decrease the number of commits per SQL*Loader session.</p> <p><b>Note:</b> See Oracle documentation on SQL*Loader for more information about valid options. RETL does not validate any options passed through this parameter. If these options are incorrect, SQL*Loader will cause RETL to fail.</p>
PROPERTY	rows	<p>Number of rows per commit.</p> <p><b>Deprecated:</b> Use rows=numrows in the loaderoptions property.</p>

Element Name	Name	Value
PROPERTY	controlfileloadoptions	<p>Optional. Options to pass in the control file LOAD command. RETL handles the following:</p> <ul style="list-style-type: none"> <li>■ LOAD keyword</li> <li>■ INFILE clause</li> <li>■ APPEND, TRUNCATE, REPLACE keywords</li> <li>■ INTO TABLE clause</li> </ul> <p>In particular, the controlfileloadoptions property can be set to 'PRESERVE BLANKS' so that trailing blanks are not stripped by SQL*Loader. The outputdelimiter property should also be specified in this case, so that the spaces used for padding a fixed-width data file are not preserved.</p>
PROPERTY	sortedindexes	<p>Name of an existing database index.</p> <p>Indicates data sent to SQL*Loader exists in the same order as the specified database index. This eliminates the need to sort the new index entries and can give considerable performance increases when the sorted order of the records is known to match the index. Multiple sort indexes can be separated by commas. May be used in direct mode only.</p>
PROPERTY	singlerow	<p>Specify yes to update indexes as each new row is inserted. Default value is no.</p> <p><b>Note:</b> The singlerow property should be Yes only when it has been shown to increase performance. In most cases, it will make the load slower. Consult the Oracle SQL*Loader documentation for more information about the SINGLEROW option and when it makes sense to use it.</p>
PROPERTY	preload	<p>This is a special query that allows the flow developer to specify a stored procedure to be run prior to SQL*Loader execution. The current implementation does not allow variables to be specified, and it does not allow data to be returned. This is useful to drop indexes prior to a load, for example.</p>
PROPERTY	postload	<p>This is a special query that allows the flow developer to specify a stored procedure to be run after SQL*Loader execution. The current implementation does not allow variables to be specified, and it does not allow data to be returned. This is useful to rebuild indexes after a load.</p>
PROPERTY	hostname	<p>Hostname or IP address</p> <p>This is an optional property. The fully-specified hostname or IP address where the database resides. This defaults to the localhost if not specified.</p> <p><b>Note:</b> This property should only be specified when it is known that connections are being made to a remote database.</p> <p>This can be specified as a default in rfx.conf for convenience.</p>

Element Name	Name	Value
PROPERTY	port	<p>This is a required property. The port on which the database resides. This defaults to 1521 if not specified in an operator or in rfx.conf.</p> <p><b>Note:</b> There are preset defaults in rfx.conf for the necessary port numbers. The default for Oracle is 1521. The Oracle utility tnsping can be used to obtain the port number.</p> <p>Verify with your database administrator the proper port for your database installation.</p>
PROPERTY	tablespace	Tablespace name
PROPERTY	control_file	<p>File name of a valid a SQL*Loader control file that RETL should specify to SQL*Loader, instead of generating a control file dynamically. Using this property provides the most control over the SQL*Loader process. See the SQL*Loader documentation for information on the control file options.</p> <p>The file specified by this property is not deleted when the load is complete, unlike the dynamically generated control file used when this property is not specified.</p>
PROPERTY	data_file	File name to use for the SQL*Loader data file. Using this property allows you to control where the file resides.
PROPERTY	jdbcconnectionstring property	<p>Provide RAC URL to connect to RAC database. Use this property only if you want to connect to RAC database.</p> <p>Example value:</p> <pre>jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS_ LIST = (ADDRESS = (PROTOCOL = TCP) (HOST = hostname1) (PORT = 1521)) (ADDRESS = (PROTOCOL = TCP) (HOST = hostname2) (PORT = 1521)) (LOAD_ BALANCE = yes )) (CONNECT_DATA = (SERVICE_NAME = serviceName)))</pre>
INPUT		The input dataset name.

## UPDATE

Element Name	Name	Value
PROPERTY	dbname	Name of the Oracle database.
PROPERTY	sid	Alias for dbname.
PROPERTY	dbuserid	username/password
PROPERTY	maxdescriptors	Maximum number of columns allowed to be written to the database.
PROPERTY	schemaowner	The database owner of the table.
PROPERTY	tablename	Name of the Oracle table to which the data is written.
PROPERTY	table	Alias for tablename.
PROPERTY	hostname	<p>Hostname or IP address</p> <p>This is an optional property. The fully-specified hostname or IP address where the database resides. This defaults to the localhost if not specified.</p> <p><b>Note:</b> This property should only be specified when it is known that connections are being made to a remote database.</p> <p>This can be specified as a default in rfx.conf for convenience.</p>
PROPERTY	port	<p>This is a required property. The port on which the database resides. This defaults to 1521 if not specified in an operator or in rfx.conf.</p> <p><b>Note:</b> There are preset defaults in rfx.conf for the necessary port numbers. The default for Oracle is 1521. The Oracle utility tnsping can be used to obtain the port number.</p> <p>Verify with your database administrator the proper port for your database installation.</p>
PROPERTY	fields	Space-separated list of fields to use for selecting records to update. Fields in the input schema but not in the fields property are the fields that are updated.
PROPERTY	batchcount	Number of records to process between commit calls. Default value is 1000.
PROPERTY	jdbcconnectionstring property	<p>Provide RAC URL to connect to RAC database. Use this property only if you want to connect to RAC database.</p> <p>Example value:</p> <pre>jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=hostname1)(PORT=1521))(ADDRESS=(PROTOCOL=TCP)(HOST=hostname2)(PORT=1521))(LOAD_BALANCE=yes))(CONNECT_DATA=(SERVICE_NAME=serviceName)))</pre>
INPUT		The input dataset name.

**DELETE**

Element Name	Name	Value
PROPERTY	dbname	Name of the Oracle database.
PROPERTY	sid	Alias for dbname.
PROPERTY	dbuserid	username/password
PROPERTY	maxdescriptors	Maximum number of columns allowed to be written to the database.
PROPERTY	schemaowner	The database owner of the table.
PROPERTY	tablename	Name of the Oracle table to which the data is written.
PROPERTY	table	Alias for tablename.
PROPERTY	hostname	<p>Hostname or IP address</p> <p>This is an optional property. The fully-specified hostname or IP address where the database resides. This defaults to the localhost if not specified.</p> <p><b>Note:</b> This property should only be specified when it is known that connections are being made to a remote database.</p> <p>This can be specified as a default in rfx.conf for convenience.</p>
PROPERTY	port	<p>This is a required property. The port on which the database resides. This defaults to 1521 if not specified in an operator or in rfx.conf.</p> <p><b>Note:</b> There are preset defaults in rfx.conf for the necessary port numbers. The default for Oracle is 1521. The Oracle utility tnsping can be used to obtain the port number.</p> <p>Verify with your database administrator the proper port for your database installation.</p>
PROPERTY	fields	Space separated list of fields to use for selecting records to delete.
PROPERTY	batchcount	Number of records to process between commit calls. Default value is 1000.
PROPERTY	jdbconnectionstring property	<p>Provide RAC URL to connect to RAC database. Use this property only if you want to connect to RAC database.</p> <p>Example value:</p> <pre>jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = TCP) (HOST = hostname1) (PORT = 1521)) (ADDRESS = (PROTOCOL = TCP) (HOST = hostname2) (PORT = 1521)) (LOAD_BALANCE = yes )) (CONNECT_DATA = (SERVICE_NAME = serviceName)))</pre>
INPUT		The input dataset name.

**INSERT**

Element Name	Name	Value
PROPERTY	dbname	Name of the Oracle database.
PROPERTY	sid	Alias for dbname.
PROPERTY	dbuserid	username/password
PROPERTY	maxdescriptors	Maximum number of columns allowed to be written to the database.
PROPERTY	schemaowner	The database owner of the table.
PROPERTY	tablename	Name of the Oracle table to which the data is written.
PROPERTY	table	Alias for tablename.
PROPERTY	hostname hostname	<p>Hostname or IP address</p> <p>This is an optional property. The fully-specified hostname or IP address where the database resides. This defaults to the localhost if not specified.</p> <p><b>Note:</b> This property should only be specified when it is known that connections are being made to a remote database.</p> <p>This can be specified as a default in rfx.conf for convenience.</p>
PROPERTY	port	<p>This is a required property. It is the port on which the database resides and defaults to 1521 if not specified in an operator or in rfx.conf.</p> <p><b>Note:</b> There are preset defaults in rfx.conf for the necessary port numbers. The default for Oracle is 1521. The Oracle utility tnsping can be used to obtain the port number.</p> <p>Verify with your database administrator the proper port for your database installation.</p>
PROPERTY	fields	Space-separated list of fields to insert into the database table. Defaults to all fields.
PROPERTY	batchcount	Number of records to process between commit calls. Default value is 1000.
PROPERTY	jdbcconnectionstring property	<p>Provide RAC URL to connect to RAC database. Use this property only if you want to connect to RAC database.</p> <p>Example value:</p> <pre>jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = TCP) (HOST = hostname1) (PORT = 1521)) (ADDRESS = (PROTOCOL = TCP) (HOST = hostname2) (PORT = 1521)) (LOAD_BALANCE = yes )) (CONNECT_DATA = (SERVICE_NAME = serviceName)))</pre>
INPUT		The input dataset name.

## PREPAREDSTATEMENT

Element Name	Name	Value
PROPERTY	dbname	Name of the Oracle database.
PROPERTY	sid	Alias for dbname.
PROPERTY	dbuserid	username/password
PROPERTY	masdescriptors	Maximum number of columns allowed to be written to the database.
PROPERTY	schemaowner	The database owner of the table.
PROPERTY	tablename	Name of the Oracle table to which the data is written.
PROPERTY	table	Alias for tablename.
PROPERTY	hostname	<p>Hostname or IP address</p> <p>This is an optional property. The fully-specified hostname or IP address where the database resides. This defaults to the localhost if not specified.</p> <p><b>Note:</b> This property should only be specified when it is known that connections are being made to a remote database.</p> <p>This can be specified as a default in rfx.conf for convenience.</p>
PROPERTY	port	<p>This is a required property. The port on which the database resides. This defaults to 1521 if not specified in an operator or in rfx.conf.</p> <p><b>Note:</b> There are preset defaults in rfx.conf for the necessary port numbers. The default for Oracle is 1521. The Oracle utility tnsping can be used to obtain the port number.</p> <p>Verify with your database administrator the proper port for your database installation.</p>
PROPERTY	statement	SQL to execute for each input record. Use ? to denote a value to be replaced by a field from the current record.
PROPERTY	fields	Space separated list of fields to use as the values for the ? placeholders in the statement property.
PROPERTY	batchcount	Number of records to process between commit calls. Default value is 1000.
PROPERTY	jdbcconnectionstring property	<p>Provide RAC URL to connect to RAC database. Use this property only if you want to connect to RAC database.</p> <p>Example value:</p> <pre>jdbc:oracle:thin:@(DESCRIPTION =(ADDRESS_LIST =(ADDRESS = (PROTOCOL = TCP) (HOST = hostname1) (PORT = 1521)) (ADDRESS = (PROTOCOL = TCP) (HOST = hostname2) (PORT = 1521)) (LOAD_ BALANCE = yes )) (CONNECT_DATA =(SERVICE_NAME = serviceName)))</pre>
INPUT		The input dataset name.

## Database Operator Examples

The following are examples of database operators.

### ORAREAD

Examples are provided below for non RAC and RAC database through THIN driver.

#### Connection to non RAC database through THIN database driver:

```
<OPERATOR type="oraread">
  <PROPERTY name="sp_prequery" value="exec pre_storedproc"/>
  <PROPERTY name="dbname" value="RETLdb"/>
  <PROPERTY name="connectstring" value="RETLdb/rpassword"/>
  <!--Note: query must be enclosed in CDATA element otherwise -->
  <!-- this query will contain invalid XML! -->
  <PROPERTY name="query">
    <![CDATA[
      select * from rtbl where col > 1
    ]]>
  </PROPERTY>
  <PROPERTY name="maxdescriptors" value="100"/>
  <PROPERTY name="datetotimestamp" value="false"/>
  <PROPERTY name="sp_postquery" value="exec post_storedproc"/>
  <OUTPUT name="test.v"/>
</OPERATOR>
```

#### Connection to RAC database through THIN database driver:

```
<OPERATOR type="oraread">
  <PROPERTY name="dbname" value="serviceName"/>
  <PROPERTY name="dbuseralias" value="dbalias1"/>
  <PROPERTY name="port" value="1521" />
  <PROPERTY name="hostname" value="hostname1.us.oracle.com" />
  <PROPERTY name="maxdescriptors" value="100"/>
  <PROPERTY name="jdbcconnectionstring" value="jdbc:oracle:thin:@(DESCRIPTION
=(ADDRESS_LIST =(ADDRESS = (PROTOCOL = TCP)(HOST = hostname1)(PORT =
1521))(ADDRESS = (PROTOCOL = TCP)(HOST = hostname2)(PORT = 1521))(LOAD_BALANCE =
yes ))(CONNECT_DATA =(SERVICE_NAME = serviceName)))/>
  <PROPERTY name="query" value="SELECT * from tableName"/>
  <OUTPUT name="import.v"/>
</OPERATOR>
```

### ORAWRITE

Examples are provided below for non RAC and RAC database through THIN driver.

#### Connection to non RAC database through THIN database driver:

Connection to non RAC database through THIN database driver:

```
<OPERATOR type="orawrite">
<INPUT name="import0.v" />
<PROPERTY name="threadModel" value="start_thread" />
<PROPERTY name="dbuserid" value="username/password" />
<PROPERTY name="dbname" value="RETLdb" />
<PROPERTY name="tablename" value="ORG_LOC_DM" />
<PROPERTY name="createtablemode" value="recreate" />
</OPERATOR>
```

**Connection to RAC database through THIN database driver:**

```

<OPERATOR type="orawrite">
<INPUT name="import.v"/>
<PROPERTY name="hostname" value="hostname1.us.oracle.com" />
<PROPERTY name="port" value="1521" />
<PROPERTY name="dbname" value="serviceName"/>
<PROPERTY name="dbuseralias" value="dbalias1"/>
  <PROPERTY name="jdbcconnectionstring" value="jdbc:oracle:thin:@(DESCRIPTION
=(ADDRESS_LIST =(ADDRESS = (PROTOCOL = TCP)(HOST = hostname1)(PORT =
1521))(ADDRESS = (PROTOCOL = TCP)(HOST = hostname2)(PORT = 1521))(LOAD_BALANCE =
yes ))(CONNECT_DATA =(SERVICE_NAME = serviceName)))/>
<PROPERTY name="createtablemode" value="recreate"/>
<PROPERTY name = "tablename" value = "tableName8"/>
</OPERATOR>

```

---



---

**Note:** For more information, see "[Appendix: Database Connections Quick Reference](#)."

---



---

**UPDATE**

Examples are provided below for non RAC and RAC database through THIN driver.

**Connection to non RAC database through THIN database driver:**

For a dataset with records reflecting new employee salaries with the following fields:

- EMPLOYEE\_ID
- SALARY
- EFFECTIVE\_DATE

To update a table named EMPLOYEE\_SALARY with the new values using EMPLOYEE\_ID as the key field, use an UPDATE operator similar to the following:

```

<OPERATOR type="update">
  <INPUT name="salary_updates.v"/>
  <PROPERTY name="dbname" value="dbname"/>
  <PROPERTY name="schemaowner" value="schemaowner"/>
  <PROPERTY name="dbuserid" value="username/password"/>
  <PROPERTY name="tablename" value="EMPLOYEE_SALARY"/>
  <PROPERTY name="fields" value="EMPLOYEE_ID"/>
</OPERATOR>

```

The SQL executed is:

```

UPDATE EMPLOYEE_SALARY SET SALARY = salary, EFFECTIVE_DATE = effective_date WHERE
EMPLOYEE_ID = employee_id

```

salary, effective\_date, and employee\_id are the values of the SALARY, EFFECTIVE\_DATE, and EMPLOYEE\_ID fields of the current record being processed by RETL.

**Connection to RAC database through THIN database driver:**

```

<OPERATOR type="update">
<INPUT name="import.v"/>
<PROPERTY name="hostname" value="hostname1.us.oracle.com" />
<PROPERTY name="port" value="1521" />
<PROPERTY name="dbname" value="serviceName"/>
<PROPERTY name="dbuseralias" value="dbalias1"/>
<PROPERTY name="fields" value="eno"/>

```

```

<PROPERTY name="tablename" value="TABLENAME"/>
<PROPERTY name="jdbcconnectionstring" value="jdbc:oracle:thin:@(DESCRIPTION
=(ADDRESS_LIST =(ADDRESS = (PROTOCOL = TCP)(HOST = hostname1)(PORT =
1521))(ADDRESS = (PROTOCOL = TCP)(HOST = hostname2)(PORT = 1521))(LOAD_BALANCE =
yes ))(CONNECT_DATA =(SERVICE_NAME = serviceName)))/>
</OPERATOR>

```

## DELETE

Examples are provided below for non RAC and RAC database through THIN driver.

### Connection to non RAC database through THIN database driver:

For a dataset with information on stores that are being closed keyed on STORE\_ID and AREA, and you want to delete the corresponding records in the STORES table. Use a DELETE operator similar to the following:

```

<OPERATOR type="delete">
  <INPUT name="closing_stores.v"/>
  <PROPERTY name="dbname" value="dbname"/>
  <PROPERTY name="schemaowner" value="schemaowner"/>
  <PROPERTY name="dbuserid" value="username/password"/>
  <PROPERTY name="tablename" value="STORES"/>
  <PROPERTY name="fields" value="STORE_ID AREA"/>
</OPERATOR>

```

The SQL executed is:

```
DELETE FROM STORES WHERE STORE_ID = store_id AND AREA = area
```

store\_id and area are the values of the STORE\_ID and AREA fields of the current record being processed by RETL.

### Connection to RAC database through THIN database driver:

```

<OPERATOR type="delete">
<INPUT name="import.v"/>
<PROPERTY name="dbname" value="serviceName"/>
<PROPERTY name="hostname" value="hostname1.us.oracle.com" />
  <PROPERTY name="port" value="1521" />
  <PROPERTY name="dbuseralias" value="dbalias1"/>
  <PROPERTY name="jdbcconnectionstring" value="jdbc:oracle:thin:@(DESCRIPTION
=(ADDRESS_LIST =(ADDRESS = (PROTOCOL = TCP)(HOST = hostname1)(PORT =
1521))(ADDRESS = (PROTOCOL = TCP)(HOST = hostname2)(PORT = 1521))(LOAD_BALANCE =
yes ))(CONNECT_DATA =(SERVICE_NAME = serviceName)))/>
  <PROPERTY name="tablename" value="tableName"/>
  <PROPERTY name="fields" value="ename"/>
</OPERATOR>

```

## INSERT

Examples are provided below for non RAC and RAC database through THIN driver.

### Connection to non RAC database through THIN database driver:

For a dataset with records containing purchase order information that you need inserted into the PURCHASE\_ORDER table. Use an INSERT operator similar to the following:

```
<OPERATOR type="insert">
  <INPUT name="purchase_orders.v"/>
  <PROPERTY name="dbname" value="dbname"/>
  <PROPERTY name="schemaowner" value="schemaowner"/>
  <PROPERTY name="dbuserid" value="username/password"/>
  <PROPERTY name="tablename" value="PURCHASE_ORDERS"/>
</OPERATOR>
```

The SQL executed is as follows:

```
INSERT INTO PURCHASE_ORDERS (PO_NUMBER, ORDER_DATE, BILL_TO_ADDRESS, COMMENT, ...)
VALUES (po_number, order_date, bill_to_address, comment, ...)
```

po\_number, order\_date, bill\_to\_address, and comment are the values of the PO\_NUMBER, ORDER\_DATE, BILL\_TO\_ADDRESS, and COMMENT fields of the current record being processed by RETL.

To leave out fields, specify only the fields you want in the fields property:

```
<OPERATOR type="insert">
  <INPUT name="purchase_orders.v"/>
  <PROPERTY name="dbname" value="dbname"/>
  <PROPERTY name="schemaowner" value="schemaowner"/>
  <PROPERTY name="dbuserid" value="username/password"/>
  <PROPERTY name="tablename" value="PURCHASE_ORDER_DATES"/>
  <PROPERTY name="fields" value="PURCHASE_ORDER ORDER_DATE"/>
</OPERATOR>
```

The SQL executed is as follows:

```
INSERT INTO PURCHASE_ORDER_DATES (PO_NUMBER, ORDER_DATE) VALUES (po_number, order_date)
```

### Connection to RAC database through THIN database driver:

```
<OPERATOR type="insert">
<INPUT name="import.v"/>
<PROPERTY name="hostname" value="hostname1.us.oracle.com" />
<PROPERTY name="port" value="1521" />
<PROPERTY name="dbname" value="serviceName"/>
<PROPERTY name="dbuseralias" value="dbalias1"/>
<PROPERTY name="jdbcconnectionstring" value="jdbc:oracle:thin:@(DESCRIPTION
=(ADDRESS_LIST =(ADDRESS = (PROTOCOL = TCP)(HOST = hostname1)(PORT =
1521))(ADDRESS = (PROTOCOL = TCP)(HOST = hostname2)(PORT = 1521))(LOAD_BALANCE =
yes ))(CONNECT_DATA =(SERVICE_NAME = serviceName)))"/>
<PROPERTY name="tablename" value="TABLENAME"/>
</OPERATOR>
```

## PREPAREDSTATEMENT

Examples are provided below for non RAC and RAC database through THIN driver.

