

Retek[®] Extract Transform and Load[™] 11.2.2

Release Notes

Corporate Headquarters:

Retek Inc.
Retek on the Mall
950 Nicollet Mall
Minneapolis, MN 55403
USA

888.61.RETEK (toll free US)
Switchboard:
+1 612 587 5000

Fax:
+1 612 587 5100

European Headquarters:

Retek
110 Wigmore Street
London
W1U 3RW
United Kingdom

Switchboard:
+44 (0)20 7563 4600

Sales Enquiries:
+44 (0)20 7563 46 46

Fax:
+44 (0)20 7563 46 10

The software described in this documentation is furnished under a license agreement, is the confidential information of Retek Inc., and may be used only in accordance with the terms of the agreement.

No part of this documentation may be reproduced or transmitted in any form or by any means without the express written permission of Retek Inc., Retek on the Mall, 950 Nicollet Mall, Minneapolis, MN 55403, and the copyright notice may not be removed without the consent of Retek Inc.

Information in this documentation is subject to change without notice.

Retek provides product documentation in a read-only-format to ensure content integrity. Retek Customer Support cannot support documentation that has been changed without Retek authorization.

The functionality described herein applies to this version, as reflected on the title page of this document, and to no other versions of software, including without limitation subsequent releases of the same software component. The functionality described herein will change from time to time with the release of new versions of software and Retek reserves the right to make such modifications at its absolute discretion.

Retek[®] Extract Transform and Load[™] is a trademark of Retek Inc.

Retek and the Retek logo are registered trademarks of Retek Inc.

This unpublished work is protected by confidentiality agreement, and by trade secret, copyright, and other laws. In the event of publication, the following notice shall apply:

©2005 Retek Inc. All rights reserved.

All other product names mentioned are trademarks or registered trademarks of their respective owners and should be treated as such.

Printed in the United States of America.

Customer Support

Customer Support hours

Customer Support is available 7x24x365 via email, phone, and Web access.

Depending on the Support option chosen by a particular client (Standard, Plus, or Premium), the times that certain services are delivered may be restricted. Severity 1 (Critical) issues are addressed on a 7x24 basis and receive continuous attention until resolved, for all clients on active maintenance. Retek customers on active maintenance agreements may contact a global Customer Support representative in accordance with contract terms in one of the following ways.

Contact Method Contact Information

E-mail support@retек.com

Internet (ROCS) rocs.retek.com
Retek's secure client Web site to update and view issues

Phone +1 612 587 5800

Toll free alternatives are also available in various regions of the world:

| | |
|----------------|--|
| Australia | +1 800 555 923 (AU-Telstra) or +1 800 000 562 (AU-Optus) |
| France | 0800 90 91 66 |
| Hong Kong | 800 96 4262 |
| Korea | 00 308 13 1342 |
| United Kingdom | 0800 917 2863 |
| United States | +1 800 61 RETEK or 800 617 3835 |

Mail Retek Customer Support
Retek on the Mall
950 Nicollet Mall
Minneapolis, MN 55403

When contacting Customer Support, please provide:

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

Overview

This RETL 11.2.2 release is a patch release containing only bug fixes. The bugs that were fixed are explained in the “Changes since RETL 11.2.1” section below.



Note: Several bug fixes catch errors in flows that were not detected in earlier versions of RETL. Before upgrading to RETL 11.2.2, Retek recommends that you perform a full regression test, paying special attention to Bug #916.

Database support

The following changes to RETL’s database support have been made with the release of RETL 11.2.2:

- Teradata 2 Release 4.1 is no longer supported by RETL. It has been replaced with support for Teradata 2 Release 5.0.
- DB2 Version 7 is no longer supported by RETL. Support will continue for DB2 Version 8.1.

See the “Compatibility matrix” section below for a list of all supported platforms and databases.

Product certification

RETL 11.2.2 has been certified with RDW 10.2.3. See the “Certification” section below for a matrix of Retek products and the certified version of RETL.

Retek strongly recommends that, before any installation or development work is performed, the RETL 11.2.2 Programmer's Guide, provided with the release, is read in its entirety. In particular, read and understand “Chapter 2 – Installation and system configuration”. In addition, pay special attention to the “Known issues” section of these Release Notes.

Product groups should take care to complete all certifications and regression/volume testing with RETL 11 versions before deployment to customer sites.

The build number for RETL 11.2.2 is 558. After a successful installation of RETL 11.2.2, running RETL with the `-v` command line option should produce the following output:

```
$ retl -v
retl 11.2.2 build 558
```

Clients that took delivery of the RETL 11.2.1 release, prior releases and/or subsequent patch releases are strongly encouraged to upgrade to the RETL 11.2.2 patch release once a full regression test has been performed. Issues identified in RETL 11 releases and prior releases will be resolved with enhancements and fixes being delivered through RETL 11 patch releases. No additional releases of RETL 10.x are planned.

RETL 11.x enhancements

- Partitioning/sizing implementation
The concept of partitioning is to parallelize as much as possible, which in RETL, ideally results in a flow uninhibited by bottlenecks. The goal is to code flows so that they can be partitioned to perform within a customer's batch window.
- Major performance enhancements
 - RETL 11 releases include a new architecture that optimizes pipeline parallelism by consolidating connected operators into the same pipeline where possible. This architecture reduces the number of threads required, minimizing context switching and thrashing among threads. The framework thus supports improved performance.
 - An enhanced funneling algorithm results in significant speed increases for flows that contain the funnel operator.
 - Internal types are now represented in native byte format, which saves space and results in increased performance; strings do not have to be converted and so on.
 - There is increased performance when handling rejects from the import operator. Users can now turn off 'verbosity' of printing rejects to the screen. This action can result in vast improvements to RETL modules that reject records in import.
 - There have been significant improvements to performance in oraread and in the implementation of datasets that connect operators.
- Error handling
 - RETL 11 releases introduce better flow debugging by giving line and column specific error messages when RETL encounters problems reading an XML flow definition. Additionally, there is exception handling in RETL 11 that allows for better error handling and debugging.
 - Due to the configuration capabilities of the error threshold in orawrite, RETL can be configured to abort when reaching the maximum # of rejected records from SQL*Loader.
 - There are overall improved error messages in 11.1 as opposed to the 10.x versions.

- Debugging and logging
 - RETL visual flow graphs
Beginning with RETL 11.2, a command-line option can be used to generate a visual graph of a RETL flow, complete with operator names, link names, and thread identifications. See the RETL Programmer's Guide for more information.
 - Enhanced performance logging
RETL 11.2 now logs significant events such as query execution, sorting via gsort, loading data via native database load utilities (for example, SQL*Loader), and more.
 - Performance hotspots/bottleneck graphs
RETL now exports an HTML graph of performance activity. This graph uses visual cues to point out bottlenecks in the flow so that flow developers, performance analysts, and services consultants can diagnose the cause of slow-performing flows.
 - RETL 11 uses the log4j logging facility for all terminal and performance log output. The log4j facility is an extremely flexible open-source package maintained by the Apache Software Foundation. See <http://logging.apache.org/log4j/docs/documentation.html> and the RETL Programmer's Guide for more information on log4j.
- New operators
 - PARSER operator
Business logic can be coded in a Java-like code script, which eliminates proliferation of too many operators, improves performance by reducing the number of operators (for example, FIELDMOD, BINOP, FILTER, SWITCH, COPY), reduces maintenance costs by decreasing complexity of flows, and aids in a better sizeability when partitioning flows.
 - DBLOOKUP operator
This operator allows the flow to join records directly to a table in the database.
 - CHANGECAPTURELOOKUP operator
Identical in functionality to changecapture, but can take unsorted input, which lends itself to better partitioning and sizing.
 - SPLITTER operator
A round-robin-based partitioning operator that allows equal distribution of records to each partition.
- Documentation
 - Much-improved Programmer's Guide
 - Performance Tuning Guide
 - Best Practices Guide
- Online help
A command-line option has been added to RETL 11 that allows flow developers to view operator syntax and usage via the command-line without having to refer to this document.
- Simplified installation and configuration
RETL 11 releases consolidate code into a single binary rather than the 26 binaries of 10.x release. This design considerably eases installation. In addition, there is less environment setup that needs to be done in order to complete an installation.