### Connection to non RAC database through THIN database driver:

The example for UPDATE above can be performed by a PREPAREDSTATEMENT similar to the following:

```
<OPERATOR type="preparedstatement">
  <INPUT name="salary_updates.v"/>
  <PROPERTY name="dbname" value="dbname"/>
  <PROPERTY name="schemaowner" value="schemaowner"/>
  <PROPERTY name="dbuserid" value="username/password"/>
  <PROPERTY name="tablename" value="EMPLOYEE_SALARY"/>
  <PROPERTY name="statement"
    <![CDATA[
      UPDATE EMPLOYEE_SALARY
      SET SALARY = ?,
      EFFECTIVE_DATE = ?
      WHERE EMPLOYEE_ID = ?
    ]]>
  <PROPERTY name="fields" value="EMPLOYEE_ID SALARY EFFECTIVE_DATE"/>
</OPERATOR>
```

The preceding example for DELETE can be performed by a PREPAREDSTATEMENT similar to the following:

```
<OPERATOR type="preparedstatement">
  <INPUT name="closing_stores.v"/>
  <PROPERTY name="dbname" value="dbname"/>
  <PROPERTY name="schemaowner" value="schemaowner"/>
  <PROPERTY name="dbuserid" value="username/password"/>
  <PROPERTY name="tablename" value="STORES"/>
  <PROPERTY name="statement" value="DELETE FROM STORES WHERE STORE_ID = ? AND
  AREA = ?"/>
  <PROPERTY name="fields" value="STORE_ID AREA"/>
</OPERATOR>
```

The example above for INSERT can be performed by a PREPAREDSTATEMENT similar to the following:

```
<OPERATOR type="preparedstatement">
  <INPUT name="purchase_orders.v"/>
  <PROPERTY name="dbname" value="dbname"/>
  <PROPERTY name="schemaowner" value="schemaowner"/>
  <PROPERTY name="dbuserid" value="username/password"/>
  <PROPERTY name="tablename" value="PURCHASE_ORDERS"/>
  <PROPERTY name="statement">
    <![CDATA[
      INSERT INTO PURCHASE_ORDERS
      (PO_NUMBER, ORDER_DATE, BILL_TO_ADDRESS, COMMENT)
      VALUES (?, ?, ?, ?)
    ]]>
  </PROPERTY>
  <PROPERTY name="fields" value="PO_NUMBER ORDER_DATE BILL_TO_ADDRESS COMMENT"/>
</OPERATOR>
```

**Connection to RAC database through THIN database driver:**

The example for PREPAREDSTATEMENT above can be performed by a PREPAREDSTATEMENT similar to the following:

```
<OPERATOR type="preparedstatement">
<INPUT name="import.v"/>
<PROPERTY name="hostname" value="hostname1.us.oracle.com" />
<PROPERTY name="port" value="1521" />
<PROPERTY name="dbname" value="serviceName"/>
<PROPERTY name="dbuseralias" value="dbalias1"/>
  <PROPERTY name="jdbcconnectionstring" value="jdbc:oracle:thin:@(DESCRIPTION
=(ADDRESS_LIST =(ADDRESS = (PROTOCOL = TCP)(HOST = hostname1)(PORT =
1521))(ADDRESS = (PROTOCOL = TCP)(HOST = hostname2)(PORT = 1521))(LOAD_BALANCE =
yes ))(CONNECT_DATA =(SERVICE_NAME = serviceName)))/>
<PROPERTY name="tablename" value="TABLENAME"/>
<PROPERTY name="statement" value="DELETE FROM TABLENAME WHERE ENO = ? AND
ENAME = ?"/>
<PROPERTY name="fields" value="ENO ENAME"/>
</OPERATOR>
```

---

---

## RETL Parallel Processing

Parallel processing is the ability to break divisible tasks into granular units of work, and to execute those units of work concurrently and independently of one another. The goal of parallelism is to achieve greater speed by executing tasks concurrently. RETL uses parallelism in these ways:

- Pipeline parallelism
- Framework parallelism and data partitioning

### Pipeline Parallelism

Pipeline parallelism can be demonstrated in manufacturing by an assembly line, where each worker can operate on tasks independently of other workers. At each stage in the assembly line, the worker is assigned a task and completes this task on the evolving product. Many tasks can be completed simultaneously (after each worker has product to work on). This is similar RETL pipeline parallelism. RETL operators effectively act as workers on an assembly line, whose goal is to produce records of a specific format. Each operator has its own assignment (for example, to sort, join, or merge), and they execute these tasks independently of other operators' tasks.

### Framework Parallelism

While pipeline parallelism is a means of increasing throughput, it only works to the point at which each worker reaches capacity. Some workers may be slower at their tasks than others, resulting in bottlenecks in the process. In the manufacturing world, a way of solving this dilemma is to increase the number of assembly lines. Now each assembly line can work independently of other assembly lines, uninhibited by bottlenecks. This is similar to RETL's framework-based parallelism. RETL implements framework-based parallelism in a concept called data partitioning. Data partitioning splits data from the source and passes each chunk of data to separate and concurrent pipelines.

Where can bottlenecks arise? In the manufacturing example, there may be only one inlet for transports to bring in raw materials, or one outlet for transports to ship out final product. These both limit the productivity of all assembly lines. The same can be said for RETL's processing. Input data being read from databases or files, and output data being written to the same, need to be parallelized to fully maximize data partitioning and be uninhibited by task bottlenecks. Starting with RETL release 11.2 and higher, RETL supports this additional enhancement to parallelize input and output tasks, by multithreading file and database reads and writes.

RETL uses both framework data partitioning and pipeline parallelism to achieve overall speed improvements in throughput.

Future versions of RETL will include support for Massively Parallel Processing (MPP), which maximizes processing on multiple interconnected machines. In the manufacturing example, this would be represented by a series of interconnected factories, each working simultaneously to produce a finished product.

## RETL Data Partitioning

This section assumes that you are familiar with the graphical representation of a flow. (See [Chapter 3, "RETL Program Flow"](#) for more information.)

RETL data partitioning, or simply partitioning, is the process RETL uses to increase the number of operators that are processing records, in an attempt to increase overall throughput. Partitioning consists of splitting the datasets into subsets, and processing each subset with additional pipelines constructed from the originals defined in the flow.

Partitioning does not happen automatically in RETL; the flow must be configured to tell RETL how to split up the data into subsets, and partitioning must be enabled.

### Enabling RETL Data Partitioning

You enable partitioning in one of two ways:

- By specifying a value greater than 1 for the `numpartitions` attribute of the `NODE` element in the RETL configuration file. This method enables partitioning for all invocations of RETL that use the configuration file.
- By specifying a value greater than 1 for the `-n` command line parameter. This method enables partitioning for just the single invocation of RETL and overrides the setting in the configuration file.

### Partition Construction

The partitioned data and the duplicated pipelines that process the partitioned data together make up a partition.

Partitions are created by RETL when one or more partitioning operators are included in the flow and configured for partitioning. A partitioning operator, also known as a partitioner, is an operator that splits the data into subsets. The operators that partition data are `IMPORT`, `DBREAD`, `HASH`, `SPLITTER`, and `ORAREAD`.

The partition pipelines are duplicated automatically by RETL from the operators specified in the flow. For example, if the original flow has a `BINOP` operator, multiple copies of the `BINOP` operator are created, one for each partition.

A partition pipeline ends when an operator is encountered that does not support partitioning, or does not support the type of partitioning performed by the partitioner. If necessary, RETL inserts a funnel in order to "unpartition" the data subsets back into a single dataset.


For a simple example, suppose you have a text file that contains records, and you need to add two fields and write the sum to another file. If you wanted to use data partitioning, the flow would look like this example:

```

<FLOW name="sumvalues">
  <!-- Read in the values. -->
  <OPERATOR type="import">
    <PROPERTY name="inputfile" value="values.txt"/>
    <PROPERTY name="schemafile" value="values-schema.xml"/>
    <OUTPUT name="import.v"/>
  </OPERATOR>
  <!-- Partition into 3 partitions. -->
  <OPERATOR type="splitter">
    <INPUT name="import.v"/>
    <PROPERTY name="numpartitions" value="3"/>
    <OUTPUT name="splitter.v"/>
  </OPERATOR>
  <!-- Add the values. -->
  <OPERATOR type="binop">
    <INPUT name="splitter.v"/>
    <PROPERTY name="left" value="VALUE1"/>
    <PROPERTY name="right" value="VALUE2"/>
    <PROPERTY name="operator" value="+"/>
    <PROPERTY name="dest" value="SUM"/>
    <OUTPUT name="binop.v"/>
  </OPERATOR>
  <!--Keep only the SUM field. -->
  <OPERATOR type="fieldmod">
    <INPUT name="binop.v"/>
    <PROPERTY name="keep" value="SUM"/>
    <OUTPUT name="fieldmod.v"/>
  </OPERATOR>
  <!-- Write the sum to sum.txt. -->
  <OPERATOR type="export">
    <INPUT name="fieldmod.v"/>
    <PROPERTY name="outputfile" value="sum.txt"/>
    <PROPERTY name="schemafile" value="sum-schema.xml"/>
  </OPERATOR>
</FLOW>

```

The graphs of the unpartitioned and partitioned flow look like this:



In this example, the SPLITTER operator is the partitioner. It distributes the records read by the IMPORT operator across three partitions. The BINOP and FIELDMOD operators are duplicated across all three partitions by RETL. Finally, a funnel is inserted before the EXPORT operator, to gather the partitioned records back into one dataset.

## Partitioning Types

The partitioners provide two different types of partitioning: keyed and non-keyed.

### Keyed Partitioning

Keyed partitioning provides two guarantees:

- Records with the same key are processed in the same partition.
- The order of records is retained within the data subset. That is, any two records in a partition are in the same order as they were in the original dataset.

Keyed partitioning is required to partition operators that require sorted input, such as GROUPBY and CLIPROWS.

### Nonkeyed Partitioning

Unlike keyed partitioning, nonkeyed partitioning makes no guarantee about which partition will process a particular record. Thus, nonkeyed partitioning is inappropriate for any operator that requires sorted input.

## Partitioners

The following are the partitioners:

- HASH
- IMPORT
- SPLITTER
- DBREAD
- ORAREAD
- GENERATOR

The number of partitions created by a partitioner is specified in the `numpartitions` property. The `numpartitions` property is common to all of the partitioners, although some of the partitioners require specification of other properties.

---

**Note:** Data partitioning must be enabled for the `numpartitions` property to have any affect. If data partitioning is not enabled, the property is ignored.

---

### HASH


The HASH operator is the only partitioner that performs keyed partitioning. The key fields are specified in the `key` property. As each record is processed, the hash code of the key fields is calculated, and records with the same hash code are sent to the same partition.

If the `numpartitions` property is not specified in the HASH operator, RETL uses the value from the `-n` command line option, if specified. If the `-n` command line option is not specified, the value from the configuration file is used.

In the following example, the HASH operator splits the data into two partitions. RETL automatically duplicates the GROUPBY operator and inserts the funnel. (The EXPORT operator in this example is not configured to be partitionable.)

```
<FLOW name="hash-example">
  <OPERATOR type="import">
    <PROPERTY name="inputfile" value="input.txt"/>
    <PROPERTY name="schemafile" value="in-schema.xml"/>
    <OUTPUT name="import.v"/>
  </OPERATOR>
  <OPERATOR type="hash">
    <INPUT name="import.v"/>
    <PROPERTY name="numpartitions" value="2">
    <PROPERTY name="key" value="KEY_FIELD">
    <OUTPUT name="hash.v"/>
  </OPERATOR>
  <OPERATOR type="groupby">
    <INPUT name="hash.v"/>
    <PROPERTY name="key" value="KEY_FIELD">
    <PROPERTY name="reduce" value="AMOUNT">
    <PROPERTY name="sum" value="TOTAL_AMOUNT">
    <OUTPUT name="groupby.v"/>
  </OPERATOR>
  <OPERATOR type="export">
    <INPUT name="groupby.v"/>
    <PROPERTY name="outputfile" value="output.txt">
    <PROPERTY name="schemafile" value="out-schema.xml">
```

</OPERATOR>



## IMPORT

IMPORT partitioning performs nonkeyed partitioning. If one input file is specified, then roughly equal portions of the file are handled by each partition. If more than one file is specified, the number of partitions must equal the number of files, and each file is handled by one partition.

To enable IMPORT partitioning, the numpartitions property must be specified. Because the main purpose of the IMPORT operator is not data partitioning, the default number of partitions does not come from the configuration file or from the command line.

In the following example, three IMPORT operators are created by RETL to read records. RETL automatically duplicates the BINOP and FIELDMOD operators and inserts the FUNNEL. (The EXPORT operator in this example is not configured to be partitionable.)

```
<FLOW name="import-example">
  <OPERATOR type="import">
    <PROPERTY name="numpartitions" value="3">
    <PROPERTY name="inputfile" value="input.txt"/>
    <PROPERTY name="schemafile" value="in-schema.xml"/>
    <OUTPUT name="import.v"/>
  </OPERATOR>
  <OPERATOR type="fieldmod">
    <INPUT name="import.v"/>
    <PROPERTY name="drop" value="UNNEEDED_FIELD">
```


```

    <OUTPUT name="fieldmod.v" />
</OPERATOR>
<OPERATOR type="binop">
  <INPUT name="fieldmod.v" />
  <PROPERTY name="left" value="VALUE1" />
  <PROPERTY name="operator" value="+" />
  <PROPERTY name="right" value="VALUE2" />
  <PROPERTY name="dest" value="SUM" />
  <OUTPUT name="fieldmod.v" />
</OPERATOR>
<OPERATOR type="export">
  <INPUT name="fieldmod.v" />
  <PROPERTY name="outputfile" value="output.txt" />
  <PROPERTY name="schemafilename" value="out-schema.xml" />
</OPERATOR>

```

**Unpartitioned**

**Partitioned**



## SPLITTER

The SPLITTER operator performs nonkeyed partitioning by sending records to the different partitions in round-robin fashion.

---



---

**Note:** Although SPLITTER will be supported in future versions of RETL, it is not guaranteed that it will use the round-robin algorithm. Flows should not rely on SPLITTER partitioning records in a round robin manner.

---




---

If the numpartitions property is not specified in the SPLITTER operator, RETL uses the value from the -n command line option. If the -n command line option is not specified, the value from the configuration file is used.

In the following example, the SPLITTER operator splits the data into two partitions. RETL automatically duplicates the FIELDMOD and BINOP operators and inserts the FUNNEL. (The EXPORT operator in this example is not configured to be partitionable.)

```
<FLOW name="splitter-example">
  <OPERATOR type="import">
    <PROPERTY name="inputfile" value="input.txt"/>
    <PROPERTY name="schemafile" value="in-schema.xml"/>
    <OUTPUT name="import.v"/>
  </OPERATOR>
  <OPERATOR type="splitter">
    <INPUT name="import.v"/>
    <PROPERTY name="numpartitions" value="3">
    <OUTPUT name="splitter.v"/>
  </OPERATOR>
  <OPERATOR type="fieldmod">
    <INPUT name="splitter.v"/>
    <PROPERTY name="drop" value="UNNEEDED_FIELD">
    <OUTPUT name="fieldmod.v"/>
  </OPERATOR>
  <OPERATOR type="binop">
    <INPUT name="fieldmod.v"/>
    <PROPERTY name="left" value="VALUE1"/>
    <PROPERTY name="operator" value="+"/>
    <PROPERTY name="right" value="VALUE2"/>
    <PROPERTY name="dest" value="SUM"/>
    <OUTPUT name="fieldmod.v"/>
  </OPERATOR>
  <OPERATOR type="export">
    <INPUT name="fieldmod.v"/>
    <PROPERTY name="outputfile" value="output.txt">
    <PROPERTY name="schemafile" value="out-schema.xml">
  </OPERATOR>
```

**Unpartitioned****Partitioned****DBREAD, ORAREAD**

The database read partitioners implement nonkeyed partitioning by allowing specification of multiple queries. Each query result set provides the data for one partition.

The numpartitions property must be set to enable partitioning, and the value must equal the number of queries. If the numpartitions property is not specified, the result sets are funneled together into one dataset, instead of feeding into multiple partitions.

For example, suppose that you want to read all of the records from the table named sourcetable. The unpartitioned DBREAD operator would have a query property like this:

```
<PROPERTY name="query">
  <![CDATA[
    select * from sourcetable
  ]]>
</PROPERTY>
```

If you wanted to partition the database read into three, the flow would have something like this:

```
<PROPERTY name="numpartitions" value="3">
<PROPERTY name="query">
  <![CDATA[
    select * from sourcetable where keyfield between 0 and 100000
  ]]>
</PROPERTY>
```

```

<PROPERTY name="query">
  <![CDATA[
    select * from sourcetable where keyfield between 100001 and 200000
  ]]>
</PROPERTY>
<PROPERTY name="query">
  <![CDATA[
    select * from sourcetable where keyfield > 200000
  ]]>
</PROPERTY>

```

It is important to perform performance analysis on each of the queries. Adding a WHERE clause to the unpartitioned SQL query may change the execution plan and cause it to run many times slower, negating any time saved by partitioning.

Data partitioning with ORAREAD works particularly well with Oracle partitioned tables, when each query reads from a different table partition. Consult Oracle database documentation for information on partitioned tables.


The stored procedure specified by the `sp_prequery` property is run before any of the queries are executed. Likewise, the stored procedure specified by the `sp_postquery` property is run after the last query completes.

In the following example, two queries are specified, so RETL creates two ORAREAD operators to read records from two different tables in an Oracle database. RETL automatically duplicates the FIELDMOD and BINOP operators and inserts the funnel.

```

<FLOW name="oraread-example">
  <OPERATOR type="oraread">
    <!-- connection properties not shown -->
    <PROPERTY name="numpartitions" value="2">
    <PROPERTY name="query">
      <![CDATA[
        select * from table_a
      ]]>
    </PROPERTY>
    <PROPERTY name="query">
      <![CDATA[
        select * from table_b
      ]]>
    </PROPERTY>
    <OUTPUT name="oraread.v"/>
  </OPERATOR>
  <OPERATOR type="fieldmod">
    <INPUT name="splitter.v"/>
    <PROPERTY name="drop" value="UNNEEDED_FIELD">
    <OUTPUT name="fieldmod.v"/>
  </OPERATOR>
  <OPERATOR type="binop">
    <INPUT name="fieldmod.v"/>
    <PROPERTY name="left" value="VALUE1"/>
    <PROPERTY name="operator" value="+"/>
    <PROPERTY name="right" value="VALUE2"/>
    <PROPERTY name="dest" value="SUM"/>
    <OUTPUT name="fieldmod.v"/>
  </OPERATOR>
  <OPERATOR type="export">
    <INPUT name="fieldmod.v"/>
    <PROPERTY name="outputfile" value="output.txt">
    <PROPERTY name="schemafilename" value="out-schema.xml">
  </OPERATOR>

```



## Operators and Partitioning Types

When a flow uses an operator other than HASH to partition the data, RETL ends the partition before any operator that requires sorted input, such as CLIPROWS or GROUPBY. This is because only the HASH operator guarantees that records with the same key are handled in the correct order by the same partition.

RETL ends the partition by inserting a funnel. Because a funnel does not put records back into a sorted order, the operator that requires sorted input will display a warning message about not having sorted input, and the results will be incorrect.

There are several ways to fix this problem. The best solution depends on the specifics of the flow. In general:


- If a SPLITTER operator is used to partition the data, change it to a HASH.
- If an IMPORT operator is used to partition the data, change the IMPORT to be unpartitioned, and insert a HASH operator after the IMPORT.
- If a database read operator is used to partition the data, replace the partitioned database read operator with multiple database read operators, and insert a SORTFUNNEL and a HASH, or combine all of the queries into one and insert a HASH after the DBREAD. In either case, make sure that all queries return sorted records.
- If the operator requiring sorted order is a JOIN, consider using a LOOKUP operator instead. The LOOKUP operator does not require sorted input. Keep in mind that the LOOKUP operator requires more memory than a JOIN and is only appropriate for small lookup table sizes.

- If the operator requiring sorted order is CHANGECAPTURE, consider using a CHANGECAPTURELOOKUP operator instead. The CHANGECAPTURELOOKUP operator does not require sorted input. Keep in mind that the CHANGECAPTURELOOKUP operator requires more memory than a CHANGECAPTURE and is only appropriate for small lookup table sizes.


As an example, suppose that a SPLITTER is used to partition imported data and a CLIPROWS operator is encountered. RETL inserts a FUNNEL before the CLIPROWS and displays the following warning message:

```
[cliprows:1]: Warning: Input to cliprows:1 does not seem to be sorted! Data should be sorted according to the proper keys or you may get unexpected results!
```

The flow graph looks like this:



To correct this flow, change the SPLITTER to a HASH:



The following operators require sorted input and support only keyed partitioning:

- CHANGECAPTURE
- CLIPROWS
- COMPARE
- DIFF
- FULL OUTERJOIN
- GROUPBY
- INNERJOIN
- LEFT OUTERJOIN
- MERGE (MERGE does not require sorted input, but it does require that records be processed in the same order, so HASH partitioning is required.)
- REMOVEDUP
- RIGHT OUTERJOIN

## Parallel Property

To be included in a partition, some operators require that the parallel property be set to "true". The parallel property is typically required of an operator in any of the following cases:

- Partitioning the operator may sometimes produce incorrect results, and so the default is not to partition the operator.
- Partitioning the operator may improve performance on some platforms and not others. For example, using a partitioned EXPORT may not always improve performance.
- The operator did not partition in earlier versions of RETL.

The following operators require the parallel property to be set to "true" to allow the operator to be partitioned:

- CLIPROWS
- DBWRITE
- DEBUG
- EXPORT
- GENERATOR
- GROUPBY
- ORAWRITE

If the parallel property is not specified, RETL ends the partition before the operator by inserting a funnel.

## Partitioned EXPORT

When the parallel property of an EXPORT operator is set to "true", RETL includes the EXPORT in the partition. The output from each partition is written to separate temporary files. When all partitions have completed processing, the temporary files are concatenated into the specified output file.

---

---

**Note:** Partitioning the EXPORT operator may or may not improve performance and should be tested before implementation.

---

---

The temporary files are created in the TEMPDIR directories specified in the configuration files. Best results are obtained when there is a temporary directory for each partition and each directory is on a separate disk controller.

The following graph is for a flow identical to the simple flow described at the beginning of this section, except that a partitioned EXPORT is specified. Note that two EXPORT operators are created but that there is only one file exported.

```
<FLOW>
  <!-- Other operators same as in original example. -->
  <!-- Write the sum to sum.txt. -->
  <OPERATOR type="export">
    <INPUT name="fieldmod.v"/>
    <PROPERTY name="parallel" value="true"/>
```

```


    <PROPERTY name="outputfile" value="sum.txt" />
    <PROPERTY name="schemafile" value="sum-schema.xml" />
  </OPERATOR>
</FLOW>

```

---

**Unpartitioned**
**Partitioned**


---



## Partitioned GENERATOR

The GENERATOR operator uses two partitioning-related attributes when generating a sequence field within a partition. Generally, these two attributes are used together to allow the sequence field to contain unique values across all partitions:


- When the `partnum_offset` property is set to true, the initial value is incremented by the zero-based partition number.
- When the `partnum_incr` property is set to true, the increment value is multiplied by the number of partitions.

For example, if there are three partitions, and the sequence field's initial value is specified as 1, and the increment value is specified as 1, then the first generator will generate values 1, 4, 7, ...; the second 2, 5, 8, ...; and the third 3, 6, 9, ...

A suggestion for debugging purposes is to set `init` to 0, `incr` to 0, `partnum_offset` to True, and `partnum_incr` to False. The generated field indicates the partition that processed the record.


## Partitioned LOOKUP and CHANGECAPTURELOOKUP

When both the source and lookup datasets of a LOOKUP or CHANGECAPTURELOOKUP operator are partitioned by a HASH operator, each partitioned LOOKUP or CHANGECAPTURELOOKUP operator uses a subset of the lookup dataset and a subset of the source dataset.



However, if the source dataset is partitioned by any other partitioner, or the lookup table is not partitioned at all, then the LOOKUP or CHANGECAPTURELOOKUP operators in each partition will share the same lookup table. A SHARE operator is inserted into the flow to accomplish the sharing and can be ignored. (The SHARE operator is a special operator used for LOOKUP partitioning.)

From the graph, it appears that the inserted HASH operator is partitioning the data; however, it actually is not. All of the lookup table records are sent to the first partition's LOOKUP or CHANGECAPTURELOOKUP operator, which loads the shared lookup table.



### Partitioned ORAWRITE

The stored procedure specified in the preload property is run before any partitioned ORAWRITE has started loading data. Likewise, the stored procedure specified in the postload property is run after the last partition completes.

### Flows with Multiple Partitioners

It is possible to use more than one partitioner in a flow. For example, both an IMPORT and a database read can be partitioned.


#### HASH Within a Partition

If a HASH is found within a partition, a funnel is automatically inserted after the HASH to gather records with the same keys into the same data subset.

**Note:** This will not leave the data sorted by the new key fields. You must enter the appropriate sort operator if the data needs to be sorted.

The following graph shows a flow that produces invalid results. The funnels inserted after the HASH operators lose the sorted order required by GROUPBY.

**Unpartitioned    Partitioned**



The following graph shows a corrected version of the above flow. Note that there is a sort in the unpartitioned flow immediately after the HASH.

**Unpartitioned    Partitioned**



### SPLITTER Within a SPLITTER Partition


If a SPLITTER is found within a partition, the partitioned data is repartitioned. For example, if the first partitioner specifies two partitions and the SPLITTER specifies three partitions, there are six partitions after the SPLITTER.

The following graph shows a flow with two partitions started by a SPLITTER expanding to six partitions.

---


#### Unpartitioned Partitioned

---



## Two HASH Operators Joined


The join operators (such as INNERJOIN, LEFTOUTERJOIN) support keyed partitioning. RETL ensures that records with identical keys are sent to the same partitioned join operator. If the HASH operators specify a different number of partitions, then the smaller numpartitions are used for both HASH operators.



### Funneled Partitions

The data from partitions started by multiple partitioners can be funneled using a single FUNNEL operator. RETL requires that both partitioners have the same number of partitions. This restriction will be lifted in a later release.

In the following example, an IMPORT creates two partitions and an ORAREAD creates two partitions. The four partitions are funneled together using the same FUNNEL.



### Data Partitioning Guidelines

Keep these guidelines in mind when designing your flow:

- Performance improvement is more likely when the inputs to the flow are partitioned than when a HASH or a SPLITTER is used to partition. The input to HASH or SPLITTER is serial, whereas IMPORT and database reads are parallel.
- Because of the expense in funneling, performance improvement is more likely when parallel database writes or EXPORT are performed than when the data is unpartitioned using a funnel or sortfunnel before the output.
- More partitions are not always better. More partitions require more threads, and there is a limit to the optimum number of threads used to process the flow. This limit depends on the hardware and the specifics of the flow. Thus, finding the optimum number of partitions can be a difficult task.

## Operator Configuration for Data Partitioning

The following table shows each operator that requires modifications to its properties to allow it to create or process partitioned data. If an operator is not listed, it either supports partitioned data without any configuration or it does not support data partitioning at all.

Operator	Property Name	Property Value
CLIPROWS	parallel	true
DBREAD	numpartitions	greater than 1
	query	Specify a query for each partition.
DBRWRITE	parallel	true
DEBUG	parallel	true
EXPORT	parallel	true
GENERATOR (generating fields for pre-existing records)	parallel	true
GENERATOR (generating new records)	numpartitions	greater than 1
GROUPBY	parallel	true
IMPORT	numpartitions	greater than 1
	inputfile	Specify either one input file or numpartitions input files.
ORAREAD	numpartitions	greater than 1
	query	Specify a query for each partition.
ORAWRITE	parallel	true
	mode	append  <b>Note:</b> Changing the mode property to "append" from "truncate" may require changes to the batch job, because the existing database rows will no longer be deleted by RETL.

Additionally, some operators only work correctly when keyed (hash) partitioning is used. For a description of keyed partitioning, see "[Partitioning Types](#)." The following operators require keyed partitioning:

- CHANGECAPTURE
- CLIPROWS
- COMPARE
- DIFF
- FULL OUTER JOIN
- GROUPBY

- INNERJOIN
- LEFTOUTERJOIN
- MERGE
- REMOVEDUP
- RIGHTOUTERJOIN

## A Final Word on Data Partitioning

Data partitioning does not necessarily correct performance problems in a flow. Just as performance tuning cannot significantly improve a poorly designed application, data partitioning cannot significantly improve a poorly designed flow.

Partitioning is not guaranteed to improve performance. In fact, carelessly adding partitioners likely degrades performance, because of the overhead of the partitioning, funneling, and additional threads.

If your flow is experiencing performance problems, do not look to partitioning first. Perform a thorough analysis of your flow to determine where the performance problem is taking place, and take appropriate steps. (See the RETL Performance Tuning Guide for more information.)

---

---

## Input and Output Operators

The following input and output operators are described in this chapter, along with tag usage example code:

- DEBUG
- NOOP
- EXPORT
- IMPORT

### DEBUG

The DEBUG operator prints all records to standard output.

### NOOP

The NOOP operator acts as a terminator for datasets that does little more than purge data.

### EXPORT

The EXPORT operator writes a RETL dataset to a flat file, either as delimited or fixed-length records, depending on the schema.

---

---

**Note:** By default, EXPORT exports records in pipe-delimited format (|). To change the output format, an export schema file should be used to specify a different format.

---

---

### IMPORT

The IMPORT operator imports a flat file into RETL, translating the data file into a RETL dataset.

## Input and Output Operators XML Specification Tables

This section describes the input and output operators XML specification tables.

### DEBUG

Element Name	Name	Value
INPUT		The input dataset name.
PROPERTY	parallel	true or false Optional property. When set to true, the DEBUG operator is included in a partition. When set to false, the DEBUG operator is not included in a partition and a funnel is inserted before the DEBUG operator to unpartition the data (default). Ignored if not partitioning.

### NOOP

Element Name	Name	Value
INPUT		The input dataset name.

### IMPORT

Element Name	Name	Value
PROPERTY	inputfile	The name of the text file to import into RETL. This property can be specified multiple times. If numpartitions is not specified or is set to 1, records are funneled into one dataset. Otherwise, a data partition is created for each input file. This property can be specified as one of the following: a filename in the current working directory, a filename with a relative path to the current working directory, or an absolute path to a file.
PROPERTY	schemafilename	The name of the schema file describing the layout of the data file. This property can be specified as one of the following: a filename in the current working directory, a filename with a relative path to the current working directory, or an absolute path to a file.
PROPERTY	rejectfile	Optional property. The name of the file where records that do not match the input schema are deposited. This property can be specified as one of the following: a filename in the current working directory, a filename with a relative path to the current working directory, or an absolute path to a file.
PROPERTY	allowedrejects	Optional property. A value greater than zero indicates how many errors are allowed before RETL exits with a critical error. <b>Note:</b> A zero indicates that any number of errors will be accepted. This is the default setting. The allowedrejects property is only valid when the rejectfile property is specified.

Element Name	Name	Value
PROPERTY	verboserejects	true or false  Optional property. A value of true means that RETL prints granular warning messages about rejected records to standard error. The default is true.  <b>Note:</b> Displaying rejected records to standard error will degrade performance, but setting verboserejects to false without specifying a rejectfile will result in lost records.
PROPERTY	degreesparallel	Optional property indicating the number of threads used to read each file. Defaults to 1. Increasing this value may increase performance.
PROPERTY	numpartitions	Optional property indicating the number of data partitions to create. Defaults to 1.  If multiple input files are specified and numpartitions is greater than 1, then numpartitions must equal the number of input files. One data partition is created for each input file.
OUTPUT		The output dataset name.

## EXPORT

Element Name	Name	Value
PROPERTY	outputfile	The output text file name. This property can be specified as one of the following: a filename in the current working directory, a filename with a relative path to the current working directory, or an absolute path to a file.
PROPERTY	outputmode	overwrite or append  Optional property that specifies whether EXPORT overwrites an existing file or appends to it. Default is to overwrite.
PROPERTY	schemafilename	The name of the schema file that describes the layout of the output file. If not specified, uses current schema (as output from the previous dataset). The default delimiter ( ) is used if no schema file is specified. This property can be specified as one of the following: a filename in the current working directory, a filename with a relative path to the current working directory, or an absolute path to a file.  <b>Note:</b> This property can be used to reorder and drop fields from the output and reformat the fixed length, delimiter, and nullvalue characteristics of the schema. However, if there are incompatible changes between the incoming schema and output schemafilename, RETL may print warning messages or throw errors.
PROPERTY	parallel	true or false  When set to true, each partition writes its data to a temporary file; then the temporary files are concatenated into the destination file.  When set to false, the EXPORT operator is not included in a partition, and a funnel is inserted before the EXPORT operator to unpartition the data. (default)  Optional property. Ignored if not partitioning.

Element Name	Name	Value
OUTPUT		The output dataset name.

## Input and Output Operators Examples

The following are input and output examples.

### IMPORT

```
<OPERATOR type="import" name="import1,0">
  <PROPERTY name="inputfile"          value="import.dat" />
  <PROPERTY name="schemafilename"     value="ImportOp.schema" />
  <OUTPUT name="test.v" />
</OPERATOR>

<OPERATOR type="import" name="import1,0">
  <PROPERTY name="inputfile"          value="import.dat" />
  <PROPERTY name="schemafilename"     value="ImportOp.schema" />
  <PROPERTY name="allowedrejects"     value="100" />
  <PROPERTY name="rejectfile"        value="import.rej" />
  <OUTPUT name="test.v" />
</OPERATOR>
```

### EXPORT

```
<OPERATOR type="export" name="export1,0">
  <PROPERTY name="outputfile"        value="output.dat" />
  <PROPERTY name="schemafilename"    value="outputOp.schema" />
  <INPUT name="test.v" />
</OPERATOR>
```

---

---

## Join Operators

RETL provides the following join operators, described in this chapter:

- INNERJOIN
- LEFTOUTERJOIN
- RIGHTOUTERJOIN
- FULLOUTERJOIN
- LOOKUP

---

---

**Note:** For INNERJOIN, LEFTOUTERJOIN, RIGHTOUTERJOIN, and FULLOUTERJOIN, the input datasets must be sorted on the key. LOOKUP does not require sorted input.

---

---

### INNERJOIN

INNERJOIN transfers records from both input datasets whose key fields contain equal values to the output dataset. Records with key fields that do not contain equal values are dropped.

### LEFTOUTERJOIN

LEFTOUTERJOIN transfers all values from the left dataset, and transfers values from the right dataset only where key fields match. The operator drops the key field from the right dataset. Otherwise, the operator writes default values.

### RIGHTOUTERJOIN

RIGHTOUTERJOIN transfers all values from the right dataset, and transfers values from the left dataset only where key fields match. The operator drops the key field from the left dataset. Otherwise, the operator writes default values.

### FULLOUTERJOIN

For records that contain key fields with identical and dissimilar content, the FULLOUTERJOIN operator transfers records from both input datasets to the output dataset.

## LOOKUP

LOOKUP is similar to the INNERJOIN operator. However, there is no need to do a sort on the input datasets. Looks up a value from one dataset whose key fields match the lookup dataset and outputs the matching values.

Remember that the lookup dataset must be small enough to fit in memory; otherwise severe performance problems may result. If in doubt as to whether the lookup dataset will fit in memory, always use one of the join operators.

## DBLOOKUP

DBLOOKUP behaves in a manner similar to the LOOKUP operator, but it is used to look up records directly in a database table. This is useful when a flow needs to join a relatively small number of records with a relatively large database table. In this sense, it behaves similar to a DBREAD with a LOOKUP. However, DBLOOKUP only pulls records from the database as needed. DBLOOKUP also attempts to maximize performance by caching database results and thus minimizing database accesses on key-matched records.

## Special Notes about Join Operators

Join operators favor data fields from the "dominant" side (usually, the first INPUT to the operator) when joining datasets that contain the same nonkey fields on each side. For example:

Left Dataset			Right Dataset		
Key	Same Nonkey Field	Nonkey field	Key	Same Nonkey Field	Nonkey field
A	B	C	A	B	C

Joined Dataset				
Operator	Key	Same Nonkey Field	Nonkey Field	Nonkey Field
INNERJOIN	A	B (from LHS)	C	D
LEFTOUTERJOIN	A	B (from LHS)	C	D
RIGHTOUTERJOIN	A	B (from RHS)	C	D
FULLOUTERJOIN	A	left_B, right_B (from LHS)	C	D
LOOKUP	A	B (from LHS)	C	D

It is important to note that FULLOUTERJOIN will rename the data fields that have the same name on each dataset, to left\_<shared column name> and right\_<shared column name>. It is recommended to drop the unneeded same-name fields from each dataset prior to joining, to make the functionality more explicit to those maintaining your flows.

## Join Operators XML Specification Tables

This section describes the join operators XML specification tables.

### INNERJOIN

Element Name	Name	Value
INPUT		Input dataset 1 - sorted on the key columns.
NPUT		Input dataset 2 - sorted on the key columns.
PROPERTY	key	Key columns to be joined.
OUTPUT		The output dataset name.

### LEFTOUTERJOIN, RIGHTOUTERJOIN, FULLOUTERJOIN

Element Name	Name	Value
NPUT		Input dataset 1. (left dataset) - sorted on the key columns
NPUT		Input dataset 2. (right dataset) - sorted on the key columns
PROPERTY	key	Key columns to be joined.
PROPERTY	nullvalue	column name=<null value> The default null value that is assigned to the null column.
OUTPUT		The output dataset name.

### LOOKUP

Element Name	Name	Value
NPUT		The first input dataset is always the data that needs to be processed.
NPUT		The second input dataset is always the lookup dataset.
PROPERTY	tablekeys	Comma separated key columns to be looked up.
PROPERTY	ifnotfound	reject, continue, drop, or fail What to do with the record if the result did not match. If not specified, this option defaults to fail.
PROPERTY	allowdups	true or false Optional property, defaults to true. This property allows the lookup to return more than 1 record if the lookup dataset has more than 1 matching set of keys.
OUTPUT		The first output is the successful lookup output dataset name.
OUTPUT		The second will contain all records not matching the tablekeys property. This needs to be specified if the ifnotfound option is set to reject.

## DBLOOKUP

Element Name	Name	Value
INPUT		The data to be processed.
PROPERTY	dbname	Required property. Name of the database.
PROPERTY	userid	Required property. Database log in name.
PROPERTY	dbtype	oracle or jdbc Required property. The type of database to which to connect.
PROPERTY	tablekeys	Required property. Comma separated key columns to be looked up in the database table or select_sql result set.
PROPERTY	table	The table to look up in. This property can be used generically instead of select_sql. Either table or select_sql must be specified.
PROPERTY	select_sql	Simple SQL SELECT statement used instead of table. Either table or select_sql must be specified. <b>Note:</b> select_sql can only contain a SELECT statement without the WHERE clause. RETL throws an error on an invalid query. Valid select_sql: SELECT colA,colB from table Invalid select_sql: SELECT colA,colB from table WHERE colA > colB
PROPERTY	hostname	Hostname or IP address This is an optional property. The fully-specified hostname or IP address where the database resides. This defaults to the localhost if not specified. <b>Note:</b> This property should only be specified when it is known that connections are being made to a remote database. This can be specified as a default in rfx.conf for convenience.
PROPERTY	port	The port on which the database resides. This defaults to 1521 if not specified in an operator or in rfx.conf. <b>Note:</b> There are preset defaults in rfx.conf for the necessary port numbers. The default for Oracle is 1521. The Oracle utility tnsping can be used to obtain the port number. Verify with your database administrator the proper port for your database installation.
PROPERTY	datetotimestamp	true or false This is an optional property that defaults to false. When set to true, ORAREAD returns DATE columns with the time included. When false, the time is not included.
PROPERTY	ifnotfound	reject, continue, drop, or fail Optional property, defaults to fail. What to do with the record if the result did not match.

Element Name	Name	Value
PROPERTY	allowdups	true or false Optional property, defaults to true. This property allows the lookup to return more than 1 record if the lookup dataset has more than 1 matching set of keys.
PROPERTY	usecache	yes or no Optional property, defaults to yes. Specifies whether the results from the database lookup should be cached.
OUTPUT		The first output is the successful lookup output dataset name.
OUTPUT		The second will contain all records not matching the tablekeys property. This needs to be specified if the ifnotfound option is set to reject.

## Join Operators Examples

The following are examples of join operators.

### INNERJOIN

```
<OPERATOR type="innerjoin" name="innerjoin,0">
  <INPUT name="test_1.v"/>
  <INPUT name="test_2.v"/>
  <PROPERTY name="key" value="LOC_KEY"/>
  <OUTPUT name="output.v"/>
</OPERATOR>
```

### LEFTOUTERJOIN

```
<OPERATOR type="leftouterjoin" name="leftouterjoin,0">
  <INPUT name="left.v"/>
  <INPUT name="right.v"/>
  <PROPERTY name="key" value="LOC_KEY"/>
  <PROPERTY name="nullvalue" value="PROD_SEASN_KEY=-1"/>
  <OUTPUT name="result.v"/>
</OPERATOR>
```

### RIGHTOUTERJOIN

```
<OPERATOR type="rightouterjoin" name="rightouterjoin,0">
  <INPUT name="right.v"/>
  <INPUT name="left.v"/>
  <PROPERTY name="key" value="LOC_KEY"/>
  <PROPERTY name="nullvalue" value="PROD_SEASN_KEY=-1"/>
  <OUTPUT name="output.v"/>
</OPERATOR>
```

## FULLOUTERJOIN

```
<OPERATOR type="fullouterjoin" name="fullouterjoin,0">
  <INPUT name="right.v" />
  <INPUT name="left.v" />
  <PROPERTY name="key" value="LOC_KEY" />
  <PROPERTY name="nullvalue" value="PROD_SEASN_KEY=-1" />
  <OUTPUT name="output.v" />
</OPERATOR>
```

## LOOKUP

```
<OPERATOR type="lookup">
  <INPUT name="dataset.v" />
  <INPUT name="lookupdataset.v" />
  <PROPERTY name="tablekeys" value="LOC_KEY,SUPP_KEY" />
  <PROPERTY name="ifnotfound" value="reject" />
  <OUTPUT name="result.v" />
  <OUTPUT name="reject.v" />
</OPERATOR>
```

```
<OPERATOR type="lookup">
  <PROPERTY name="tablekeys" value="LOC_KEY" />
  <PROPERTY name="ifnotfound" value="continue" />
  <PROPERTY name="allowdups" value="true" />
  <INPUT name="dataset.v" />
  <INPUT name="lookupdataset.v" />
  <OUTPUT name="result.v" />
</OPERATOR>
```

## DBLOOKUP

```
<OPERATOR type="dblookup">
  <INPUT name="dataset.v" />
  <PROPERTY name="table" value="mytab" />
  <PROPERTY name="tablekeys" value="D" />
  <PROPERTY name="ifnotfound" value="fail" />
  <PROPERTY name="hostname" value="dbhostname" />
  <PROPERTY name="userid" value="myuserid" />
  <PROPERTY name="password" value="mypassword" />
  <PROPERTY name="dbname" value="mydatabase" />
  <PROPERTY name="dbtype" value="oracle" />
  <OUTPUT name="joined.v" />
</OPERATOR>
```

---

---

## Mathematical Operators

RETL provides the following mathematical operators, described in this chapter:

- BINOP
- GROUPBY
- GROUPBY on multiple partitions

### BINOP

BINOP performs basic algebraic operations on two fields. Addition, subtraction, multiplication, and division are supported for all numeric types. The left and right operands can also be defined as constants. Note that division by zero produces a null value. Only addition is supported for string fields; this becomes concatenation. BINOP is not supported for date fields.

### GROUPBY

The input to the GROUPBY operator is a dataset to be summarized; the output is a dataset containing one record for each group in the input dataset. Each output record contains the fields that define the group and the output summaries calculated by the operator. One simple summary is a count of the number of records in each group. Another kind of summary is a statistical value calculated on a particular field for all records in a group.

---

---

**Note:** If all rows in a column are null, the GROUPBY operator provides a null column total. GROUPBY requires sorted input only when keys are specified. Otherwise, sorted input is not necessary.

---

---

### GROUPBY on Multiple Partitions

When running RETL using multiple partitions, run the GROUPBY operator with the parallel property set to "true" to get the performance increase from the multiple partitions. A HASH/SORT needs to be performed before the parallel GROUPBY. If the same KEYS are used to GROUPBY that were used for the HASH/SORT, data with the same KEY set is not spread to multiple partitions. If the next operator is a serial operator, there is an implied FUNNEL operator called by RETL after the GROUPBY to gather the data from the different partitions together into a single dataset.

If the next operator is a parallel operator, RETL does not call a FUNNEL operator, and the previous HASH/SORT is maintained. If you are not using the same KEYS to GROUPBY, a SORTFUNNEL is required to collect the data from the multiple partitions back into a single, sorted dataset. A final serial GROUPBY must then be performed, because data may have been spread across multiple partitions and not included in the GROUPBY calculation. If the next operator is a serial operator, nothing more needs to be done. If the next operator is a parallel operator, another HASH/SORT must be performed, because the HASH/SORT from the parallel GROUPBY is not maintained, because a serial GROUPBY was performed after that.

## Mathematical Operators XML Specification Tables

This section describes mathematical operators XML specification tables.

### BINOP

Element Name	Name	Value
INPUT		Input dataset name.
PROPERTY	left	Column name indicating the left operand. Only one of left and constleft can be specified.
PROPERTY	constleft	Constant value indicating the left operand. Only one of left and constleft can be specified.
PROPERTY	right	Column name indicating the right operand. Only one of right and constright can be specified.
PROPERTY	constright	Column name indicating the right operand. Only one of right and constright can be specified.
PROPERTY	dest	The result destination field name. This can be the same field name as was specified in left or right properties, to reuse existing fields.
PROPERTY	destype	Optional field. The type of the destination field. The default value is the type of the left field.
PROPERTY	operator	+ for addition - for subtraction * for multiplication / for division
OUTPUT		The output dataset name.

### GROUPBY

Element Name	Name	Value
INPUT		Input dataset name.
PROPERTY	parallel	true or false true—If running in multiple partitions. false—If running in a single partition (default).

Element Name	Name	Value
PROPERTY	key	Key columns to group by. If key is not specified, groupby considers all input data as one group. <b>Note:</b> If the key property is not specified, GROUPBY does not require sorted input.
PROPERTY	reduce	Column to be used in the calculation specified in the following min, max, sum, first, last, or count properties.  More than one operation (min and max) property can be specified after a reduce property. The column specified is used in all of the operations until the next reduce property is specified.
PROPERTY	min	Name of the new column that will hold the minimum of the values in the column specified in the preceding reduce property.
PROPERTY	max	Name of the new column that will hold the maximum of the values in the column specified in the preceding reduce property.
PROPERTY	sum	Name of the new column that will hold the sum of the values in the column specified in the preceding reduce property.
PROPERTY	first	Name of the new column that will hold the first value in the column specified in the preceding reduce property.
PROPERTY	last	Name of the new column that will hold the last value in the column specified in the preceding reduce property.
PROPERTY	count	Name of the new column that will hold the count of the values in the column specified in the preceding reduce property.
OUTPUT		The output dataset name.

## Mathematical Operators Examples

The following are examples of mathematical operators.

### BINOP

```
<OPERATOR type="binop">
  <INPUT name="import.v" />
  <PROPERTY name="left" value="field1"/>
  <PROPERTY name="operator" value="+"/>
  <PROPERTY name="constright" value="2"/>
  <PROPERTY name="dest" value="destfield"/>
  <OUTPUT name="binop.v" />
</OPERATOR>

<OPERATOR type="binop">
  <INPUT name="import.v" />
  <PROPERTY name="left" value="field1"/>
  <PROPERTY name="operator" value="-"/>
  <PROPERTY name="right" value="field2"/>
  <PROPERTY name="dest" value="destfield"/>
  <OUTPUT name="binop.v" />
</OPERATOR>
```

## GROUPBY

```

<OPERATOR type="groupby">
  <INPUT name="input_1.v"/>
  <PROPERTY name="parallel" value="true"/>
  <PROPERTY name="key" value="sex"/>
  <PROPERTY name="reduce" value="age"/>
  <PROPERTY name="min" value="min_age"/>
  <PROPERTY name="max" value="max_age"/>
  <PROPERTY name="first" value="first_age"/>
  <PROPERTY name="last" value="last_age"/>
  <PROPERTY name="count" value="count_age"/>
  <PROPERTY name="sum" value="sum_age"/>
  <PROPERTY name="reduce" value="birthday"/>
  <PROPERTY name="min" value="min_birthDay"/>
  <PROPERTY name="max" value="max_birthDay"/>
  <PROPERTY name="first" value="first_birthDay"/>
  <PROPERTY name="last" value="last_birthday"/>
  <PROPERTY name="count" value="count_birthday"/>
  <OUTPUT name="output.v"/>
</OPERATOR>

```

The following example shows what happens if the dataset was previously hashed and sorted on a different key from the group by key in the GROUPBY operator, and if the GROUPBY was run in multiple partitions.