- True platform independence
The new architecture of the RETL 11 releases has been written entirely in Java. As a result, RETL may now run on platforms that previously would have required significant efforts to port to. See the Programmer's Guide and the Compatibility Matrix for supported platforms. As a direct result of the platform independence, significant RETL resources in platform maintenance will be freed up to develop new features going forward.
- True database independence
With the RETL 11 releases, generic 'dbread'/'dbwrite' operators have been introduced. These operators use JDBC-compliant technology to connect to databases in a truly database-independent manner. Additionally, previous 10.x releases were unable to connect to different database types in the same flow. RETL 11 releases allow full interoperability among different databases in the same flow (for example, among DB2, Oracle, and Teradata all in the same flow). See the RETL Programmer's Guide for more information on this topic.
- Backwards compatibility with previous releases
A requirement of RETL 11 versions is that they be backward-compatible with the 10.x versions. However, RETL 11 versions are stricter on enforcement of valid XML flow interfaces, input schemas, and data fields. In some instances, the RETL 11 product will find data and/or flow errors that may have previously been unreported. The result may be more discarded or rejected records and/or error messages than previously identified.

Compatibility matrix

The following table shows the OS/database combinations that are supported and have been certified to work properly.

| OS | Version | Arch | DB | Version |
|-------|---------|-------|----------|---------|
| AIX | 5.1 | 32/64 | none | none |
| AIX | 5.1 | 64 | Oracle | 9.2 |
| AIX | 5.1 | 64 | DB2 | 8.1.3 |
| AIX | 5.1 | 64 | Teradata | 2r5.0 |
| AIX | 5.2 | 32/64 | none | none |
| AIX | 5.2 | 64 | Oracle | 9.2 |
| AIX | 5.2 | 64 | DB2 | 8.1.3 |
| AIX | 5.2 | 64 | Teradata | 2r5.0 |
| HP-UX | 11.11 | 32/64 | none | none |
| HP-UX | 11.11 | 64 | Oracle | 9.2 |
| HP-UX | 11.11 | 64 | Oracle | 9.0.1 |
| HP-UX | 11.11 | 32 | Oracle | 8.1.7 |
| HP-UX | 11.11 | 64 | Teradata | 2r5.0 |
| HP-UX | 11.23 | IA64 | None | None |
| HP-UX | 11.23 | IA64 | Oracle | 9.2 |

| OS | Version | Arch | DB | Version |
|---------|---------|-------|----------|---------|
| Solaris | 8 | 32/64 | none | none |
| Solaris | 8 | 64 | Oracle | 9.2 |
| Solaris | 8 | 32/64 | Oracle | 9.0.1 |
| Solaris | 8 | 32 | Oracle | 8.1.7 |
| Solaris | 8 | 32 | DB2 | 8.1.3 |
| Solaris | 9 | 32/64 | none | none |
| Solaris | 9 | 9 | Oracle | 9.2 |
| Solaris | 9 | 32/64 | Oracle | 9.0.1 |
| Solaris | 9 | 32 | Oracle | 8.1.7 |
| Solaris | 9 | 32/64 | Teradata | 2r5.0 |
| Solaris | 9 | 64 | DB2 | 8.1.3 |



Note: RETL 11 only requires one binary and installation for all database/platform combinations listed above.