In this case additional SORTCOLLECT and GROUPBY operators need to be used to guarantee the correct result.

```

<OPERATOR type="hash">
  <PROPERTY name="key" value="ZIP_CODE"/>
  <PROPERTY name="key" value="NAME"/>
  <OPERATOR type="sort">
    <PROPERTY name="key" value="ZIP_CODE"/>
    <PROPERTY name="key" value="NAME"/>
    <OUTPUT name="output1.v"/>
  </OPERATOR>
</OPERATOR>

<OPERATOR type="groupby">
  <INPUT name="output1.v"/>
  <PROPERTY name="parallel" value="true"/>
  <PROPERTY name="key" value="sex"/>
  <PROPERTY name="reduce" value="age"/>
  <PROPERTY name="count" value="count_age"/>
  <PROPERTY name="reduce" value="birthday"/>
  <PROPERTY name="max" value="max_birthday"/>
  <OPERATOR type="sortcollect">
    <PROPERTY name="key" value="sex"/>
    <OPERATOR type="groupby">
      <PROPERTY name="parallel" value="true"/>
      <PROPERTY name="key" value="sex"/>
      <PROPERTY name="reduce" value="age"/>
      <PROPERTY name="count" value="count_age"/>
      <PROPERTY name="reduce" value="birthday"/>
      <PROPERTY name="max" value="max_birthday"/>
      <OUTPUT name="output2.v"/>
    </OPERATOR>
  </OPERATOR>
</OPERATOR>

```

---

---

## Other Operators

This chapter describes these RETL operators and their usage:

- COMPARE
- SWITCH
- CHANGECAPTURE
- COPY
- DIFF
- CLIPROWS
- PARSER
- EXIT

### COMPARE

COMPARE performs a field-by-field comparison of records in two presorted input datasets. This operator compares the values of top-level, non-vector data types such as strings. All appropriate comparison parameters are supported (for example, case sensitivity and insensitivity for string comparisons). The assumption is that the original dataset is used as a basis for the comparison to the second dataset. The comparison results are recorded in the output dataset.

COMPARE appends a compare result to the beginning of each record with a value set as follows:

Code	Other
-1	First dataset FIELD is LESS THAN second dataset FIELD.
0	First dataset FIELD is the SAME AS second dataset FIELD.
-2	Record does not exist in the first dataset.
1	First dataset FIELD is GREATER THAN second dataset FIELD.
2	Record does not exist in the second dataset.

For example:

Original Key	Compare Key
A	A
B	A
C	E
D	

For this example, the output is as follows:

Key	Result	Compare Result
A	0	Keys match.
B	1	Original key is greater than compare key.
C	-1	Original key is less than compare key.
D	2	Original key does not exist in the compare dataset.

Note that the Compare Result column above is only for illustration. Only the Result column above would be added to the actual output dataset.

## SWITCH

SWITCH assigns each record of an input dataset to an output dataset, based on the value of a specified field.

## CHANGECAPTURE and CHANGECAPTURELOOKUP

Use CHANGECAPTURE or CHANGECAPTURELOOKUP to compare two datasets. The assumption is that the flow is comparing an original dataset with a modified copy (changed) dataset. The output is a description of what modifications have been made to the original dataset. Here are the changes that are captured:

- Inserts: Records added to the original dataset.
- Deletes: Records deleted from the original dataset.
- Edits: Records matching the given keys fields whose other fields have values have been modified from the original records.
- Copies: Records not changed from the original.

The operator works based on these conditions:

- Two inputs (the original and changed datasets), one output containing the combination of the two (edits are taken from the modified side) and a new field to show what change (if any) occurred between the original and the changed dataset.
- The two input datasets must have the same schemas
- For CHANGECAPTURE, the two datasets must be previously sorted on specified key fields. CHANGECAPTURELOOKUP does not require sorted input.

CHANGECAPTURE and CHANGECAPTURELOOKUP identify duplicate records between two datasets by comparing specified key fields. It then compares the value of the fields between the two datasets. The operator adds a field to the output records that contain one of four possible values, representing the conditions noted earlier (inserts, deletes, edits, or copies). The actual values assigned to each of the possibilities can be changed from the default if desired. It is possible to assign the name of this field and to indicate that records assigned into one or more of the above categories be filtered from the output dataset, rather than being passed along with the change code.

To eliminate the need for sorted input, CHANGECAPTURELOOKUP reads the entire changed dataset into memory into a lookup table. Thus, the memory requirements of CHANGECAPTURELOOKUP are much higher than for CHANGECAPTURE.

Here is an example to illustrate:

Original Dataset			Changed Dataset		
Key	Value	Other	Key	Value	Other
John	5	corn	John	5	corn
Jill	10	beans	George	7	olives
Allie	2	pizza	Allie	5	pizza
Frank	14	42fas	Frank	14	Bogo

The resulting OUTPUT would look like this:

Key	Value	Other	Change	CodeField
John	5	Corn	Copy	0
Jill	10	Beans	deletion	2
George	7	Olives	insertion	1
Allie	5	Pizza	Edit	3
Frank	14	Bogo	Copy	0

Note that the Change column is just for illustration. Only the CodeField would be added to the actual output dataset.

CHANGECAPTURE is often used in conjunction with the SWITCH operator.

## COPY

The COPY operator copies a single input dataset to one or more output datasets.

## DIFF

The DIFF operator performs a record-by-record comparison of two versions of the same dataset (the "before" and "after" datasets) and outputs one dataset that contains the difference between the compared datasets. DIFF is similar to CHANGECAPTURE; the only differences are the default change codes and default drops.

CHANGECAPTURE is the recommended operator to use for these types of operations, and DIFF is provided simply for backward compatibility.

## CLIPROWS

The CLIPROWS operator performs a record-by-record comparison of a sorted dataset to "clip" a number of rows from a group of records. For each group value of the input keys, it returns up to the first or last n entries in the group. It sends to its OUTPUT dataset the given records for further processing and discards the rest.

## PARSER

The PARSER operator allows business logic to be coded in a Java-like script. As a result, it is possible to:

- Encapsulate several operators' functionality into one operator (for example, BINOP, FILTER, COPY, NOOP)
- Improve performance by reducing the number of operators in the flow
- Improve maintenance costs by decreasing complexity of flows

Additionally, flows with the PARSER operator are more readily sizable with partitioning.

Supported parser constructs:

Conditional operators:

- ==, !=, >, !>, >=, !>=, <, <=, !<, !<=

Assignment operators:

- =, +=, -=, /=, \*=, %=

Mathematic operators:

- +, -, \*, /, %

Statements:

- if/else

Datatypes and representations:

- Strings: "test1", "-345"
- Longs: 1, 2, 1000, -10000
- Doubles: 1.0, 1e2
- Nulls: null, ""
- Fields: RECORD.FIELDNAME (for example, RECORD.ITEM\_LOC)  
(RECORD must be in all capitals.)
- Dataset output:
  - A record can be sent to the nth output by using RECORD.OUTPUT[n-1] = true;. (The OUTPUTs are indexed starting at 0, just like Java arrays.) Additionally, RECORD.OUTPUT[n-1] = false; allows the nth OUTPUT to be turned off.
  - RECORD.OUTPUT[\*] = true; sends the record to all defined OUTPUTs. While RECORD.OUTPUT[\*] = false; will turn all OUTPUTs off.

- A record can be conditionally sent to an OUTPUT, making PARSER act like a filter or switch operator.
- A record can be sent to more than one OUTPUT, making PARSER act like a copy operator. When this is done, the allowcopy property must be set to true.
- A record can be sent to no OUTPUT and consumed like NOOP. All OUTPUT settings must be set to false, individually or by using the \* designation.
- RECORD.OUTPUT[n] can be used as a BOOLEAN value in comparisons, so that you can check whether a RECORD is set to go to a specific OUTPUT.  

```
if ( RECORD.OUTPUT[0] == true ) { ...
```

#### Comparisons:

- Comparisons are made based upon the left-hand-side datatype.

#### Assignments:

- In an assignment, the right-hand side will be converted to the left-hand side datatype. An error will be thrown for invalid conversions.
- Strings can be concatenated with the + and += operators.
- Assignments must end with a semicolon.

#### Field usage:

- Fields can be used in assignments and comparisons by using RECORD.FIELDNAME to refer to the FIELDNAME field of the current record.
- Fields can only be used in an assignment to set other fields. For example, `-1 = RECORD.ITEM_LOC` is an invalid assignment.
- Multiple fields can be assigned or compared, but all fields come from the current record.

#### Code Blocks:

- All statements and expressions appear within a code block, which must be surrounded by curly braces {}. The braces are not optional for single line code blocks as they are in Java.

#### Comments:

- PARSER supports both single line comments starting with // and block comments delimited by /\* and \*/.

## EXIT

The EXIT operator aborts the RETL flow after a specified number of records have been processed by the operator.

## Other Operators XML Specifications

This section provides XML specifications for other operators.

### COPY

Element Name	Name	Value
INPUT		The input dataset name.
OUTPUT		The output dataset name. OUTPUT can be specified multiple times.

### COMPARE

Element Name	Name	Value
INPUT		Input dataset 1.
INPUT		Input dataset 2.
PROPERTY	key	Column name to be compared between the two datasets defined above.
OUTPUT	name.v	Output dataset.

### CLIPROWS

Element Name	Name	Value
INPUT		The input dataset name.
PROPERTY	key	Required property. Indicates that the given column is part of the group definition. Multiple keys can be specified. If you are using partitioning with CLIPROWS, make sure you are using a HASH operator with the same key values. That is, the cliprows keys and hash keys must be the same otherwise incorrect results will be returned.
PROPERTY	which	first or last Required property. Indicates that the given column is part of the group definition. Multiple keys can be specified.
PROPERTY	count	Required property. Indicates the number of rows to clip.
PROPERTY	parallel	true or false Optional property. When set to true, the CLIPROWS operator will be included in a partition. When set to false, the CLIPROWS operator will not be included in a partition, and a funnel will be inserted before the CLIPROWS operator to unpartition the data.(default). Ignored if not partitioning.
OUTPUT		The output dataset name.

## DIFF

Element Name	Name	Value
INPUT		First input dataset name. This dataset represents the original dataset.
INPUT		Second input dataset name. This dataset represents the modified dataset.
PROPERTY	key	Records that match key fields are considered to be the same when determining existence in one dataset versus the other. After the two records are matched via the keys basis, the operator can determine if there are copies or edits. If a record exists in the original dataset but not the changed dataset, it is considered to be a delete. Likewise, if a record exists in the changed dataset but not the original, it is an insertion. There can be multiple instances of this property, allowing the use of composite keys.
PROPERTY	allvalues	true or false true—Compares all the values between the datasets. false—Does not compare all values. If this property is false, then the value property should be set. <b>Note:</b> If the values property is not set when allvalues is false, DIFF acts as if allvalues is set to true.
PROPERTY	value	After the key property has been used to join datasets, the field specified by the value property is used to determine differences among the records in the joined datasets.  Must be present in both before and after schemas. Determines if the record is a copy or a delete. The fields are processed one by one in the order they are present in the record.
PROPERTY	codefield	The field name to set the change code field to. Defaults to codefield.
PROPERTY	copycode	The code that is written to the change code field for a copy. The default value is 2.
PROPERTY	editcode	The code that is written to the change code field for an edit. The default value is 3.
PROPERTY	deletecode	The code that is written to the change code field for a delete. The default value is 1.
PROPERTY	insertcode	The code that is written to the change code field for an insert. The default value is 0.
PROPERTY	dropcopy	true or false If true, copy records are dropped from the output. The default is false.
PROPERTY	dropedit	true or false If true, edit records are dropped from the output. The default is false.
PROPERTY	dropdelete	true or false If true, delete record are dropped from the output. The default is false.

Element Name	Name	Value
PROPERTY	dropinsert	true or false If true, insert records are dropped from the output. The default is false.
PROPERTY	sortascending	true or false Deprecated.
OUTPUT		The output dataset name.

## CHANGECAPTURE and CHANGECAPTURELOOKUP

Element Name	Name	Value
INPUT		First input dataset name. This dataset represents the original dataset.
INPUT		Second input dataset name. This dataset represents the modified dataset.
PROPERTY	allvalues	true or false true—Compares all the values between the datasets. false—Does not compare all values. If this property is false, then the value property should be set. <b>Note:</b> If the values property is not set when allvalues is false, CHANGECAPTURE and CHANGECAPTURELOOKUP act as if allvalues is set to true.
PROPERTY	key	Records that match key fields are considered to be the same when determining existence in one dataset versus the other. After the two records are matched via the keys basis, the operator can determine if there are copies or edits. If a record exists in the original dataset but not the changed dataset, it is considered to be a delete. Likewise, if a record exists in the changed dataset but not the original, it is an insertion. There can be multiple instances of this property, allowing the use of composite keys.
PROPERTY	value	This property indicates the fields to use to base the comparisons after records have been matched on key fields. When the given values match, the records are considered to be copies. When the values differ, the records are considered to be edits. There can be multiple instances of this property, allowing the use of composite values.  This is an optional property. When not specified, all common fields are used in the comparison. This is identical to setting the allvalues property to true.
PROPERTY	dropcopy	true or false If true, copy records are dropped from the output. The default is true.
PROPERTY	dropedit	true or false If true, edit records are dropped from the output. The default is false.

Element Name	Name	Value
PROPERTY	dropdelete	true or false If true, delete records are dropped from the output. The default is false.
PROPERTY	dropinsert	true or false If true, insert records are dropped from the output. The default is false.
PROPERTY	copycode	The code that is written to the change code field for a copy. The default value is 0.
PROPERTY	insertcode	The code that is written to the change code field for an insert. The default value is 1.
PROPERTY	deletecode	The code that is written to the change code field for a delete. The default value is 2.
PROPERTY	editcode	The code that is written to the change code field for an edit. The default value is 3.
PROPERTY	codefield	The field name to set the change code field to.
PROPERTY	sortascending	true or false Deprecated.
OUTPUT		The output dataset name.

## SWITCH

Element Name	Name	Value
INPUT		The input dataset name.
PROPERTY	switchfield	Field name to determine the switch.
PROPERTY	cases	value1=0, value2=1 This property assigns the switch value based on a comma-separated list of mappings from value to output dataset. <b>Note:</b> Output datasets are numbered starting from 0, starting from first specified to the last.
PROPERTY	discard	Value to be discarded or dropped.
PROPERTY	ifnotfound	allow, fail, ignore This property determines what to do with the record if the case values are not found. allow—The record is put on the last output dataset. If the number of output datasets matches the number of values in the cases property, then records that match the last case value and those that do not match any of the case values will share the same dataset. fail—Halts the process (default). ignore—Drops the record. The record is not put on any output dataset.

Element Name	Name	Value
INPUT		The valid output dataset. Number of outputs should correspond to number of casevalues. For example, value1=0, value2=1 should have two outputs (or three if discard is specified)
INPUT		The reject output dataset.

## PARSER

Element Name	Name	Value
INPUT		The input dataset name.
PROPERTY	expression	Language expression the operator should use to process records. Supports general Java-like syntax and must be surrounded by CDATA tags. See the examples that follow.
PROPERTY	allowcopy	true or false Set allowcopy to true if the expression can copy the same record to more than one OUTPUT.
OUTPUT		First output dataset name. At least one OUTPUT is required.
OUTPUT		Multiple OUTPUTs can be specified. To copy the current record to a specific OUTPUT, use RECORD.OUTPUT[n] = true; where n is between 0 and the number of OUTPUTs minus 1.

## EXIT

Element Name	Name	Value
INPUT		The input dataset name.
PROPERTY	records	When the number of records processed by the EXIT operator reaches the number of records specified in the records property, RETL aborts the flow. An error message is displayed indicating that the EXIT operator terminated the flow. The default value is 0, which is interpreted to mean that the EXIT operator should not terminate the flow.
OUTPUT		The output dataset name.

## Other Operators Examples

This section provides additional examples of operators.

### COPY

```
<OPERATOR type="copy">
  <INPUT name="input.v"/>
  <OUTPUT name="copy1.v"/>
  <OUTPUT name="copy2.v"/>
</OPERATOR>
```

## SWITCH

```
<OPERATOR type="switch">
  <INPUT name="import1.v" />
  <PROPERTY name="switchfield" value="COLOR" />
  <PROPERTY name="casevalues" value=" RED=0, BLUE=1" />
  <PROPERTY name="discard" value="YELLOW" />
  <PROPERTY name="ifnotfound" value="allow" />
  <OUTPUT name="red.v" />
  <OUTPUT name="blue.v" />
  <OUTPUT name="not_red_blue_or_yellow.v " />
</OPERATOR>
```

## COMPARE

```
<OPERATOR type="compare">
  <INPUT name="import1.v"/>
  <INPUT name="import2.v"/>
  <PROPERTY name="key" value="NAME"/>
  <OUTPUT name="compare.v"/>
</OPERATOR>
```

## CHANGECAPTURE

```
<OPERATOR type="changepcapture">
  <INPUT name="import1.v" />
  <INPUT name="import2.v" />
  <PROPERTY name="key" value="emp_age" />
  <PROPERTY name="value" value="emp_name" />
  <PROPERTY name="codefield" value="change_code" />
  <PROPERTY name="copycode" value="0" />
  <PROPERTY name="editcode" value="3" />
  <PROPERTY name="deletcode" value="2" />
  <PROPERTY name="insertcode" value="1" />
  <PROPERTY name="dropcopy" value="false" />
  <PROPERTY name="dropedit" value="true" />
  <PROPERTY name="dropdelete" value="true" />
  <PROPERTY name="dropinsert" value="true" />
  <PROPERTY name="allvalues" value="true" />
  <PROPERTY name="sortascending" value="false" />
  <OUTPUT name="changepcapture.v" />
</OPERATOR>
```

## CHANGECAPTURELOOKUP

```
<OPERATOR type="changepcapturelookup">
  <INPUT name="import1.v" />
  <INPUT name="import2.v" />
  <PROPERTY name="key" value="emp_age" />
  <PROPERTY name="value" value="emp_name" />
  <PROPERTY name="codefield" value="change_code" />
  <OUTPUT name="changepcapture.v" />
</OPERATOR>
```

## CLIPROWS

```
<OPERATOR type="cliprows">
  <INPUT name="import1.v"/>
  <PROPERTY name="key" value="LOC"/>
  <PROPERTY name="which" value="first"/>
  <PROPERTY name="count" value="2"/>
  <OUTPUT name="cliprows.v"/>
</OPERATOR>
```

## DIFF

```
<OPERATOR type="diff">
  <INPUT name="import1.v"/>
  <INPUT name="import2.v"/>
  <PROPERTY name="key" value="emp_name"/>
  <PROPERTY name="allvalues" value="true"/>
  <OUTPUT name="diff.v"/>
</OPERATOR>
```

## PARSER

```
<OPERATOR type="parser">
  <INPUT name="sales.v"/>
  <PROPERTY name="expression">
    <![CDATA[
      if (RECORD.SALES_AMT == null)
      {
        RECORD.SALES_AMT = 0;
      }
    ]]>
  </PROPERTY>
  <OUTPUT name="newsales.v"/>
</OPERATOR>
```

## EXIT

```
<!-- Validates that CODE > 0. Terminates the flow on the first record found that
has an invalid CODE. -->
<OPERATOR type="filter">
  <INPUT name="data.v"/>
  <PROPERTY name="filter" value="CODE GT 0"/>
  <PROPERTY name="rejects" value="true"/>
  <OUTPUT name="valid_codes.v"/>
  <OPERATOR type="exit">
    <PROPERTY records="1"/>
  </OPERATOR>
</OPERATOR>
```

## PARSER

The expression value must be surrounded by CDATA tags:

```
<PROPERTY name="expression">
  <![CDATA[
    if (RECORD.D > 500000)
    {
      RECORD.D *= 10;
    }
    else
    {
```

```

        RECORD.D *= 5;
    }
]]>
</PROPERTY>

```

Enclose date and string constants in double quotation marks:

```

<PROPERTY name="expression">
  <![CDATA[
    if (RECORD.A == "20020101")
    {
      RECORD.A = "20030101";
    }
    if (RECORD.B == "abc")
    {
      RECORD.B = "DWS";
    }
    else if (RECORD.B == "def")
    {
      RECORD.B = "CP";
    }
    else
    {
      if (RECORD.B == "hij")
      {
        RECORD.B = "2uy";
      }
      RECORD.B = "JC";
    }
  ]]>
</PROPERTY>

```

Example of ==, <, <=, >, >=, +=, -=, \*=, /=, and %= operators:

```

<PROPERTY name="expression">
  <![CDATA[
    if (RECORD.D == 1)
    {
      RECORD.D += 1;
    }

    if (RECORD.D < 10)
    {
      RECORD.D -= 2;
    }

    if (RECORD.D <= 1)
    {
      RECORD.D *= 3;
    }

    if (RECORD.D > 1)
    {
      RECORD.D /= 4;
    }

    if (RECORD.D >= 1)
    {
      RECORD.D %= 5;
    }
  ]]>

```

```

        if (RECORD.DAY_DT != "20030629")
        {
            RECORD.DAY_DT = "20030725";
        }
    ]]>
</PROPERTY>
Setting RECORD(n) sends the record to the n+1 OUTOUT, making PARSER act like a
FILTER or SWITCH operator.
<PROPERTY name="expression">
    <![CDATA[
        if (RECORD.A == 1)
        {
            RECORD.OUTPUT[0] = true;
        }
        else
        {
            RECORD.OUTPUT[1] = true;
        }
    ]]>
</PROPERTY>
Sending the record to multiple outputs makes PARSER act like a copy operator. The
allowcopy property must be set to "true" or you receive an error.
<PROPERTY name="allowcopy" value="true"/>
<PROPERTY name="expression">
    <![CDATA[
        // Either the code below
        RECORD.OUTPUT[0] = true;
        RECORD.OUTPUT[1] = true;

        // Or the simpler syntax in the commented code below
        //RECORD.OUTPUT[*] = true;
    ]]>
</PROPERTY>
PARSER can act like a filter and a copy at the same time. In this example, some
records are sent to the first OUTPUT and some are sent to the second OUTPUT. All
records are sent to the third OUTPUT.
<PROPERTY name="allowcopy" value="true"/>
<PROPERTY name="expression">
    <![CDATA[
        if (RECORD.A == 1)
        {
            RECORD.OUTPUT[0] = true;
        }
        else
        {
            RECORD.OUTPUT[1] = true;
        }
        RECORD.OUTPUT[2];
    ]]>
</PROPERTY>
PARSER can act like a FILTER, COPY, and a NOOP at the same time. In this example,
some records are sent to the first OUTPUT, some will be sent to both OUTPUTs, and
some are sent to no OUTPUTs.
<PROPERTY name="allowcopy" value="true"/>
<PROPERTY name="expression">
    <![CDATA[
        if (RECORD.A == 1)
        {
            RECORD.OUTPUT[0] = true;
        }
    ]]>

```

```
    else if (RECORD.A == 2)
    {
        RECORD.OUTPUT[*] = true;
    }
    else
    {
        RECORD.OUTPUT[*] = false;
    }
  ]]>
</PROPERTY>
```



---



---

## Structures and Data Manipulation Operators

RETL provides the following operators for the manipulation of structures and data, described in this chapter:

- CONVERT
- FIELDMOD
- FILTER
- GENERATOR
- REMOVEDUP

### CONVERT

The convert operator is used to convert an existing datatype to a new datatype from an input dataset. The following paragraph describes the syntax of the CONVERT property.

- The root tag <CONVERT>, containing one or more <CONVERTFUNCTION> or <TYPEPROPERTY> tags
- The <CONVERT> tag requires the following attributes:

Conversion Name	Description
destfield	Destination field name or new name
sourcefield	Original field name or old field name
newtype	New data type

- The <CONVERTFUNCTION> tag allows conversion of a column from one data type to a different data type and has the following attribute:

Conversion Name	Description
name	The value of the name attribute of the conversion function determines what conversion is done. The name generally follows the form: <typeTo>_from_<typeFrom>. Conversions are defined between all numeric values and from numbers to strings as well. The table below shows some example conversions.

- The <TYPEPROPERTY> tag allows turning the column into nullable or not null and has the following attributes:

Conversion Name	Description
name	nullable
value	true or false

## Conversion Functions

There are many conversion functions to allow conversion between types. Much of this is done with default conversions between types (see "[Appendix: Default Conversions](#)"). When using a default conversion, you simply specify 'default' as the name of the conversion.