Certification

The following table lists products that have been certified to work with a specific RETL version.

| Product | RETL Version |
|-------------|--------------|
| RMS 11.0.2 | RETL 11.2.1 |
| RPM 11.0.1 | RETL 11.2.1 |
| ReIM 11.0.2 | RETL 11.2.1 |
| RDW 11.0 | RETL 11.2.1 |
| RDW 10.2.3 | RETL 11.2.2 |

Installation

See the RETL Programmer's Guide for installation and configuration instructions.

Package contents

The following is a brief description of the contents of retl.11.2.2.tar.gz:

bin/

Includes executables required for running retl on different platforms.

| File | Description |
|----------------------|--|
| gsort.SunOS | gsort for SunOS |
| gsort.HP-UX | gsort for HP-UX |
| gsort.AIX | gsort for AIX |
| retl | RETL executable |
| rfx | Symbolic link to retl provided for backwards compatibility |
| verify_ret1 | The RETL verification tool |
| README.verify_ret1 | The README for verify_ret1 |
| install_graphing.ksh | Installs the GraphViz graphing package |
| install_jre.ksh | Installs the Java Runtime Environment |
| test_funcs.ksh | Used by verify_ret1 to run test flows |



Note: There is only one binary (retl/rfx) required for all database and platform combinations

lib/

These are jar files that are needed in order to run retl.

| File | Description |
|-------------------|--------------------------|
| retl.jar | RETL runtime library |
| db2java.zip | DB2 JDBC Driver |
| log4j-1.2.9.jar | log4j Logging Facility |
| ojdbc14.jar | Oracle JDBC driver |
| teradata.jar | NCR Teradata JDBC driver |
| xercesImpl.jar | Xerces XML Parser |
| xml-apis.jar | XML APIs |
| xmlParserAPIs.jar | XML APIs |
| tools.jar | Future enhancements |

| File | Description |
|----------------|---------------------|
| activation.jar | Future enhancements |

etc/

This directory contains the documentation associated with RETL.

| File | Description |
|-------------|--|
| rfx.conf | The default retl configuration file. |
| logger.conf | The default logger configuration file. |

samples/

This directory contains the samples packaged with RETL. See the README in the samples/ directory for more information about the samples packaged with RETL. These are the samples that are run after the install, by verify_retl.

graphing/

This directory contains the Graphviz application suite for graphical visualization of RETL flows. See the "Producing graphical output of flows with RETL" section of the RETL Programmer's Guide.

JRE/

This directory contains the Java Runtime Environment (JRE) required by RETL.

Known issues

General known issues

- RETL does not handle arbitrary precision math.
- The export operator defaults to use '|' -delimited fields if a 'schemafilename' property is not specified. This issue can cause problems if the incoming data contains the '|' character. The recommendation is to use an export schemafilename if it is possible for data to contain a '|' character.
- Join operators may use a large amount of memory when the distribution of key-equivalent records is low and the number of records is high. It is important to test join operations with appropriate data volumes on each input side to verify the maximum amount of memory your application may require.

Differences between the 10.x and 11.x versions of RETL

- RETL 11 versions are stricter on the enforcement of valid XML flow interfaces, input schemas, and data fields. In some instances, the RETL 11 product finds data and/or flow errors that may have previously been unreported. The result may be more discarded or rejected records and/or error messages than previously identified.
- An unintended feature of 10.x allowed input data files to contain more fields than are specified in the schemafile. RETL 11 does not allow more fields to exist in the data file and throws an exception should this happen. This issue is logged as bug#632 and will be fixed in a future version of RETL 11.
- Output field order between 10.x and 11.x may be different for those flows that do not specify a 'schemafile' property in the export operator.
- RETL 11.x versions require additional properties for dbread/dbwrite operators. 'hostname' and 'port' are now required properties and may be specified in rfx.conf for convenience. See the RETL Programmer's Guide for more information on syntax and usage.
- When upgrading an Oracle database, a separate instance of RETL should be used in order to test the new database. This recommendation is due to the fact that only one port can be specified in rfx.conf at a time. The other option is to use RETL with the '-c' option to use a different config file. See the RETL Programmer's Guide for more information.
- RETL 11.x versions may require more physical memory for certain flows than 10.x releases.