Here are the other (nondefault) conversion functions:

Conversion Name	Description
make_not_nullable	Converts nullable field to not null. Requires specification of a functionarg tag. The value specified in the nullvalue functionarg is used to replace any null fields that are found. See " <a href="#">Structures and Data Manipulation Operators Examples</a> ."
make_nullable	Converts non-nullable field to nullable field. Requires specification of a functionarg tag. The value specified in the "nullvalue" functionarg is used as the field's null value. See " <a href="#">Structures and Data Manipulation Operators Examples</a> ."
string_length	Converts a string to a UINT32 with a value equivalent to the string's length.
string_from_int32	Converts to STRING from INT32.
string_from_int64	Converts to STRING from INT64.
string_from_int16	Converts to STRING from INT16.
string_from_int18	Converts to STRING from INT8.
int32_from_dfloat	Converts to INT32 from DFLOAT rounding the result to the nearest integer value.
int64_from_dfloat	Converts to INT64 from DFLOAT rounding the result to the nearest integer value.
string_from_date	Converts to STRING from DATE.
dfloat_from_string	Converts to DFLOAT from STRING.
string_from_dfloat	Converts to STRING from DFLOAT.

---



---

**Note:** RETL reports invalid conversions, but RETL does not handle certain conversions well. Take care to ensure that the data is well-formed by the time a conversion takes place. Overflow and underflow conversion errors are not detected by RETL. For example, converting the INT16 "-12" to a UINT16 will result in undefined behavior.

---



---

## FIELDMOD

FIELDMOD is used to remove, duplicate, drop and rename columns within RETL. There are several optional parameters that you can use in the fieldmod operator. Here are some notes on how to use the fieldmod operator:

- Use the keep property when you need to drop a large number of columns and/or ignore columns that you do not know about (allowing your flows to be more resilient to change). Any columns that are not specified in the keep property are dropped. The column names are separated by spaces. Multiple keep properties can be specified.
- If you want to keep most of columns and just drop a few columns, use the drop property. The columns to drop are separated by spaces. The drop property is processed after the keep property, which means that a column is dropped if it is specified on both the keep and drop properties. Multiple drop properties can be specified. Duplicated columns in single or multiple drop properties will generate an error (for example, "Delete of field failed (field 'Emp\_Age' not found)"), since the column has been dropped.
- The rename and duplicate properties affect the dataset after keep and drop properties have been processed. If you attempt to rename or duplicate a column that has been dropped, RETL generates an error. For example, "Error: the column Emp\_Name has been dropped and can not be duplicated."

## FILTER

FILTER is used to filter a dataset into two categories, those that meet the filter criterion and those that do not.

Possible uses are as follows:

- To filter records where a field must be equivalent to a certain string
- To filter on records where a certain field is less than 10, or filter on records where two fields are equal (somewhat like the WHERE clause of a SQL statement)

## GENERATOR

GENERATOR is used to create new datasets for testing or for combining with other data sources. It can act as a stand-alone operator that generates the dataset from scratch, or as an operator that can add fields to an existing dataset. The following paragraph describes the syntax of the SCHEMA property, which in the following example is represented within the CDATA tag.

The SCHEMA property should specify the XML that will generate the fields. The basic outline is:

- The root tag <GENERATE>, containing one or more <FIELD> tags
- The <FIELD> tag specifies the name and type of the field as attributes, and it can contain exactly one field generation tag. Field generation tags include:
  - <SEQUENCE> tag with the following attributes:

Name	Description
init	Starting number in sequence. The default is 0.
incr	Increment to add for each record in sequence. The default is 1.
limit	Max number, once reached, will start from beginning. Optional, unlimited if not specified.
partnum_offset	(true/false) Adds the partition number to the init value for parallel operation. The default is false.
partcount_incr	(true/false) Multiplies the incr value by the number of partitions. The default is false.

- The <CONST> tag allows the specification of a single attribute, value, which is the value that the field will take on.
- The <VALUELIST> tag has no attributes, but has multiple <CONST> tags within it, specifying the list of values that should be cycled through

The following are the only data types that can be used within the GENERATOR operator:

- int8
- int16
- int32
- int64
- dfloat
- string
- date

## REMOVEDUP

The REMOVEDUP operator performs a record-by-record comparison of a sorted dataset to remove duplicate records. Duplicates are determined based upon the key fields specified.

## Structures and Data Manipulation Operators XML Specification Tables

This section describes the structures and data manipulation operators XML specification tables.

### CONVERT

Element Name	Name	Value
INPUT		The input dataset name.
PROPERTY	convertspec	Describes the new column data structure to be converted. See the convert specification and the example for more detail.
OUTPUT		The output dataset name.

### FIELDMOD

Element Name	Name	Value
INPUT		The input dataset name.
PROPERTY	keep	Space-separated list of the columns to be kept. The named fields are retained and all others are dropped. This property can be specified multiple times. Optional
PROPERTY	drop	Space-separated list of the columns to be dropped. The named fields are dropped and all other fields are retained. This property can be specified multiple times. Optional
PROPERTY	rename	Column to be renamed. The format is: new_column_name=existing_column_name
PROPERTY	duplicate	Column to be duplicated. The format is: new_column=existing_column Duplicate one column to another column. The source column is separated by an equal sign from the target column name.
OUTPUT		The output dataset name.

## FILTER

Element Name	Name	Value																								
INPUT		The input dataset name.																								
PROPERTY	filter	Filter expression																								
		The following operations are supported. See the <a href="#">"Filter Expressions"</a> section for more detail.																								
		<table border="1"> <thead> <tr> <th>Operation</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>NOT</td> <td>Logical NOT</td> </tr> <tr> <td>AND</td> <td>Logical AND</td> </tr> <tr> <td>OR</td> <td>Logical OR</td> </tr> <tr> <td>IS_NULL</td> <td>True if field is null.</td> </tr> <tr> <td>IS_NOT_NULL</td> <td>True if field is not null.</td> </tr> <tr> <td>GE</td> <td>Logical &gt;=</td> </tr> <tr> <td>GT</td> <td>Logical &gt;</td> </tr> <tr> <td>LE</td> <td>Logical &lt;=</td> </tr> <tr> <td>LT</td> <td>Logical &lt;</td> </tr> <tr> <td>EQ</td> <td>True if right and left fields are equal.</td> </tr> <tr> <td>NE</td> <td>True if right and left fields are not equal.</td> </tr> </tbody> </table>	Operation	Description	NOT	Logical NOT	AND	Logical AND	OR	Logical OR	IS_NULL	True if field is null.	IS_NOT_NULL	True if field is not null.	GE	Logical >=	GT	Logical >	LE	Logical <=	LT	Logical <	EQ	True if right and left fields are equal.	NE	True if right and left fields are not equal.
Operation	Description																									
NOT	Logical NOT																									
AND	Logical AND																									
OR	Logical OR																									
IS_NULL	True if field is null.																									
IS_NOT_NULL	True if field is not null.																									
GE	Logical >=																									
GT	Logical >																									
LE	Logical <=																									
LT	Logical <																									
EQ	True if right and left fields are equal.																									
NE	True if right and left fields are not equal.																									
PROPERTY	rejects	<p>true or false</p> <p>Optional filter value, default value false. If true, then the operator will expect a second OUTPUT to be specified into which rejected (filtered) records will be deposited.</p>																								
OUTPUT		The first output dataset specified contains all of the records for which the filter expression evaluates to true.																								
PROPERTY		The second output is required if the rejects property (above) is specified. This output dataset is the set of records for which the filter expression evaluates to false.																								

## GENERATOR

Element Name	Name	Value
INPUT		This property is optional. If specified, this is the input dataset. If not specified, numrecords must be specified.
PROPERTY	parallel	<p>true or false</p> <p>true—GENERATOR should be run in parallel.</p> <p>false—GENERATOR should be run in a single partition (default).</p>

Element Name	Name	Value
PROPERTY	numpartitions	Optional property indicating the number of data partitions to create.
PROPERTY	numrecords	This optional field is used when no INPUT is specified to determine how many records are generated. If an input is specified this value is ignored. The default value is 1.
PROPERTY	schema	Describes the new column data structure to be generated. See the GENERATOR specification and the example for more detail.
OUTPUT	name.v	The output dataset name.

## REMOVEDUP

Element Name	Name	Value
INPUT		The input dataset name.
PROPERTY	key	Column names of columns to use to detect duplicate records. More than one column property can be specified.
PROPERTY	keep	first or last Either keeps the first or last data if key fields are the same. It defaults to first. <b>Note:</b> Specify first as the value for the keep property when sort order makes it possible. This increases REMOVEDUP performance.
OUTPUT		The output dataset name.

## Filter Expressions

Filter expressions consist of fields, constants, and comparison operators, much like the WHERE clause of a SQL SELECT statement.

Simple filter expressions are in one of the following forms:

*operand comparison-operator operand*

or

*field null-comparison-operator*

where:

*operand* can be the name of a field or a constant.

*comparison-operator* is one of the following:

- GE (greater than or equal to)
- GT (greater than)
- LE (less than or equal to)
- LT (less than)

- EQ (equal to)
- NE (not equal to)

null-comparison-operator is one of the following:

- IS\_NULL
- IS\_NOT\_NULL

For example:

```
CUST_KEY EQ -1
CUST_DT_OF_BIRTH IS_NOT_NULL
```

Simple filter expressions can be chained together using AND and OR to form more complex expressions. The complex expression is always evaluated from left to right. That is, AND and OR are at the same order of evaluation.

---

---

**Note:** This order of evaluation is different from the order used by SQL.

---

---

Parentheses are intended to alter the order of evaluation. Subexpressions within parentheses are evaluated first.

For example:

```
TRAN_CODE EQ 30 OR TRAN_CODE EQ 31
START_DATE LE DAY_DT AND END_DATE GE DAY_DT
CUST_DT_OF_BIRTH EQ '20050101' OR CUST_DT_OF_BIRTH IS_NULL
((TRAN_CODE EQ 30 AND TRAN_TYPE EQ 'A') OR (TRAN_CODE EQ 31 AND TRAN_TYPE EQ 'B'))
```

- Note that without the parentheses, the filter expression above would not evaluate as intended because the order of evaluation is left to right and would be equivalent to

```
((TRAN_CODE EQ 30 AND TRAN_TYPE EQ 'A') OR TRAN_CODE EQ 31) AND TRAN_TYPE EQ 'B')
```

- In this case, a record with TRAN\_CODE = 30 and TRAN\_TYPE = 'A' would evaluate to false.

## Structures and Data Manipulation Operators Examples

This section provides examples of structures and data manipulation operators.

### CONVERT

```
<OPERATOR type="convert">
  <INPUT name="input.v"/>
  <PROPERTY name="convertspec">
    <![CDATA[
      <CONVERTSPECS>
        <CONVERT destfield="SUPP_IDNT" sourcefield="SUPPLIER"
          newtype="string">
          <CONVERTFUNCTION name="string_from_int64"/>
        </CONVERT>
        <CONVERT destfield="LOC_IDNT" sourcefield="LOCATION"
          newtype="string">
          <CONVERTFUNCTION name="string_from_int32"/>
        </CONVERT>
      </CONVERTSPECS>
    ]]>
  </PROPERTY>
</OPERATOR>
```

```

        <TYPEPROPERTY name="nullable" value="true"/>
    </CONVERT>
</CONVERTSPECS>
]]>
</PROPERTY>
<OPERATOR type="debug"/>
</OPERATOR>

<OPERATOR type="convert">
  <INPUT name="inv_sbc_lw_dm.v"/>
  <PROPERTY name="convertspec">
    <![CDATA[
      <CONVERTSPECS>
        <CONVERT destfield="SBCLASS_KEY"
          sourcefield="SBCLASS_KEY">
          <CONVERTFUNCTION name="make_not_nullable">
            <FUNCTIONARG name="nullvalue" value="-1"/>
          </CONVERTFUNCTION>
        </CONVERT>
      </CONVERTSPECS>
    ]]>
  </PROPERTY>
  <OUTPUT name="converted.v"/>
</OPERATOR>

```

## FIELDMOD

```

<OPERATOR type="fieldmod">
  <INPUT name="input.v"/>
  <PROPERTY name="keep" value="Emp_Name Emp_Age"/>
  <PROPERTY name="keep" value="Emp_ID"/>
  <OUTPUT name="output.v"/>
</OPERATOR>

<OPERATOR type="fieldmod">
  <INPUT name="input.v" />
  <PROPERTY name="rename" value="Emp_Name=EmpName"/>
  <OUTPUT name="output.v"/>
</OPERATOR>

<OPERATOR type="fieldmod">
  <INPUT name="input.v"/>
  <PROPERTY name="drop" value="Emp_Title"/>
  <OUTPUT name="output.v" />
</OPERATOR>

<OPERATOR type="fieldmod">
  <INPUT name="input.v" />
  <PROPERTY name="duplicate" value="Emp_Name=Name"/>
  <OUTPUT name="output.v"/>
</OPERATOR>

```

## FILTER

```
<OPERATOR type="filter">
  <INPUT name="input.v"/>
  <PROPERTY name="filter" value="SSN EQ '123456789' AND
    SALARY GT '124.22' AND DOB LT '19981010'"/>
  <PROPERTY name="rejects" value="true"/>
  <OUTPUT name="valid.v"/>
  <OUTPUT name="reject.v"/>
</OPERATOR>
```

## GENERATOR

```
<OPERATOR type="generator">
  <PROPERTY name="numrecords" value="5"/>
  <PROPERTY name="schema">
    <![CDATA[
      <GENERATE>
        <FIELD name="DM_REC'D_LOAD_DT" type="date"
          nullable="false">
          <CONST value="19800101"/>
        </FIELD>
        <FIELD name="LOC_KEY" type="int8" nullable="true">
          <SEQUENCE init="1" incr="1"/>
        </FIELD>
      </GENERATE>
    ]]>
  </PROPERTY>
</OPERATOR>

<OPERATOR type="generator">
  <INPUT name="input.v"/>
  <PROPERTY name="schema">
    <![CDATA[
      <GENERATE>
        <FIELD name="F_FULL_PO_COUNT" type="int64"
          nullable="true">
          <CONST value="1"/>
        </FIELD>
        <FIELD name="F_PART_PO_COUNT" type="int64"
          nullable="true">
          <CONST value="0"/>
        </FIELD>
      </GENERATE>
    ]]>
  </PROPERTY>
  <OPERATOR type="debug"/>
</OPERATOR>
```

## REMOVEDUP

```
<OPERATOR type="removedup">
  <INPUT name="input.v"/>
  <PROPERTY name="key" value="SKU_KEY"/>
  <PROPERTY name="key" value="LOC_KEY"/>
  <PROPERTY name="keep" value="LAST"/>
  <OUTPUT name="output.v"/>
</OPERATOR>
```

---

---

## RETL Schema Files

RETL can process datasets directly from a database or from file-based systems. Depending on the data source, one of two operators can perform this function:

- For reading data directly from database tables, use the ORAREAD operator. See [Chapter 4, "Database Operators"](#) for more information about how to use ORAREAD.
- For situations in which you expect RETL to process data in a file format, use the IMPORT operator. This operator imports a disk file into RETL, translating the data file into a RETL dataset.

### Schema File Requirements

RETL stores information about each dataset that describes the data structures (the metadata). The means by which you supply metadata is dependent of whether you interface data using the DBREAD or IMPORT operator. When datasets are created using DBREAD, the metadata is read along with the data directly from the database table. In this case, RETL requires no further metadata. If the IMPORT operator is used for dataset input, RETL requires the use of a schema.

The schema ensures that datasets are consistent from the source data to the target data. RETL reads a schema file that is specific to the dataset that is being imported. RETL uses the schema file to validate the input dataset structure. For example, if the dataset ultimately populates a target table in a database, the schema file dictates the type and order of the data, as well as the data characteristics such as the character type, length, and nullability.

Two types of schema files are defined for use with an incoming text data file:

- Delimited file type
- Fixed length file type

## Schema File XML Specification

Element Name	Name	Value/Description
RECORD		RECORD is the root tag of the schema file.
	type	delimited or fixed Required field. Specified whether the schema is fixed records or delimited records.
	final_delimiter	Any character Required field. End of record delimiter that is used within the input file to distinguish the end of a record. This is extremely important to help distinguish between records in the event that the input data file is corrupt. Generally this is a newline character for readability.
	len	1..n Required only if the type attribute is fixed. This is the total length of the fixed length record. Note that this must be equivalent to the sum of all aggregate field lengths.
	sortkeys	Space separated list of keys Optional property to specify what fields the records are sorted by. This should be used only when it is known that the records will always be sorted by specified keys. This property must be specified in conjunction with "sortorder".
	sortorder	asc or desc Optional property to specify the sort order of the record. Specify asc for ascending sorted order and "desc" for descending sorted order. See the sortkeys attribute.

Element Name	Name	Value/Description
FIELD		FIELD is a child element of RECORD.
	name	The name of the field.
	delimiter	field delimiter This optional attribute is for delimited fields only. It specifies the end-of-field delimiter, which can be any character, but the default value is the pipe character ( ). The delimiter can be specified in hexadecimal notation by prefixing 0x to the ASCII value of the character. This allows specification of a delimiter that is highly unlikely to appear in the data, such as 0x02.
	datatype	Any one of the following: int8, int16, int32, int64, uint8, uint16, uint32, uint64, dfloat, sfloat, string, date, time, or timestamp. See <a href="#">Appendix A, "Appendix: RETL Data Types"</a> (especially the note in RETL data type to Oracle data type, ORAWRITE).
	nullable	true or false Optional field describing whether null values are allowed in the field. The default value is false. If true then a nullvalue MUST be specified.
	nullvalue	Required if nullable is "true". This is the text string that an IMPORT operator uses when reading flat files to determine whether a field is null. This should be a unique value that is not otherwise be found within the dataset. For fixed-length fields, this value must be equal in length to the field length, so that the record length remains fixed. <b>Note:</b> The null string ("") can specified for a fixed-length field, rather than specifying a number of blanks equivalent to the length of the field. Null values must be a valid value for the type specified. For example, 00000000 as a nullvalue for a date is invalid.
	maxlength	Specifies the maximum allowable length of the string field in characters. Required only for delimited string fields.
	len	Specifies the length of the field in characters. Required for all fixed length fields.
	strip	leading, trailing, both, or none Optional attribute used for delimited string fields, to determine whether and how white space (tabs and spaces) should be stripped from input field. The default value is None.

## Delimited Record Schema

When you create a schema for use with a delimited file type, follow these guidelines:

- The delimiter can be any symbol.

---

---

**Note:** Make sure that the delimiter is never part of your data.

---

---

- To define the field's nullability, set `nullable` to `true` or `false`.
- If the field is nullable, set the default `nullvalue`. This is often set to `""` (the null string).
- For the string data type, specify the `maxlength`.

## Example Delimited Schema File

```
<RECORD type="delimited" final_delimiter="0x0A">
  <FIELD name="colname1" delimiter="|" datatype="int8" nullable="false"/>
  <FIELD name="colname2" delimiter="|" datatype="int16" nullable="true"
nullvalue="" />
  <FIELD name="colname3" delimiter="|" datatype="int32"
nullable="true" nullvalue="" />
  <FIELD name="colname4" delimiter="|" datatype="int64"
nullable="true" nullvalue="" />
  <FIELD name="colname5" delimiter="|" datatype="dfloat"
nullable="true" nullvalue="" />
  <FIELD name="colname6" delimiter="|" datatype="string" maxlength="5"
nullable="false" />
  <FIELD name="colname7" delimiter="|" datatype="date"
nullable="false" />
  <FIELD name="colname7" delimiter="|" datatype="timestamp"
nullable="false" />
  <FIELD name="colname8" delimiter="|" datatype="uint8"
nullable="false" />
  <FIELD name="colname9" delimiter="|" datatype="uint16"
nullable="true" nullvalue="" />
  <FIELD name="colname10" delimiter="|" datatype="uint32"
nullable="true" nullvalue="" />
  <FIELD name="colname11" delimiter="|" datatype="uint64"
nullable="true" nullvalue="" />
  <FIELD name="colname12" delimiter="|" datatype="sfloat"
nullable="true" nullvalue="" />
</RECORD>
```

## Fixed-Length Record Schema

When you create a schema for use with fixed-length records, follow these guidelines:

- Specify the length (len) of the record.
- Specify a len value for every field. The total of all fields must be equivalent to the record length.
- To define the field's nullability, set nullable to true or false.
- If the field is nullable, set the default nullvalue. Remember that since this is a fixed length field the null value must be the correct length for the field.
- If possible, specify the final\_delimiter to ensure that input files are at least record-delimited, so that RETL can recover in the event that data is corrupt or invalid.

---



---

**Note:** Make sure that the final\_delimiter is never part of your data.

---



---

### Example Fixed-Length Schema File

```
<RECORD type="fixed"      len="99"          final_delimiter="0x0A">
  <FIELD name="colname1"  len="2"          datatype="int8"
nullable="false"         />
  <FIELD name="colname2"  len="5"          datatype="int16"
nullable="false" />
  <FIELD name="colname3"  len="10"         datatype="int32"
nullable="false" />
  <FIELD name="colname4"  len="18"         datatype="int64"
nullable="false" />
  <FIELD name="colname5"  len="6"          datatype="dfloat"
nullable="true" nullvalue="NULVAL" />
  <FIELD name="colname6"  len="5"          datatype="string" nullable="false"
/>
  <FIELD name="colname7"  len="8"          datatype="date"
nullable="false"         />
  <FIELD name="colname7"  len="14"         datatype="timestamp"
nullable="false"         />
  <FIELD name="colname8"  len="2"          datatype="uint8"
nullable="false"         />
  <FIELD name="colname9"  len="4"          datatype="uint16"
nullable="true" nullvalue="XXXX" />
  <FIELD name="colname10" len="5"          datatype="uint32"
nullable="true" nullvalue=" " />
  <FIELD name="colname11" len="18"         datatype="uint64"
nullable="false" />
  <FIELD name="colname12" len="18"         datatype="sfloat"
nullable="false" />
</RECORD>
```

## nullvalue Considerations

Because nullvalues can be changed between import and export schemas, it is possible to lose data if you are not careful. However, this property can also be used to assign default values to null fields on export. No matter how you use them, take care when selecting values for a field's nullvalue. Unless you are assigning a default value to a null field on export, do not use a value that can appear as a valid value within a given dataset.

To illustrate consider the following schema (named 1.schema):

```
<RECORD type="delimited" final_delimiter="0x0A">
  <FIELD name="colname1" delimiter="|" datatype="int8" nullable="false" />
  <FIELD name="colname2" delimiter="|" datatype="dfloat" nullable="true"
nullvalue="" />
</RECORD>
```

Also consider this schema (named 2.schema), where the null value has been modified for the second field:

```
<RECORD type="delimited" final_delimiter="0x0A">
  <FIELD name="colname1" delimiter="|" datatype="int8" nullable="false" />
  <FIELD name="colname2" delimiter="|" datatype="dfloat" nullable="true"
nullvalue="0.000000" />
</RECORD>
```

When running through the following flow:

```
<FLOW>
  <OPERATOR type="import">
    <PROPERTY name="inputfile" value="1.dat" />
    <PROPERTY name="schemafile" value="1.schema" />
  <OUTPUT name="import.v" />
  </OPERATOR>
  <OPERATOR type="export">
    <INPUT name="import.v" />
    <PROPERTY name="outputfile" value="2.dat" />
    <PROPERTY name="schemafile" value="2.schema" />
  </OPERATOR>
</FLOW>
```

Where 1.dat contains the following data:

```
2|
3|4.000000
4|5.000000
5|0.000000
```

The following is the resulting output to 2.dat:

```
2|0.000000
3|4.000000
4|5.000000
5|0.000000
```

By using an export schema that has valid values to represent the nullvalue for the second field, this flow has assigned a default value to all of the second field's null values.