Backwards compatibility issues with 10.x versions

- Retek does *not* recommend upgrading to RETL 11 unless the product using RETL has been certified on RETL 11.

Known bugs in RETL 11.2.2

- Bug #511 – The port is not configurable for DB2 operators. This bug prevents RETL from running on a different server than the database.
- Bug #611 – Tests for table existence can return false matches if the user does not have sufficient privileges on a table.
- Bug #659 – A pipe delimiter is used when paging datasets to disk. This bug causes problem re-importing if there are pipes in the data.
- Bug #693 – TERAWRITE does not handle timestamp fields.
- Bug #744 – SORT does not work with newlines in data.
- Bug #759 – Invalid INPUT element results in StackOverflowError in flow parsing.
- Bug #781 – Duplicate OUTPUT causes warning and then NullPointerException.
- Bug #795 – Too few inputs to INNERJOIN throws an exception instead of showing an error.
- Bug#820 – EXPORT file schema can be broken when the input data is cached for numeric values and the export field length is less than the imported field length.
- Bug #828 – DB2WRITE NullPointerException when used with a dataset.

- Bug #840 – CLIPROWS displays incorrect message about possible incorrect results when the input data is sorted in descending order.
- Bug#856 – If a record is altered by a BINOP operator after the record has been exported and then re-exported, the second EXPORT has incorrect fields because the record's cache is not cleared by the binop.
- Bug #917 – TERAFOREAD does not support multiple queries/partitioning.
- Bug #940 – An exception occurs when partitioning DB2WRITE with method=direct.
- Bug #941 – Parallel direct datafile ignored when partitioning.
- Bug #959 – Loops in INPUTs and OUTPUTs cause infinite loop in schema resolution.
- Bug #1069 – REMOVEDUP should warn if dataset is not sorted.
- Bug #1101 – RFX_SHOW_SQL output cannot be shown without RFX_DEBUG output.
- Bug #1155 – removedup property of SORT does not work correctly when numsort > 1.
- Bug #1209 – IllegalArgumentException exception in BINOP when destination is incorrect.
- Bug #1210 – A NullPointerException is thrown in BINOP when the destination field already exists but is NULL.
- Bug #1223 – RQST00000066282: FILTER fails when using the AND operation and IS_NOT_NULL on the left hand side of the expression.

Changes

Changes since RETL 11.2.1

Bug #785 – Number of INPUTs to SORT not validated

In RETL versions 11.0 through 11.2.1, multiple INPUTs to SORT were allowed but only the first INPUT was used. This bug allowed silent “loss” of records.

In RETL 11.2, an error is reported:

```
E348: Operator sort:1 allows only one input.
```

A flow that ran without error prior versions will now report an error if this issue is encountered. If the records in the second (and subsequent) INPUT(s) to CLIPROWS are not needed, the flow should be changed to seed the INPUT(s) to a NOOP instead.

Bug #805 – Generator sequence attributes not fully validated

In RETL versions 11.0 through 11.2.1, no warning was displayed when the schema for a SEQUENCE field contained an invalid attribute.