Then if you change the input and schema files, as in the following flow:

```
<FLOW>
  <OPERATOR type="import">
    <PROPERTY name="inputfile" value="2.dat" />
    <PROPERTY name="schemafile" value="2.schema" />
    <OUTPUT name="import.v" />
  </OPERATOR>
  <OPERATOR type="export">
    <INPUT name="import.v" />
    <PROPERTY name="outputfile" value="3.dat" />
    <PROPERTY name="schemafile" value="1.schema" />
  </OPERATOR>
</FLOW>
```

Output file 3.dat will look like this:

```
2|
3|4.000000
4|5.000000
5|
```

---

---

**Note:** Null values are always considered the smallest value in a field. That is, for any value n, null < n. In sorted order, null values precede all other values in ascending order and follow all other values in descending order.

---

---

## Configure RETL to Print Schema Output in Schema File Format

One error-prone area in the development of RETL flows is when developers break up and put together flows. The schema files are the interface between flows that tell RETL how to read and write files. Writing schema files manually is tedious and error-prone. It is highly recommended that developers use the `-sSCHEMAFILE` command line option to speed up development. For more information about this feature, see the "[rfx Command Line Options](#)" section in Chapter 2, "[Installation and System Configuration](#)."



---

---

## Common Operator Properties

This chapter describes properties common to all operators.

### Common Operator XML Specification

Element Name	Name	Value
PROPERTY	progind	Progress indicator string displayed as records are processed.
PROPERTY	progfreq	How often to display the progind string. The progind string is displayed every progfreq records.



Traditional application integration is done in a point-to-point manner. Developers are given an integration tool and told to integrate to a defined file/database specification or to an API. The end result may be functionally valid according to the single specification, but implementation may be cumbersome, non-reusable, and subject to heavy maintenance costs. The design itself may not have accounted for all the business process needs, invalidating the entire implementation.

Consequently, best practice guidance is required to help ensure the interface is developed and deployed to the maximum advantage in terms of costs and benefits. This chapter includes simple and proven practices that can be used when developing and designing integration flows using the RETL tool.

RETL moves large amounts of information in bulk. However, managing the variety of disparate systems and pulling data together to form a complete picture can be very complex tasks.

As a result of business process, integrated system, database and interface complexity, the RETL tool must be used correctly to ensure good performance, maintainability, and quality.

This chapter describes the following activities—and best practices that should be considered before, during, and after each activity:

1. **Project Initiation/Design/functional specification**—The best practices to follow during requirements elicitation, interface analysis, and design.
2. **Code/Implementation/Test**—The best practices to follow when setting up the environment, during flow construction and all phases of testing. This will be the bulk of the chapter.
3. **Review/product handoff**—Activities to follow when reviewing the project and its deliverables. This also provides guidance for handing off interface deliverables to operations and support personnel.

This content in this chapter will expand to include tips from experts in domains related to RETL.

## Prerequisites

When designing and developing RETL integration flows, the following core skill sets are required.

### Interface Design Skills

A functional/technical architect should design the interfaces. The designer should have the following:

- An understanding of source and target data sources; in-depth knowledge of the data models for each application and of how a transformation may be performed to integrate the two applications
- Ability to understand how each application and business process works; ability to quickly interpret data needs for an application, based on the business processes that need support.
- General understanding of schedule, data dependencies, batch jobs, and volume of the application in question.

### ETL Technical Skills

ETL technical skills. The ETL coder should have the following:

- Strong Unix experience, including in-depth Korn shell scripting.
- Previous RETL experience, including familiarity with the *Oracle Retail Extract, Transform, and Load Programmer's Guide* and the ability to write RETL flows.
- Strong database experience, including familiarity with SQL statements, database operations, and utilities (for example, SQL\*Loader).

## Project Initiation/Design/Functional Specification Best Practices

Before development begins, a solid, generic, and re-usable interface design must be in place. The following practices will help ensure such a design:

### Ask Discovery Questions

An important part of the functional design is to ask relevant questions that can answer the following:

#### General Integration Questions

- *What type of data integration is needed?* Is there potential for real-time or near-real-time integration? If there is potentially a need for more synchronous/real-time integration, consider using a different technology, such as the Oracle Retail Integration Bus (RIB) or direct access.
- *What business process is to be integrated?* For any project manager or flow developer to understand how to construct the data flow, they must understand at a high-level the business process it is that they will support. (Why do the users in system X need the data created in system Y?) Without this understanding, RETL users may not get a complete picture of what it is they need to integrate, nor would they have the opportunity to ask questions that may reveal additional interface capabilities, variations, or options for business users.

### Application Domain Questions

- *What are the targeted versions for each application involved?* This is important to establish as a basis for integration development against each product and the business features and functions available in a particular version.
- *What is the source and target for each module?* For example, database to file, file to file, and file to database.
- *On what DBMS and version does the source and target application reside, if any?* If the source/target database isn't supported by RETL, consider using a different technology or provide feedback to the RETL team to include that database support in a future release.
- *What types of transformations might be expected for each module?* This affects the complexity of each module.
- *Are there any constraints on sources and targets?* For example, does the target database table need to maintain indexes? Are there referential integrity issues? There are implications when developing a flow against a transaction database or table that maintains indexes or referential integrity.
- *What future applications might re-use these interfaces?* Designs should be generic and promote easy modification, should any future applications need to re-use the same interfaces.

### Data-related and Performance Questions

- *How much data is passing through each module?* High volume modules must be further scrutinized when coding and performance testing.
- *What is the size of the source and target database tables/files?* This is to consider performance early in the project.
- *What is the frequency in which each module will be run?* For example, nightly.
- *What is the time frame in which each module is expected to run in?* For the entire set of modules? This span must be realistic, as it will provide a baseline for performance.
- *Is data transferred over the network?* There will be a negative effect on performance if massive amounts of data must be transferred over the network.
- *Are there potentially any high volume transformations that could be easily done inside the database?* If there are, these transformations may be done inside the database to start with so as to eliminate rewrites later on. The idea is to maximize and balance each technology.

## Map Out The Movement of Data Visually

Use a tool to visually map out the movement of data from each source to each target. Use a tool such as Visio or even simply Word to diagram the source, transformation processing/staging, and target layers. (Use a database template or the 'basic flowchart' in Visio to map these layers out). This should serve as a basis for design discussions among application groups.

## Define Concrete Functional Requirements for Each Module

As part of the design process, a logical description of how each module will extract, transform, and load its data should be completed. Again, all stakeholders should sign off on this.

## Define Concrete Functional Designs for Each Module

As part of the design process, the following should be clearly defined:

- Source and Target locations (such as table name and file name)
- Source and Target sizes/volumes
- Designs and metadata definitions for each input and output
- Mapping of the transformation process on how to get from input format to output format
- Name of the script/module

## Design a Test Plan Early in the Process

Designs are not complete until the unit, system, integration and performance test plan designs are complete. This can be completed once the data and test environment needs are identified. Activities to complete for the test plan include the following:

- Identification of the different combinations of data and the expected results.
- Identification of degenerate cases and how errors should be handled.
- Identification of how many and what type of test cases that need to be built.
- Identification of configuration management and environment migration procedures to ensure that quality control individuals know what code set it is that they are testing.
- Identification of test environment needs to ensure that unit, system, integration and performance test can be done without negatively impacting each other.
- Prioritization of relative importance of execution for each test case based on most likely scenarios or highest risk code, for example.
- Determine a go/no-go threshold for which tests must pass in order to release.

## Design for Future Usage and Minimize Impact of Potential Changes

The design should take into account future applications that may need to live off the same interfaces. Relevant questions to ask are the following: What is the superset of data that an application needs or may need in the future? How can this flow be designed so that any future changes will not break the interface specification?

## Agree on Acceptance Criteria

The overall project should set requirements based on a number of acceptance criteria set by the customer. A partial list of these may involve the following:

- *Functional criteria*—the inputs, outputs, and expected functionality of each module
- *Test criteria*—the types and how intensive must the testing be, and the threshold in acceptance of defects.
- *Performance criteria*—the performance requirements for each module, and the overall batch window in which the entire suite of modules must run in.
- *Documentation*—Identify clearly the end deliverables needed to hand off the interface deliverable to operations and support staff. A well-constructed interface cannot be used and maintained successfully without documentation or communication deliverables.

## Document Design Assumptions, Issues, and Risks

Any design assumptions, issues and risks that have been purposely un-addressed should be documented in an Assumptions/Issues/Risks section of the design documentation deliverables.

## Code/Implementation/Test Best Practices

This section presents best practices within functional areas:

- Korn shell module best practices—the .ksh module containing the RETL flow
- RETL flow best practices—the XML flow that RETL runs
- Database best practices
- Testing best practices

Following this core set of best practices when developing and testing flow modules will provide the basis for high quality, maintainable, performant, and re-usable flows.

### Korn Shell Best Practices

The following are best practices for Korn shell.

#### Execute Commands Using \$(command) and Not command

This reduces the likelihood of typographical errors when entering backspaces instead of single quotes and promotes better readability.

#### Ensure set -f is set in a Configuration File

The setting, set -f, disables filename generation. This should be set in the application's config file (for example, rdw\_config.env). If filename generation isn't turned off, a flow that contains an asterisk (\*) in it (for example, a flow that contains a query, "select \* from <table>," may end up causing the module to incorrectly expand \* to include filenames and incorrectly pass these to RETL. This is because many modules may pass RETL a flow through standard input. See best practice, "Write flow to intermediate file and then call RETL on that file."

#### Write Flow to an Intermediate File and then Call RETL on that File

It is recommended to 'cat' or write to a flow file first before calling RETL, rather than calling RETL and passing a flow through standard input. This aids in debugging, helps support when it is necessary to recreate an issue, and also prevents various nuances that may exist when sending the flow to RETL via standard input.

However, care should also be taken to protect source directories since these generated .xml flow files may contain sensitive information such as database log in information. See the "[Secure/Protect Files and Directories that may Contain Sensitive Information](#)" best practice for more on the topic of security.

**Good**

```
cat > ${PROGRAM_NAME}.xml << EOF
...
EOF
${RFX_EXE} ${RFX_OPTIONS} -f ${PROGRAM_NAME}.xml
```

**Bad**

```
${RFX_HOME}/bin/${RFX_EXE} -f - << EOF
...
EOF
```

**Secure/Protect Files and Directories that may Contain Sensitive Information**

Any directories that may contain sensitive information such as database log ins and passwords should be protected by the proper UNIX permissions. Namely, any temporary directories that RETL uses, any source directories in which flows may be written to, and any configuration files and directories should have strict permissions on them. Only the userid that runs the flows should have access to these files and directories.

**Make Function Calls Only a Few Layers Deep**

Use variables for portions of code that can be re-used (for example, \${DBWRITE}), or create entire shell script functions for methods that can be re-used (for example, simple\_extract "\${QUERY}" "\${OUTPUT\_FILE}").

This can be taken to the extreme however, when functions call functions, which call functions. In general, function calls should go only a few layers deep at most, and only more when it makes sense from an abstraction and overall maintainability point of view.

**Separate Environment Data from the Flow**

Variables that can change (for example, database log in information) should go in a configuration file (for example, rdw\_config.env). The idea here is to separate the configurable environment data from the generic flow, which lends to more maintainable and configurable code base.

**Enclose Function Parameters in Double Quotes**

When calling a function in ksh, care should be taken to place all arguments in double quotes. For example,

```
do_delete "${TABLE1}" "${TABLE2}" "${JOIN_KEYS}"
```

If function parameters aren't placed in quotes, there is the potential for misalignment of positions for variables and the function will likely not operate properly, if at all.

**Set Environment Variable Literals in Double Quotes**

When setting an environment variable, always use double quotes. For example,

```
export TARGET_TABLE="INV_ITEM_LD_DM"
```

If environment variable values aren't placed in quotes, there is the potential for losing the portion of the value after a space. or having the shell interpret it as a command.

## Use Environment Variables as \${VARIABLE} Rather than \$VARIABLE

Environment variables should be used as \${VARIABLE} so as to promote readability, especially when they exist in the context of other non-variables and text.

## Follow Module Naming Conventions

There are a number of naming conventions to follow, and in general, this should be the standard for consistency among applications:

1. RETL modules/scripts (.ksh) should generally conform to the following convention:

<application name><script type>\_<program name>.ksh,

where

<application name> is the abbreviated application name,

<script type> is e for an extract program, t for a transform program, or l for a load program. A program that contains all of the above will not have a script type associated with it.

<program name> is the name of the program.

For example:

rmse\_daily\_sales.ksh (<rms><e>\_<daily\_sales>.ksh)

2. The program name in general doesn't have strict naming conventions, but it should be representative of what the program actually does. Each word should be separated by an underscore (\_).
3. Custom modules/flows should be prefixed by the short client name.
4. Application log files should follow the format:

<application name><date>.log

5. Error files should follow the format:

<application name>.<module name>.<unique id>.<date>

where

<application name> is the abbreviated application,

<module name> is name of the module,

<unique id> is optional. If multiple copies of the module can be run concurrently, this should be a reasonable unique identifier. For example, input filename or thread #.

For example:

rdw.slsildmdm.sls\_100\_1.txt.20020401

(<rdw>.<slsildmdm>.<datafile #>.<date>)

### Log Relevant Events in Module Processing

As a general rule, it is better to log too much detail than too little. The following should be completed:

1. Log the start and finish times for any significant processing events that occur in a module.
2. Log the number of records in the input and output files.
3. Redirect standard output and standard error to the log/error file. For example:

```
exec 1>>$ERR_FILE 2>&1
```

### Place Relevant Log Files in Well-known Directories

Each application log file should be placed in a well-known directory for consistency and ease of accessibility. For example, RETL performance log files (as specified in `rfx.conf`) are placed in temporary directories by default. `rfx.conf` should be changed so that RETL performance log files places these log files in well-known directories rather than in temporary directories.

### Use .ksh Templates

Insert template on a standard flow (for example, library sourcing first, then variable definitions, function initionsdefs, and so on.)

### Document Each Flow's Behavior

Each flow should be documented as to its required inputs, required outputs, and behavior in between.

## RETL Flow Best Practices

The following are best practices for RETL flow.

- **Parameterize the Call to and any Options Passed to the RETL Binary**

Use the `#{RETL_EXE}` environment variable to replace the call to the `rfx` executable and use `#{RETL_OPTIONS}` to replace the command line option `'-f -'`. `"rfx -f -"` should not be used to call RETL - it should be `"#{RETL_EXE} #{RETL_OPTIONS}"`. `#{RETL_EXE}` and `#{RETL_OPTIONS}` should be defined in the configuration file (for example, `rdw_config.env`). The default call in the configuration file should be `rfx -f -`

- **Perform all Field Modifications as Soon as Possible in the Flow**

Field modifications such as renaming and type conversion should be performed as soon as possible in the flow to aid in tracking a field through the flow.

- **Use Care when Choosing and Managing your Temp Space**

RETL temp space should be spread out on multiple disks and, as part of a normal maintenance schedule, should be cleaned periodically as RETL may leave miscellaneous files in temp space for tracking and debugging purposes.

- **Turn on Debug Code When in Development.**

- **Turn Off Debug Code When in Production**

RETL will print out extra messages in debug mode. The following environment variables can be set to produce extra verbosity:

```
export RFX_DEBUG=1
```

```
export RFX_SHOW_SQL=1
```

Additionally, there are certain warnings that may be printed to the screen in this mode that will not be printed to the screen when verbose debugging is turned off.

Do not run with these options on in production as they produce additional files and may leave around temporary/intermediate files.

- **Make Use of RETL Visual Graphs**

In RETL 10.x versions and 11.2+, RETL offers a graphing option (--graphviz-files in 10.x or -g in 11.2+). This produces a visual graph of the flow and shows how each operator is connected inside the flow. For more information, see ["Producing Graphical Output of Flows with RETL"](#) in Chapter 4, ["RETL Program Flow."](#)

- **Delineate Input records from a File by a Newline ('\n')**

Although RETL doesn't mandate having a newline record delimiter, it is recommended to use this for debugging and supportability purposes. However, it is also important to make sure that input data cannot contain the newline character.

- **Run RETL With the -s SCHEMAFILE Option When in Development**

The -s SCHEMAFILE command-line argument will tell RETL to print input and output schemas of each operator in a schemafile format that can be used in an import or export operator. This can be very useful when debugging and 'breaking up' flows into smaller portions. For example, the call to RETL would be as follows:

```
rfx -f theFlow.xml -s SCHEMAFILE
```

- **Use the Latest Version of RETL Available**

Future versions of RETL are always improving in error handling, debugging, logging, and performance. It is recommended to use the latest version of RETL that has been certified with the product being used.

- **Develop the RETL Flow First, Outside of the Shell Module**

It is easier to develop the RETL flow outside the Korn shell module (.ksh) first, and then when a stable cut of the flow has been produced, the flow can be parameterized and retrofitted back into the Korn shell module (.ksh)

- **Specify Inputs, Properties, and Outputs, in that Order**

In order to conceptually map out the flow of data from one operator to another, place the INPUTs first, the OUTPUTs last, and the properties in between.

For example:

```
<OPERATOR type="changeproperty">
  name="before_dataset.v" /
<INPUT name="after_dataset.v" /> name="key" value="CMPY_IDNT" / name="value"
value="CMPY_DESC" / name=":"changes.v" /</OPERATOR>
```

- **Group Business Logic Together with Nested Operators**

- **Keep Connected Operators in Close Physical Proximity to Each Other in a Flow**

- **Nest Operators No More Than a Few Layers Deep**

When multiple operators are used to perform a logical step in the flow, nest the operators together.

In the example below, the operators split the `after_dataset.v` dataset into separate datasets for inserts, deletes, and edits. The switch operator is embedded within the `changecapture`.

```
<!-- Split up after_dataset.v into inserts, deletes, and edits. -->
<OPERATOR type="changecapture" >
<INPUT name = "before_dataset.v"/>
<INPUT name = "after_dataset.v"/>
<PROPERTY name="codefield" value="change_code" />
<PROPERTY name = "key" value = "CMPY_IDNT"/>
<PROPERTY name = "value" value = "CMPY_DESC"/>
<OPERATOR type="switch">
<PROPERTY name="switchfield" value="change_code"/>
<PROPERTY name="casevalues" value="1=0, 2=1, 3=2"/>
<OUTPUT name="inserts.v"/>
<OUTPUT name="deletes.v"/>
<OUTPUT name="edits.v"/>
</OPERATOR>
</OPERATOR>
```

As the operator nesting gets deeper, the flow becomes unreadable, so don't nest operators more than a few layers deep.

- **Document Fixed Schema File Positions**

In order to debug flows and to aid in documenting interface fixed-length files, fixed-length schema files should include positions as a comment before each field definition. However, this also adds additional overhead to schema file maintenance, and could prompt for typographical errors. This should be taken into account when making a decision to document a module's schema files.

Example:

```
<RECORD type="fixed" len="9" final_delimiter="0x0A">
<!-- start pos 1 --> <FIELD name="FIELD1" len="4" datatype="int16" .../>
<!-- start pos 5 --> <FIELD name="FIELD2" len="4" datatype="int16" .../>
<!-- end pos 9 -->
</RECORD>
```

- **Use Valid XML Syntax**

10.x versions of RETL allowed users to enter invalid XML syntax. Later versions are more strict about valid XML syntax, and will throw errors if syntax is invalid. In particular, when calling RETL's filter operator, utilize 'GT' instead of >, LT instead of <, LE instead of <=, GE instead of >=, NOT ... EQ ... instead of <> and EQ instead of =. Also see the best practice Wrap dbread query property in a CDATA element.

- **Wrap dbread query Properties in a CDATA Element**

This is important to make sure flows are using valid XML syntax. In general, ALL query properties of the dbread operators should be wrapped in a CDATA element. This tells the XML parser to treat the enclosed text as a single chunk of text rather than to parse it as XML. If this rule is not followed, RETL may not run (see best practice, "Use valid XML syntax," for more info).

**Bad**

```
#{DBREAD}
<PROPERTY name="query" value="SELECT * FROM TABLE WHERE COL1 > COL2"/>
...
```

**Good**

```

${DBREAD}
<PROPERTY name="query">
<![CDATA[
SELECT * FROM TABLE WHERE COL1 > COL2
]]>
</PROPERTY>

```

- **Write Flows to be Insulated from Changes to a Database Table**

In general, flows should be written so that any changes to database tables don't affect the functionality of the flow. For example, the FIELDMOD operator can be used with the "keep" property to keep only fields that are needed in the flow.

- **Avoid Implicit Schema Changes**

Certain operators that take more than one input (such as funnel) can implicitly change schemas for a field's nullability or max length, for example. Any desired schema changes should be explicitly made with the CONVERT operator otherwise undefined behavior may result.

- **Test Often**

Testing often minimizes integration issues, increases quality, aids in easier defect diagnosis, and improves morale.

- **Do RETL Performance Tuning**

Consider the following practices guidelines.

- Achieve the optimal balance between the tools, hardware, and database.
- Use the database for what it does best - significant data handling operations such as sorts, inter-table joins, and groupby's should be handled inside the database when possible.
- Use RETL for what it does best - non-heterogeneous (for example, file, different db vendor) joins, data integration and movement from one platform to another and multiple source locations.
- Performance test and profile early and continuously in the development cycle. Don't delay this activity until the end of the project.
- Do not performance tune a flow before the flow functions according to design.
- Design a flow to make full use of parallelism. RETL internal parallelism coupled with standard Oracle Retail multi-threading, or running multiple RETL processes concurrently, can maximize the use of parallelism.
- Develop for clarity, maintenance, and re-use, but keep in mind performance at all times.

- **Name Flows Appropriately in the <FLOW> Element**

Flows should include a 'name' attribute in the FLOW element of the XML. This eases debugging in RETL's logging mechanism. In general, the flow should be named the same as the module in which the flow is run. For example,

```

<FLOW name="modulename.flw">
...
</FLOW>

```

- **Use SQL MERGE Statements When Doing Updates**

As of this writing, RETL does not have update capabilities built in. The recommended approach when doing updates is to use the Oracle MERGE statement. This statement is akin to a delete followed by insert for each updated record, although MERGE is much faster. The flow's processing would be as follows: insert records to a temp table, and then MERGE into the target table.

- **Document Complex Portions of Code**

As with general programming best practices, any complex portions of code that could be misunderstood should be thoroughly documented. This aids in ease of maintenance and helps future personnel to understand the original intent of the code.

- **Define Field Names Without White Spaces**

Many RETL operator properties allow specification of a list of fields. White space within a field name will cause RETL to misinterpret the field list.

- **Choose a Delimiter That Cannot be Part of the Data**

If a data can even possibly contain a delimiter, a different delimiter should be chosen. Never choose a delimiter that can be part of the data. For the most part, pipe-delimiters (!) and semi-colon delimiters (;) should work. If no reasonable delimiter can be found, a fixed format file should be used.

## **Review/Product Handoff**

A final product review and handoff to support is necessary to complete a successful RETL project.

### **Involve Support Personnel Early in the Project**

This will help the learning curve and make the impact less dramatic when the project must be transitioned to support.

Identify clearly the end deliverables needed to hand off the interface deliverable to operations and support staff. A well-constructed interface cannot be used and maintained successfully without documentation or communication deliverables.

### **Assign a Long-term Owner to the Project / Product / Interface**

Assigning ownership will allow the product to mature and develop properly and it will reduce internal contention about areas of responsibility when multiple applications/domains are involved.

Assigning ownership of the interface following handoff will also ensure that appropriate operations and support personnel are involved and trained throughout the life of the development effort.

---

---

## Sort, Merge, and Partitioning Operators

RETL provides the following sort and merge operators, described in this chapter:

- COLLECT and FUNNEL
- SORTCOLLECT and SORTFUNNEL
- HASH
- SORT
- MERGE

### COLLECT and FUNNEL

Both operators combine records from input datasets as they arrive. This can be used to combine records that have the same schema from multiple sources. Note that these operators are sometimes used implicitly by RETL to rejoin datasets that have been divided during partitioning in parallel processing environments (where the number of RETL partitions is greater than one).

### SORTCOLLECT and SORTFUNNEL

Like COLLECT and FUNNEL these operators combine records from input datasets. These operators maintain sorted order of multiple datasets already sorted by the key fields. Records are collected in sorted order using a merge sort algorithm. If incoming records are unsorted, FUNNEL followed by the SORT operator should be used instead.

---

---

**Note:** The sort order (ascending or descending) of the input datasets to SORTFUNNEL must be the same and must match the order property of SORTFUNNEL. Otherwise, SORTFUNNEL may produce unexpected results.

---

---

### HASH

The HASH operator examines one or more fields of each input record, called hash key fields, to assign records to a processing node. Records with the same values for all hash key fields are assigned to the same processing node. This type of partitioning method is useful when grouping or sorting data to perform a processing operation.

## SPLITTER

The round-robin partitioning operator splits records among outputs in an ordered record-by-record fashion. This is an alternative to the HASH operator and distributes records more evenly than HASH.

## SORT

SORT sorts the records in the RETL dataset based on one or more key fields in the dataset.

## MERGE

MERGE merges input dataset columns record-by-record. Each output record is the union of the fields for each incoming record in the INPUT datasets. The number of records contained in the input datasets must match. Columns are ordered by the order of the input datasets: first input dataset, second input dataset, and so on.

## Sort and Merge Operators XML Specification Tables

This section describes the sort and merge operators XML specification tables.

### COLLECT/FUNNEL

---



---

**Note:** These operators are aliases for each other. FUNNEL is used as a general rule, but either operator will work.

---



---

| Element Name | Name   | Value   |
|--------------|--------|---|
| INPUT        |        | Input dataset name. This can be specified 1 or more times.  |
| PROPERTY     | method | monitor, poll, or cycle<br>Optional property. Defaults to monitor. Specifies the method used to determine which INPUT has a record ready to be funneled. Changing this property from the default can have major performance and CPU usage implications. |
| OUTPUT       |        | The output dataset name.  |

### HASH

| Element Name | Name          | Value  |
|--------------|---------------|--|
| INPUT        |               | Input dataset.   |
| PROPERTY     | key           | Key columns to be hashed.  |
| PROPERTY     | numpartitions | Number of data partitions to create. Defaults to the number of partitions specified in rfx.conf. |
| OUTPUT       |               | The output dataset name.   |

## SPLITTER

| Element Name | Name          | Value   |
|--------------|---------------|---|
| INPUT        |               | The input dataset name.   |
| PROPERTY     | numpartitions | Optional property indicating the number of data partitions to create. Defaults to the number of partitions specified in rfx.conf. |
| OUTPUT       |               | The output dataset name.  |

## SORT

| Element Name | Name          | Value  |
|--------------|---------------|--|
| INPUT        |               | Input dataset name.  |
| PROPERTY     | key           | <p>Key columns to sort by.</p> <p>When specifying multiple sort keys, the order that they are specified is very important.</p> <p>The records are ordered by the first key field. If there is more than one record with the same first key field, those records are ordered by the second key field, and so on.</p> <p>This behavior replicates the sorting that happens in an SQL ORDER BY clause.</p>  |
| PROPERTY     | order         | <p>desc or asc</p> <p>Set to desc for descending sort order and asc for ascending sort order. Default value is asc.</p>  |
| PROPERTY     | removedup     | <p>true or false</p> <p>If set, then duplicate records will be removed. Default is false.</p>  |
| PROPERTY     | tmpdir        | <p>Name of directory to use for temporary files.</p> <p><b>Use of this property is not recommended.</b> Temporary directories should be specified in the rfx.conf. This property allows one to override the temporary directory to use for the sort command.</p> <p><b>Note:</b> If used, the directory should be periodically cleaned as the TMPDIR from the rfx.conf file. See <a href="#">Chapter 2, "Installation and System Configuration"</a>.</p> |
| PROPERTY     | delimiter     | <p>Character to use as a delimiter when writing out the records to a temporary file. Defaults to " " (pipe character).</p> <p>If your data contains the pipe character, set this property to a character that does not occur in your data.</p>   |
| PROPERTY     | numsort       | <p>Integer greater than 0, indicating the number of threads used to sort the data. Defaults to 1.</p> <p>Increasing the number of threads used to sort the input may improve performance.</p>  |
| PROPERTY     | numpartitions | Alias for numsort.   |
| OUTPUT       | name.v        | The output dataset name.   |

## SORTCOLLECT/SORTFUNNEL

**Note:** These operators are aliases for each other. SORTFUNNEL is used as a general rule, but either operator will work.

| Element Name | Name  | Value  |
|--------------|-------|--|
| INPUT        |       | Input dataset name. This can be specified one or more times.   |
| PROPERTY     | key   | Key columns to sort by. Specify multiple keys by using multiple key property instances. For example:<br><pre>&lt;PROPERTY name="key" value="k1" /&gt; &lt;PROPERTY name="key" value="k2" /&gt;</pre> |
| PROPERTY     | order | desc or asc<br>Optional property that describes the sort order of the incoming key fields. Default value is "asc", for ascending order. Specify desc if keys are sorted in descending order.         |
| OUTPUT       |       | The output dataset name.   |

## MERGE

Element Name	Name	Value
INPUT		Input dataset name. This can be specified one or more times.
OUTPUT		The output dataset name.

## Sort, Merge, and Partitioning Operators Tag Usage Examples

The following are usage examples.

### COLLECT

```
<OPERATOR type="collect">
  <INPUT name="test_1.v" />
  <INPUT name="test_2.v" />
  <INPUT name="test_2.v" />
  <OUTPUT name="output.v" />
</OPERATOR>
```

### HASH

```
<OPERATOR type="hash">
  <INPUT name="left.v" />
  <PROPERTY name="key" value="LOC_KEY" />
  <OPERATOR type="sort" >
    <PROPERTY name="key" value="LOC_KEY" />
    <OUTPUT name="result.v" />
  </OPERATOR>
</OPERATOR>
```

## **SORTCOLLECT**

```
<OPERATOR type="sortcollect">  
  <PROPERTY name="key" value="LOC_KEY" />  
  <INPUT name="test_1.v" />  
  <INPUT name="test_2.v" />  
  <OUTPUT name="output.v" />  
</OPERATOR>
```

## **MERGE**

```
<OPERATOR type="merge">  
  <INPUT name="test_1.v" />  
  <INPUT name="test_2.v" />  
  <OUTPUT name="output.v" />  
</OPERATOR>
```



---



---

## Appendix: RETL Data Types

Numbers work differently within RETL than they do within databases. This can cause a good deal of confusion. In short, RETL does not currently support arbitrary precision math and therefore has a limited amount of precision to deal with.

### RETL Data Type Properties

This table shows the size and precision that each RETL type has available:

RETL Type	Max Size	Max Precision	Validation/Numeric Range
DFLOAT	25	15	Min= -1.7976931348623157E+308 Max= 1.7976931348623157E+308
SFLOAT	25	6	Min=-3.40282347e+38 Max=3.40282347e+38
INT8	4	4	Min=-128 Max=127
INT16	6	6	Min=-32768 Max=32767
INT32	11	11	Min=-2147483648 Max=2147483647
INT64	20	20	Min=-9223372036854775808 Max=9223372036854775807
UINT8	3	3	Min=0 Max=255
UINT16	5	5	Min=0 Max=65535
UINT32	10	10	Min=0 Max=4294967295
UINT64	20	20	Min=0 Max=18446744073709551615
DATE	8	8	Is specified in YYYYMMDD where: YYYY=0000-9999 MM=01-12 DD=01-31

RETL Type	Max Size	Max Precision	Validation/Numeric Range
TIME	8	8	Is specified in HH24MISS format where: HH24=00-23 MI=00-59 SS=00-59
TIMESTAMP	16	16	Is specified in YYYYMMDDHH24MISS format where: MM=01-12 DD=01-31 YYYY=0000-9999 HH24=00-23 MI=00-59 SS=00-59

Note the Max Precision column above. For integral types, the precision is roughly equivalent to the size (the caveat being the min and max values in the table); however, for floating point types (sfloat and dfloat), the precision indicates how many total digits to the left of the decimal point (assuming that the float is written in scientific form) you can look at before rounding errors are encountered.

Thus, a database column defined as NUMBER(12,4) will have a maximum length of 12, with 4 digits of precision right of the decimal place. For this datatype, an SFLOAT field can be safely used to represent all numbers in that range.

## RETL Data Type/Database Data Type Mapping

The following tables show the mapping RETL uses when reading from or writing to a database.

Database data types that are merely synonyms for other data types are supported. For example, a REAL in Oracle is actually a synonym for FLOAT(63), which is supported.

---

**Note:** For integral RETL data types, not all possible values of the RETL data type can be stored in the mapped Oracle data type. For example, int8 can hold the value 100, but the mapped Oracle data type, NUMBER(2,0) cannot.

When the data source is an Oracle table (directly or indirectly) the incomplete range handling does not pose a problem because ORAREAD chooses RETL data types that map back to the same or larger Oracle data types in ORAWRITE.

If the data source is a flat file and the destination is an Oracle table, make sure to choose the RETL data type that maps to an Oracle data type that can hold the entire range of values. Otherwise, ORAWRITE (through SQL\*Loader) will reject a record if it contains a field with a value outside of the Oracle data type's range.

---

## RETL Data Type to Oracle Data Type (ORAWRITE)

RETL Data Type	Oracle Data Type	Range
int8	NUMBER(2,0)	[-99, 99]
int16	NUMBER(4,0)	[-9999, 9999]
int32	NUMBER(9,0)	[-999999999, 999999999]
int64	NUMBER(18,0)	[-9999999999999999, 9999999999999999]
uint8	NUMBER(2,0)	[0, 99]
uint16	NUMBER(4,0)	[0, 9999]
uint32	NUMBER(9,0)	[0, 999999999]
uint64	NUMBER(19,0)	[0, 99999999999999999]
dfloat	NUMBER	[-1.7976931348623157E+308, 1.7976931348623157E+308]
sfloat	NUMBER	[-3.40282347e+38, 3.40282347e+38]
string	VARCHAR	N/A
date	DATE	[1/1/0001, 12/31/9999]
time	DATE	[00:00:00, 23:59:59]
timestamp	DATE	[00:00:00 1/1/0001, 23:59:59 12/31/9999]

## Oracle Data Type to RETL Data Type (ORAREAD)

Oracle Data Type	RETL Data Type
CHAR, VARCHAR	string
DATE	date or timestamp, depending on value of datetotimestamp property
FLOAT	precision = 0: dfloat
NUMBER	scale > 0: dfloat
INTEGER	precision > 19: dfloat
DECIMAL	precision = 19: uint64
	precision >= 10, < 19: int64
	precision >= 5, < 10: int32
	precision >= 3, < 5: int16
	otherwise: int8



---

---

## Appendix: Database Configuration and Troubleshooting Guide

Since RETL database operators work with database servers and utilities, validating RETL requires some background in database setup and administration. Specifically this involves database setup, RETL database log in ID privileges, and database connectivity.

### RETL Database Configuration and Maintenance Notes

RETL uses specific tools to access the different databases. RETL developers should be sure that their database administrator is aware of what RETL does, and how the database needs to be set up to support RETL activities. The following are notes particular to certain databases and platforms.

#### Debugging Database ETL Utilities

The DBWRITE operators use the database specific import/export utilities. If you run into problems specific to the database or database operators, take a look at the log files or control files generated by the database operators. See the topic "How do I tell what commands are being sent to the database?" in "[Appendix: Troubleshooting Guide](#)." You may be able to better diagnose these problems by checking the utility's limitations, restrictions, environment, and privileges to see if there is a configuration error, either with the database or within the RETL flow.

#### Database Semaphore Problems

Semaphores from database utilities and ODBC instances invoked by rfx may not be properly cleaned up if an rfx process or job is killed from UNIX. To find and clean up semaphores, use `ipcs -s` and `ipcrm -s` respectively.

#### Runaway Loader Processes

Sometimes when rfx is killed, some ancillary database utilities are left running. These can be killed with the normal UNIX kill command. The name of the Oracle utility is `sqlldr`.

## Troubleshooting RETL with Your Database

The following RETL/ database setup validation steps will help you work with your DBA and system administrator to verify and troubleshoot the RETL/database installation.

In the following list, note the commands that should be executed from the command line. These are written in Courier font. Additionally, variables such as machine names and IP addresses that should be replaced by values particular to your environment are surrounded in <> like this: <variable name>.

1. Run `verify_retl`.

If you have not already done so run the `verify_retl` script and correct any problems as directed by the script. When this is running properly proceed to step 2.

2. Verify that your databases are set up properly.

For all databases go through RETL Database Configuration Notes to make sure that the database is setup properly.

3. Verify database connectivity.

Check the connection between the rfx local machine and the remote database machine, if the database and rfx are not on the same machine. We recommend that RETL be located on the database machine for better performance.

```
ping <database machine name>
```

If this fails, try pinging the IP address directly:

```
ping <database machine IP address>
```

If this works, then your DNS is not set up properly. Contact your network administrator to fix the problem by adding the IP address to the DNS.

If this second ping command fails then you need to contact your network administrator to determine why you cannot contact the machine. Until this problem is resolved, RETL will not work.

4. Check Oracle database server connectivity using the following command:

```
tnsping <Oracle Server Name>
```

If this fails, contact to your DBA for TNS name setup.

5. Verify database log in ID/password.

Log in to the database by using the following:

```
sqlplus userid/password@oracle_database_server
```

Contact to your DBA if this step fails; your userid/password are not set up correctly.

6. Check database user ID privileges.

Contact your DBA to ensure that the RETL user has privileges to select, insert, update, create table, and so on. Because RETL database operators invoke database utilities to extract and load (see "[Chapter 4, "Database Operators"](#)" for details), your RETL DB user needs privileges to meet all of the utilities requirements.

These are the privileges that each of the operators requires:

Operators	Privilege
ORAREAD	select
ORAWRITE	select, insert, update, create table, drop table, load

It is recommended that the RETL/database log in ID is set up with create/drop table privileges during RETL/database setup phase. so that we can run the RETL/database test flows to verify that the database is set up properly for RETL.

7. Run RETL/database testing scripts.

Note that these scripts require that the rfx userid has full privileges to run. If you do not have the appropriate privileges to change tables, you can still run the read scripts to verify that you can read from the database without problems (go to step 10).

```
$RFX_HOME/samples/verify_db/oracle.ksh <database name> <userid> <password>
```

These scripts:

- Create a RETL\_test table with two columns and two rows on the database owned by the given userid.
- Read the two rows from the database and verify that the rows contain the correct information.
- Compare the results from read and verify what was supposed to be written.

In short, this test verifies both RETL read operator and RETL write operator appropriate to the database type.

If the scripts are working properly you see something like the following:

```
Testing a write to etlsun9i
Reading data back out of etlsun9i
Comparing data received from etlsun9i and expected output...
Test Passed!
```

If both tests succeed, congratulations! Your database seems to be properly configured. Skip the following steps.

If the scripts are not working properly you will see the following:

```
Testing a write to etlsun9i
Exception in operator [oraread:1]
Error Message
Reading data back out of etlsun9i
Exception in operator [orawrite:1]
Error Message
Comparing data received from <database name> and expected output...
diff: filename: No such file or directory
Changes were detected! Test failed!
```

If either test fails, you need to move onto the following steps to try and diagnose the problem. If the ORAWRITE operator passed but the ORAREAD operator did not, move onto Step 9. If both of the tests failed, then move onto Step 8.

**8. Debug an ORAWRITE failure.**

This step verifies that the database write operator is working properly.

```
orawrite.ksh <database_name> <userid> <password>
```

This flow works if the last line of output shows "Flow ran successfully". Go to Step 9.

Here are some common failures and their causes:

```
Error Connecting to the database()
```

You have probably typed an invalid database name, userid, or password. See "When running rfx I get the following error: Error connecting to the database" in "[Appendix: Troubleshooting Guide](#)." Verify that these parameters are correct (See Step 4 above) and try again.

If you still cannot find the problems, you can try running step 10 on a table that you know already exists.

**9. Debug an ORAREAD failure.**

This step will verify that the database read operator is working properly.

```
oraread.ksh <database name> <userid> <password>
```

If the last line of output shows "Flow ran successfully," your RETL/database setup is working for reads-congratulations! You can run step 10 if you would like to try extracting data from your own tables.

If it fails, there are several possible reasons (generally the same as those for the write operator). Here are some common errors:

```
Error Connecting to the database()
```

You have probably typed an invalid database name, user ID, or password. See "When running rfx I get the following error: 'Error connecting to the database'" in "[Appendix: Troubleshooting Guide](#)." Verify that these parameters are correct (see Step 4) and try again.

If you are still running into problems and there are valid tables that you know work, try Step 10. Otherwise review the RETL Database Configuration Notes, to verify that everything is set up correctly, and try again.

**10. Run the ORAREAD script against user tables.**

The database read scripts could be used to access any table to which you have read privileges by simply specifying the table name as the last parameter.

```
oraread.ksh <database> <userid> <password> <table>
```

If the last line of output shows "Flow ran successfully," your RETL/database setup is working for reads on this table.

If you still cannot find out the problems, contact Oracle Retail Support:

<https://metalink.oracle.com>

---

---

## Appendix: Troubleshooting Guide

This troubleshooting guide provides RETL installation, development, and operation suggestions for common support issues. The suggested resolutions are in response to product and documentation feedback from RETL users to the RETL product team.

### **When trying to run rfx I get this error: ksh: rfx: cannot execute.**

RETL is not installed or configured properly. Try running `verify_retl` from the command line to get more information about the problem. Reinstalling RETL may also resolve the problem.

### **Are there any suggestions for how best to develop and debug flows?**

Yes, there are several important steps you should take when developing and debugging to keep you on the right track:

1. Do not optimize until you have a working flow. Specifically, set `numpartitions` to 1. Do not use the HASH operator. Worry about tuning for performance after your initial development is complete.
2. When building new flows, start small and work your way up. Add another operator only after you have tested the smaller flow. Back up your work often, so that you can look at the differences between flows to narrow problems to smaller segments.
3. As you are coding, copy and paste from examples that you know work, to avoid syntax and XML errors if you are using a text editor.
4. When you run into problems, try breaking up your flows. Insert DEBUG or EXPORT operators to show that your flow is working to a certain point (think binary search here). Move these operators further down the flow until you figure out what is causing the flow to break. A feature that has been added since version 10.3 is the ability of RETL to print schema files for each dataset connecting operators (by adding `-sSCHEMAFILE` to the command line options). This greatly aids when breaking up and debugging flows, by printing out each input/output schema in schema file format. See the section "[rfx Command Line Options](#)" in "[Chapter 2, 'Installation and System Configuration'](#)" for more information about the `-sSCHEMAFILE` option.
5. As of release 10.3, you can look at the log file to see when operators start and stop and how many records are flowing through them. Turn on more details by specifying a higher level in the LOGGER section of the configuration (see "[Configuration](#)" in [Chapter 2, 'Installation and System Configuration'](#)" for more details).

---

**When running rfx I get the following error: Error connecting to the database.**

You have probably typed an invalid database name, userid, or password. Verify that these parameters are correct (See steps 1-4 of "[Troubleshooting RETL with Your Database](#)" in "[Appendix: Database Configuration and Troubleshooting Guide](#)" and try again.

**How do I tell what commands are being sent to the database?**

Previous versions of RETL streamed database commands to the console by default. To see these commands starting from version 1.7, you should ensure that the environment variable RFX\_SHOW\_SQL is defined within your environment. Assuming that you are using ksh, you can do this by putting the following statement into your .profile or .kshrc, for example:

```
export RFX_SHOW_SQL=1
```

**I got an error message about operator "sort:3". How do I figure out which operator "sort:3" is?**

RETL names operators internally based upon the type and their position within the specified XML flow. Assuming that numpartitions is set to 1, sort:3 is the fourth sort operator in the XML flow (RETL starts counting at 0).

This becomes a bit more complex if you are using partitioning (numpartitions>1). In short, partitioning splits data and creates parallel data streams within the flow. In effect, this multiplies the number of operators being used, depending on the position within the flow and other factors.

There are two ways to determine which operator RETL is referring to in this case. The easiest way is to set numpartitions to 1 and count the operators of that type within the flow. If you cannot do this, then run the flow again, except specify the -g option. This generates two .dot files that contain the flow before and after partitioning. In order to view these, you need to download a copy of dotty, which is included with the graphviz package available from AT&T Bell Labs:

<http://www.research.att.com/sw/tools/graphviz/download.html>

**I've hashed and sorted my data, but for some reason my output data is not sorted. What is wrong?**

Your partitioning is probably set to a value greater than 1. Hashing your data allows RETL to break up your data and parallelize your flow into separate data streams. Later in the flow, RETL will rejoin these separate data streams. By default, RETL will do this by using the COLLECT or FUNNEL operator. However, these operators do not keep the dataset in sorted order, hence your unsorted output. You can correct this problem by explicitly using a SORTCOLLECT to rejoin the data whenever you are rejoining sorted data.

This problem is also avoided by not partitioning (set partitioning to 1, or do not specify a HASH). As always, we recommend that you keep flows as simple as possible until or unless you need the performance boost that partitioning can give you.

**I'm using sortcollect (or sortfunnel), but my data is not sorted! What is wrong?**

SORTFUNNEL performs a merge sort. It requires sorted input and can merge multiple inputs maintaining sorted order. It cannot perform a full sort of the data. You should use SORT to do a full sort of a single data stream. SORTFUNNEL can be used to merge two data sources that are already sorted on the same keys; the result is a single data stream sorted on the same key.

---

### **I am missing fields within my flow. Where did they go?**

Often the cause of this problem is that a schema file is incorrect (search for missing quotes, extra quotes), or the database schema is not what you expected. If your flow is not working correctly, take a look at the schemas that are output to stdout when RETL first starts up (specify the `-s` option if necessary to make RETL output the schemas). Look carefully at the schemas displayed for each operator within the flow. Make sure that these display ALL of the fields that you expect to see for this part of the flow. If you are missing fields or fields are of the wrong type, trace backwards to where the schema originates, to try to find the source of the problem.

If you are using partitioning and are seeing operators that you do not expect, refer to the preceding topic: 'I got an error message about operator "sort:3". How do I figure out which operator "sort:3" is?' This section helps you understand exactly what operators RETL is seeing so that you can more easily find the source of your problem.

### **How do I filter values from two different tables?**

Currently, FILTER assumes that all values required for the filter are in one record. This makes it a bit more difficult to compare values from multiple tables. The way to handle this is to join the two tables, and then filter the values after the join.

### **How do I filter out a set of dynamic values from a data stream?**

How do I filter out a set of values? For example, I want to remove any records whose FIRST\_NAME column does not have a value in the set {Joe, Susan, Mary} and keep all the rest.

This is easily done with RETL using the LOOKUP operator. Simply use the filter set as the key for the lookup. Import a flat file or table with a single column describing the values to filter (FIRST\_NAME = {Joe, Susan, Mary}), and use that as the input for the lookup table of the lookup. Any matches are discarded and any misses are processed.

The actual lookup would look something like this:

```
<OPERATOR type="lookup">
  <PROPERTY name="tablekeys" value="FIRST_NAME"/>
  <PROPERTY name="ifnotfound" value="continue"/>
  <INPUT name="records_to_filter.v"/>
  <INPUT name="set_names_to_drop.v"/>
  <OUTPUT name="records_in_set.v"/>
  <OUTPUT name="records_not_in_set.v"/>
</OPERATOR>
```

### **How do I translate from database types (for example, NUMBER(12,4)) to RETL types?**

See "[Appendix: RETL Data Types](#)" for tables that show the conversion from database data types to RETL data types and vice versa.

### **What RETL data types should I use when importing a file?**

We recommend that you use the biggest numeric type you can (int64, uint64, or dfloat) to ensure that data is properly preserved and the flow works for as many different sets of data as possible. If you know your data very well and are sure that you can use a small type, you may get a bit of a performance boost by making this a restriction. RETL warns you on the IMPORT if your data cannot be represented by the type that you have chosen.

See "[Appendix: RETL Data Types](#)" for a table that shows the range, size and precision for each datatype.

---

## Why do I receive a warning about the output schema length being less than the input schema length?

For example:

```
WARN - W115: export:1: Output schema length for field LOC_IDNT (10) is less than
the input schema length (25).
Data WILL NOT be truncated to match the output schema and the output data file may
be improperly formatted.
```

This warning is saying that the entire range of values that can be stored in the RETL data type requires more characters than are specified in the export schema. To prevent data loss, the entire value will be written to the file, thereby breaking the file's fixed-length format.

The warning can be encountered when exporting NUMBER fields read from a database. If you have a column defined as NUMBER(10), you need 11 characters (10 + 1 for the negative sign) to hold all possible values. However, RETL stores NUMBER(10) values internally as int64. See "[RETL Data Type/Database Data Type Mapping](#)" in "[Appendix: RETL Data Types](#)." The range of values for an int64 (9223372036854775808 to 9223372036854775807) requires 25 characters.

This warning can be safely ignored if you know that the data length will not exceed the width specified in the export schema. Exporting NUMBER columns from a database as outlined above is one example.

### Is there a common set of properties for all database operators?

Originally, RETL was developed to try to remain consistent with each of the database vendors' terminology. However, this made it difficult to write code that was portable between the different databases and to understand how to write flows for different databases. Starting in release 10.2, the following common properties have been added to each of the database operators:

- userid
- password
- tablename

Over time, we will attempt to merge the database operators so that RETL is database independent without the use of shell variables.

### In my query, I am explicitly selecting the timestamp field as YYYYMMDDHH24MISS, but the data in the output file reads "invalid date." What am I doing wrong?

Currently RETL cannot read DATE, TIME and TIMESTAMP types directly out of the database. Because there is no conversion from string to these fields, if you must access these fields as DATE, TIME, or TIMESTAMP, you will have to export your data to a file and then import it as a date field.

### When attempting an ORAWRITE, I get the following error: "SQL\*Loader-951: Error calling once/load initialization."

When you use the ORAWRITE operator with the method property set to "direct", you may get the following error message in your SQL\*Loader log file:

```
SQL*Loader-951: Error calling once/load initialization
ORA-26028: index name.name initially in unusable state
```

This is because SQL\*Loader cannot maintain the index that is in IU state prior to the beginning of a direct path load. Fix this problem by re-creating the table index.

---

### How do I request help when all else fails?

You should contact Oracle Retail Support at <https://metalink.oracle.com> when requesting help with a flow. Follow these guidelines to get a faster response:

1. Minimize the scope of the problem. Submit the minimal flow with the minimal set of data and minimal schema that still reproduces the problem. This helps to ensure that the support team does not waste time trying to understand a long complicated flow or the data involved, allowing us to focus on the true problem.
2. If you have a database operator (for example, ORAREAD, ORAWRITE), separate the data from the database. Usually the database is not the problem; therefore, removing the database operator from the equation makes our job easier. Replace the DB operator (for example, ORAREAD, ORAWRITE) with an IMPORT or EXPORT operator. Use the exact same schema as in the database. Extract a sample set of data into a flat file for the IMPORT operator. This should allow the support team to reproduce your problem here and debug it if necessary. This drastically reduces the amount of time it would take to look into the problem otherwise.
3. If removing the database operator removes the problem, then look very carefully to make sure that there is not a syntax error in the database operator. If you cannot find anything there, look at the IMPORT schema, and make sure that it matches the database schema exactly (including null fields). Finally, look at the flat file that the IMPORT operator is using. Make sure that you are using the correct delimiter, that the delimiter is not within the data, and that the data is valid. Try setting the rejectfile property to see if there are any rejected records. If all this looks good, then there may be a problem with the DB operator. Contact us and let us know what you have found.
4. Submit the flow and all data files associated with the flow, so that the support team can run "rfx -f <yourflowname>" to reproduce the problem, along with the log file from verify\_retl so that they can reproduce your environment.

### I've upgraded and now I get WARNINGS when I never did before. What's wrong!?

Each version of RETL brings better identification of errors, and what might have been overlooked previously is now being caught. WARNING messages are generated if the property names or values do not match the spelling/case noted in the Programmer's Guide. It is very important to run the latest version of RETL and correct the warnings, since they indicate that parameters are incorrect. If incorrect, they can produce unexpected results or corrupt data. In the future, RETL will stop processing and display error messages, rather than just providing warnings, to stop the flow until the problem is fixed. Warnings are being used as an intermediate step, to minimize the impact to existing product.



---



---

## Appendix: Default Conversions

These are the default conversions for use by the CONVERT operator.

### Default Conversions from UINT8

Conversion Name	Description
default	Converts to INT8 from UINT8.
default	Converts to UINT16 from UINT8.
default	Converts to INT16 from UINT8.
default	Converts to UINT32 from UINT8.
default	Converts to INT32 from UINT8.
default	Converts to UINT64 from UINT8.
default	Converts to INT64 from UINT8.
default	Converts to FLOAT from UINT8.
default	Converts to DFLOAT from UINT8.

### Default Conversions from INT8

Conversion Name	Description
default	Converts to UINT8 from INT8.
default	Converts to UINT16 from INT8.
default	Converts to INT16 from INT8.
default	Converts to UINT32 from INT8
default	Converts to INT32 from INT8.
default	Converts to UINT64 from INT8.
default	Converts to INT64 from INT8.
default	Converts to FLOAT from INT8.
default	Converts to DFLOAT from INT8.

## Default Conversions from UINT16

Conversion Name	Description
default	Converts to UINT8 from UINT16.
default	Converts to INT8 from UINT16.
default	Converts to INT16 from UINT16.
default	Converts to UINT32 from UINT16.
default	Converts to INT32 from UINT16.
default	Converts to UINT64 from UINT16.
default	Converts to INT64 from UINT16.
default	Converts to FLOAT from UINT16.
default	Converts to DFLOAT from UINT16.

## Default Conversions from INT16

Conversion Name	Description
default	Converts to UINT8 from INT16.
default	Converts to INT8 from INT16.
default	Converts to UINT16 from INT16.
default	Converts to UINT32 from INT16.
default	Converts to INT32 from INT16.
default	Converts to UINT64 from INT16.
default	Converts to INT64 from INT16.
default	Converts to FLOAT from INT16.
default	Converts to DFLOAT from INT16.

## Default Conversions from UINT32

Conversion Name	Description
default	Converts to UINT8 from UINT32.
default	Converts to INT8 from UINT32.
default	Converts to UINT16 from UINT32.
default	Converts to INT16 from UINT32.
default	Converts to INT32 from UINT32.
default	Converts to UINT64 from UINT32.
default	Converts to INT64 from UINT32.
default	Converts to FLOAT from UINT32.
default	Converts to DFLOAT from UINT32.

## Default Conversions from INT32

Conversion Name	Description
default	Converts to UINT8 from INT32.
default	Converts to INT8 from INT32.
default	Converts to UINT16 from INT32.
default	Converts to INT16 from INT32.
default	Converts to UINT32 from INT32.
default	Converts to UINT64 from INT32.
default	Converts to INT64 from INT32.
default	Converts to FLOAT from INT32.
default	Converts to DFLOAT from INT32.

## Default Conversions from UINT64

Conversion Name	Description
default	Converts to UINT8 from UINT64.
default	Converts to INT8 from UINT64.
default	Converts to UINT16 from UINT64.
default	Converts to INT16 from UINT64.
default	Converts to UINT32 from UINT64.
default	Converts to INT32 from UINT64.
default	Converts to INT64 from UINT64.
default	Converts to FLOAT from UINT64.
default	Converts to DFLOAT from UINT64.

## Default Conversions from INT64

Conversion Name	Description
default	Converts to UINT8 from INT64.
default	Converts to INT8 from INT64.
default	Converts to UINT16 from INT64.
default	Converts to INT16 from INT64.
default	Converts to UINT32 from INT64.
default	Converts to INT32 from INT64.
default	Converts to UINT64 from INT64.
default	Converts to FLOAT from INT64.
default	Converts to DFLOAT from INT64.

## Default Conversions from SFLOAT

Conversion Name	Description
default	Converts to UINT8 from SFLOAT.
default	Converts to INT8 from SFLOAT.
default	Converts to UINT16 from SFLOAT.
default	Converts to INT16 from SFLOAT.
default	Converts to UINT32 from SFLOAT.
default	Converts to INT32 from SFLOAT.
default	Converts to UINT64 from SFLOAT.
default	Converts to INT64 from SFLOAT.

---



---

**Note:** Conversions to integer types truncate decimal values.

---



---

## Default Conversions from DFLOAT

Conversion Name	Description
default	Converts to UINT8 from DFLOAT.
default	Converts to INT8 from DFLOAT.
default	Converts to UINT16 from DFLOAT.
default	Converts to INT16 from DFLOAT.
default	Converts to UINT32 from DFLOAT.
default	Converts to INT32 from DFLOAT.
default	Converts to UINT64 from DFLOAT.
default	Converts to INT64 from DFLOAT.

---



---

**Note:** Conversions to integer types truncate decimal values.

---



---

---



---

## Appendix: Data Partitioning Quick Reference

This appendix describes how to set up a flow to have multiple data partitions. For more detailed information, see Chapter 7, "[RETL Parallel Processing](#)."

### Partitioning the Data

Data partitioning does not happen automatically in RETL; you must configure `rfx.conf` and the operators in the flow.

In `rfx.conf`, set the `numpartitions` attribute in the `NODE` element to a value greater than 1. This tells RETL to allow partitioning and also provides a default `numpartitions` for `HASH` and `SPLITTER`.

The following table shows the operators that partition data and what must be done to configure the operator:

Operator	Configuration
DBREAD	Set the <code>numpartitions</code> property to a value greater than 1. Specify a query for each partition.
GENERATOR (generating new records)	Set the <code>numpartitions</code> property to a value greater than 1. Set <code>partnum_incr</code> and <code>partnum_offset</code> if appropriate.
HASH	No configuration required, but <code>numpartitions</code> can be used to override the <code>numpartitions</code> from <code>rfx.conf</code> .
IMPORT	Set the <code>numpartitions</code> property to a value greater than 1. Specify either one input file or <code>numpartitions</code> input files.
ORAREAD	Set the <code>numpartitions</code> property to a value greater than 1. Specify a query for each partition.
SPLITTER	No configuration required, but <code>numpartitions</code> can be used to override the <code>numpartitions</code> from <code>rfx.conf</code> .

## Continuing the Partition

Some operators must be configured to continue a data partition. If the operator is not configured to continue the partition, RETL will end the data partition by inserting a FUNNEL to collect the partitioned records. To configure an operator to continue a data partition, set the parallel property to true.

The following operators require the parallel property to be set to true.

- CLIPROWS
- DBWRITE
- DEBUG
- EXPORT
- GENERATOR (generating new fields)
- GROUPBY
- ORAWRITE

## Hash Partitioning

Some operators only work correctly if the data is partitioned using the HASH operator. These are typically operators that work with key fields. Using a HASH operator to perform the data partitioning ensures that records with the same key end up in the same data partition.

The following operators require a data partition started by a HASH operator:

- CHANGECAPTURE
- CLIPROWS
- COMPARE
- DIFF
- FULLOUTERJOIN
- GROUPBY
- INNERJOIN
- LEFTOUTERJOIN
- MERGE
- REMOVEDUP
- RIGHTOUTERJOIN

## Ending the Data Partition

RETL will continue a data partition until an operator is encountered that does not support data partitioning (or is configured to not continue a data partition). At this point, RETL inserts a FUNNEL to unpartition the records back into a single dataset.

Because FUNNEL and SORTFUNNEL do not support partitioning, you can insert a FUNNEL or a SORTFUNNEL at any point to end a partition.

---

---

## Appendix: Frequently Asked Questions

This appendix provides answers to frequently asked questions, such as the following.

### **Does RETL call stored procedures?**

RETL can call stored procedures only within the ORAWRITE operator via the sp\_prequery and sp\_postquery, which run before and after the ORAWRITE operation. These stored procedures can take static variables (not data from the flow) as parameters. We anticipate expanding on this feature in the future, to allow lookups based upon flow data and support for other database operators.

### **Can RETL handle text translations (for example, between EBCDIC and ASCII)?**

Currently, RETL does not do textual translation of this type.

### **Can RETL handle data translations (for example, between numbers and strings)?**

Data translations between different fields can be done with the LOOKUP and JOIN operators.

### **Can clients obtain RETL performance metrics on RETL?**

Generally, clients are interested in how a particular RETL flow runs, rather than how RETL itself performs at a fundamental level. These are generally handled on a case-by-case basis for the different product groups, because flows are very different from product to product, and the environment is different for hardware and configuration. Contact your product group to determine if they have benchmarking numbers specific to your group. The RETL group does benchmark internal performance results, and these are made available through the Product Enhancement Review Board. Contact your product group for inclusion to this board.

### **Is it possible to join data from two different databases?**

Yes. This is generally done with two ORAREAD operators and then a JOIN operator.

### **Does RETL have error handling capabilities?**

There are many other areas where RETL does error handling, and this capability has been expanding.

Flow-related problems. These are problems that an RETL flow developer might run into while developing flows using RETL. There are many different areas that RETL detects; because these are generally found only during development, RETL usually exits with an error message describing one of these scenarios:

- Schemas—Mismatched input schemas (for example, for FUNNEL), invalid schema specified (for example, record length does not equal sum of field lengths).
- Mismatched datasets—RETL enforces one-to-one correspondence between inputs and outputs.
- Validation of operators, properties and property values.

Data-related problems. These are found after the flow is parsed and execution begins. In these cases invalid or non-conforming data may be siphoned into a separate file for later processing:

- Improperly formatted fields via the IMPORT operator (invalid types, lengths, etc).
- Improperly formatted records via the IMPORT operator (missing fields, etc).
- Invalid schemas.

Additionally, flow developers can insert data validation to verify that various database constraints are upheld, through a combination of LOOKUP, JOIN, and FILTER operators.

### **Does RETL create temp files? How much space do they use?**

RETL generates temporary files dynamically to sort and store intermediate data. These files are generally prefixed with "rfx", but this is not always the case. Therefore, place separate rfx files into their own tmp (for example, /tmp/rfx) directory, so that they can be easily purged on a regular basis.

How much space rfx will consume depends upon the complexity of the flows and how large the source system is. A good starting point is twice the size of the data being manipulated. A more conservative number is three times the size of the data being moved. Many people start with 10GB and move up from there if needed.

However, a more precise way of calculating the minimum size of temporary space for a particular flow would be as follows:

Total temp space for a particular flow =

$$\begin{aligned} & ((\text{size of data file}) * (\# \text{ sort operators}) * 2) \\ & + ((\text{size of data file}) * (\# \text{ operators that need to expand datasets})) \\ & + ((\text{size of data file}) * (\# \text{ database write operators})) \end{aligned}$$

where:

*# sort operators* is the number of sort operators in a flow.

*# operators that need to expand datasets* are the operators that page to disk in a "diamond" flow that contains circular loops.

*# database write operators* is the number of database write operators, such as ORAWRITE.

For example, on a flow IMPORT->SORT->ORAWRITE on a 2Gig input file would be as follows:

$$(2\text{Gig} * 1 * 2) + (2\text{Gig} * 0) + (2\text{Gig} * 1) = 6 \text{Gig}$$

---



---

**Note:** It is recommended that system administrators clean up temporary files in the RETL temporary directory on a regular basis.

---



---

---

## How do I participate in the definition of RETL requirements and future direction?

The Product Enhancement Review Board (PERB) is the forum to assist Oracle Retail product strategy in defining future enhancements to RETL. If you are not part of the PERB, contact your point person or team lead to inquire about joining.

## How do I remove warnings about sorting/sort order, such as the following?

```
[<operator>]: Warning: mismatched sort order - <operator> specifies "ascending"
sort order and for the first INPUT to <operator> records are in "descending".
These should match - otherwise you may get unexpected results!
```

```
[<operator>]: Warning: Input to <operator> does not seem to be sorted! Data should
be sorted according to the proper keys or you may get unexpected results!
```

Certain operators require that incoming data be sorted. To remove these error messages, place a sort operator that sorts by the required keys in the required order, before the operator in question. If you know that your data is always going to be sorted according to a particular set of keys, you can specify this in the schema file (see the schema file documentation for more details on this). This last feature should be used with care, because if data turns out to not be in the presumed sorted order, certain operators may produce unexpected results.

## Which operators require sorted input?

You should read this document in its entirety to assess the need for sorted input. However, the following is a list of operators that do currently require sorted input:

- INNERJOIN
- LEFTOUTERJOIN
- RIGHTOUTERJOIN
- FULLOUTERJOIN
- SORTFUNNEL
- REMOVEDUP
- COMPARE
- CHANGECAPTURE
- DIFF
- GROUPBY (only if "key" properties are specified)

## How should I handle errors that occur when RETL tries to connect to the database.

### For example:

```
--- SQLException caught ---
```

```
java.sql.SQLException: Io exception: Connection
refused(DESCRIPTION=(TMP=) (VSNNUM=153092608) (ERR=12505) (ERROR_
STACK=(ERROR=(CODE=12505) (EMFI=4))))
```

```
Message: Io exception: Connection
refused(DESCRIPTION=(TMP=) (VSNNUM=153092608) (ERR=12505) (ERROR_
STACK=(ERROR=(CODE=12505) (EMFI=4))))
```

```
XOPEN SQL State: null
```

```
Database Vendor Driver/Source Error Code: 17002
```

---

The hostname or port of the database read/write operator may be incorrect or not specified, or the database is down or nonexistent. Make sure that the database instance is up and running and located on the hostname/port specified in the flow or rfx.conf. If you are using an Oracle database, check tnsnames.ora, or use tnsping to figure out the hostname/port.

**Why is my database field being translated to a RETL dfloat type?**

If the SQL query contains a union, the precision of the NUMBER/DECIMAL fields is reported as 0. To ensure that the RETL data type is large enough to hold the data, dfloat is used.

**RETL is not handling international data correctly. What's wrong?**

RETL relies on the locale being set correctly in the environment. Make sure that all locale-related environment variables are set and exported correctly.

**What is taking up so much memory?**

The LOOKUP and CHANGECAPTURELOOKUP operators store the entire lookup table in memory. Thus, if the lookup table has a large number of rows, the memory requirements will be extensive. Consider sorting the data and using an INNERJOIN instead of the LOOKUP, or a CHANGECAPTURE instead of the CHANGECAPTURELOOKUP.

**Considering RETL is written in Java, can I run RETL on other platforms that have a JRE?**

See the Compatibility Matrix in the Release Notes for supported platforms. While running RETL on platforms other than these is not supported, it is technically feasible that the core framework and operators should work on any platform that has a JRE version 1.4 or higher. We would certainly like to hear from you when/if you have successfully deployed on other platforms. If these types of requests happen with much frequency or urgency, Oracle Retail will take this into account when setting priorities for efforts for platform proliferation in future releases. Note that additional platform support comes at much lower cost than with previous versions of RETL (10.x).

**With the introduction of RETL 11.x, JDBC technology has been introduced as a mechanism for connecting and reading from and writing to databases. Does this mean I can connect to any database?**

See the Supportability Matrix in the Release Notes for supported databases. While running RETL against databases other than these is not supported, it is possible that "snapping in" a JDBC-compliant driver by using the generic dbread/dbwrite operators will work.

**What changes are required to my RETL 10.x flows so that RETL 11.x and 12.x can run them?**

RETL requirements specify that it must be backward-compatible in the XML flow interface. As a result, only minor changes will need to be made. Read this document in its entirety to assess the need for these changes.

**If the database in question is non RAC, must we make there that the jdbcconnectionstring property is not specified?** (We would like to maintain a single config.xml (if possible) that connects to both RAC and non-RAC databases. For the non RAC database I would prefer to leave jdbcconnectionstring with some value that I know RETL will not use, such as "jdbc:oracle:thin.")

Yes, if the database is non RAC, make sure there is no jdbcconnectionstring property specified.

---

**If the database in question is RAC, I would like to confirm the following in jdbcconnectionstring:**

```
<PROPERTY name="jdbcconnectionstring" value="jdbc:oracle:thin:@(DESCRIPTION
=(ADDRESS_LIST =(ADDRESS = (PROTOCOL = TCP)(HOST = hostname1)(PORT =
1521))(ADDRESS = (PROTOCOL = TCP)(HOST = hostname2)(PORT = 1521))(LOAD_BALANCE =
yes ))(CONNECT_DATA =(SERVICE_NAME = serviceName))" />
```

- **Must all hosts be declared in a URL, or is a single node also acceptable?**

You can connect to the database by giving a single host. But if you want to use the RAC feature, all hosts must be declared.

- **Is this setting mandatory? LOAD\_BALANCE = YES.**

LOAD\_BALANCE = YES is optional.

- **Does SERVICE\_NAME represent the node independent global dbname? For instance, it should not be qaolr131 or qaolr132?**

Yes, this is a global name, such as SID for the database.

- **How should I obtain the RAC URL?**

Using the entire string from tnsnames.ora is recommended.



---

---

## Appendix: Database Connections Quick Reference

This appendix is a quick reference on how to set up a flow to have database connections using thin clients for ORAREAD/ORAWRITE operators. For more detailed information, see Chapter 6, "[Database Operators](#)."

### Setting the Environment Variables

The variables that must be configured to connect RETL to an Oracle database using Oracle Thin driver are as follows.

- Set ORACLE\_HOME environment variable.
- Add ORACLE\_HOME/lib to LD\_LIBRARY\_PATH env variable.
- Set CLASSPATH environment variable to use Oracle JDBC driver from ORACLE\_HOME/bin.

---

---

**Note:** ORACLE\_HOME refers to the folder where Oracle Client or Database has been installed.

Example: /home/oracle/product/10.2.0. This folder will have sub-folders such as bin, sqlj, network, rdbms, jdbc, jlib, and lib.

---

---

### ORAREAD

ORAREAD performs a read from Oracle databases. RETL uses JDBC technology to read data from Oracle databases. RETL can connect to Oracle database as a thin client.

The following example illustrates the connection to the database using thin database drivers in case for the ORAREAD operator.

**Connection to database through thin database driver:**

```
<OPERATOR type="oraread">
  <PROPERTY name="sp_prequery" value="exec pre_storedproc"/>
  <PROPERTY name="dbname" value="RETLdb"/>
  <PROPERTY name="connectstring" value="username/password"/>
  <!--Note: query must be enclosed in CDATA element otherwise -->
  <!-- this query will contain invalid XML! -->
  <PROPERTY name="query">
    <![CDATA[
      select * from rtbl where col > 1
    ]]>
  </PROPERTY>
  <PROPERTY name="maxdescriptors" value="100"/>
  <PROPERTY name="datetotimestamp" value="false"/>
  <PROPERTY name="sp_postquery" value="exec post_storedproc"/>
  <OUTPUT name="test.v"/>
</OPERATOR>
```

## ORAWRITE

ORAWRITE performs a load to Oracle databases. RETL uses SQL\*Loader (sqlldr) to load the records. RETL can connect to Oracle database using thin client. By default, the operator will create the URL for database connection as a thin client.

The following example illustrates the connection to the database using thin database drivers for the ORAWRITE operator.

**Connection to database through thin database driver:**

```
<OPERATOR type="orawrite">
  <INPUT name="import0.v" />
  <PROPERTY name="threadModel" value="start_thread" />
  <PROPERTY name="dbuserid" value="username/password" />
  <PROPERTY name="dbname" value="RETLdb" />
  <PROPERTY name="tablename" value="ORG_LOC_DM" />
  <PROPERTY name="createtablemode" value="recreate" />
</OPERATOR>
```