In RETL 11.2.2, the following warning is displayed:

```
W144: Sequence field in operator "generator:1" contains an unrecognized attribute "attribute"
```

Bug #808 – SQL error information not reported on some SQL errors

In RETL versions 11.0 through 11.2.1, reporting of some SQL errors did not include the error information returned by the database. For example, when a table creation failed, the following was displayed:

```
Exception in [orawrite:1] : Cannot create the table: tablename
```

In RETL 11.2.2, error information provided by the database is included in the error message:

```
Exception in [orawrite:x] : Cannot create the table: xxx.xxxxxxxx  
Cannot execute the query CREATE TABLE etl.bug_xxx (FIELD_A  
VARCHAR(2) NOT NULL, FIELD_A VARCHAR(2) NOT NULL)  
ORA-00957: duplicate column name
```

Bug #849 – Sort order and keys properties lost by LOOKUP

In RETL versions 11.0 through 11.2.1, if the first INPUT to LOOKUP was sorted, the sort order and sort keys attributes were lost. (The records retained the same order; only the attributes were incorrect.) This bug caused incorrect warnings later in the flow when subsequent operators reported that sorted input was required for correct functionality.

In RETL 11.2.2, the sort order and sort keys are copied from the first INPUT to LOOKUP.

Bug #916 – Invalid number specification ignored in BINOP, FILTER, PARSER, TERAFERead, SWITCH, JOIN, and CONVERT

In RETL versions 11.0 through 11.2.1, conversion of text in flow properties to RETL data types was not strictly checked. For example, if a BINOP was used to add a constant value to an int8 field and the constant was specified as “1.0”, the error converting “1.0” to an int8 would be lost and the constant would actually set to 0.

In RETL 11.2.2, invalid conversions now produce an error, such as

```
E346: Error converting "1.0" to RETL type "int8"
```

Bug #938 – Incorrect delimiter used to handle null last field in ORAWRITE

In RETL versions 11.0 through 11.2.1, if ORAWRITE’s “outputdelimiter” property was specified and the last field was null, the field’s delimiter (‘|’ by default) was used instead of the specified delimiter. The incorrect delimiter was written to the database as the value of the field instead of null.

In RETL 11.2.2, the specified delimiter is used.

Bug #1081 – Backwards compatibility bug disallowing “len” and “maxlength” attributes in schema

In RETL 11.2.1, an error was displayed if both the “len” and “maxlength” attributes were specified in a schema. This bug caused a backwards compatibility issue because some schemas specified both attributes with equal values.

In RETL 11.2.2, “len” and “maxlength” can both be specified, but the values must be the same. If the values differ, the following error is displayed:

```
E328: If both "len" and "maxlength" are specified then they must be equal.
```

Bug #1105 – FIELDMOD duplicate does not check for pre-existing field

In RETL versions 11.0 through 11.2.1, the duplicate property did not check to see whether the new field's name matched the name of an existing field. Records would contain two fields with the same name. This bug would cause inconsistent behavior in any following operators that referred to the new or pre-existing field.

In RETL 11.2.2, the following error message is displayed:

```
E345: Field "destination" already exists while trying to duplicate field with "source =destination"
```

Bug #1111 – PARSER does not support floating-point constants

In RETL 11.2 and 11.2.1, the PARSER operator did not support floating-point constants. This capability was added to RETL 11.2.2.

Bug #1163 – GENERATOR cannot start a partition

In RETL 11.2 and RETL 11.2.1, the GENERATOR operator could not be partitioned when creating records. (GENERATOR could be partitioned when generating fields for pre-existing records.)

This functionality has been added to RETL 11.2.2.

Bug #1175 – SORT rejects import of dfloat or sfloat read from fixed-width file

In RETL versions 11.0 through 11.2.1, if a dfloat or an sfloat were read from a file with a fixed width schema and the width was not large enough to handle the 6 digits of precision that RETL uses by default, the record would be rejected when being re-imported from disk within the SORT operator.

For example, if a dfloat with a value of 1200.0 was read from a fixed width file and the width for the field was 6, the re-import by SORT would cause the record to be rejected because the field was written out as 1200.000000, which has a length longer than 6.

In RETL 11.2.2, dfloats and sfloats are written to disk respecting the maximum length specified in the import file.

Bug #1177 – FIELDMOD does not perform deep enough copy when duplicating a field

In RETL versions 11.0 through 11.2.1, a reference to a field was copied when the duplicate property was specified. Any change to one field would also show up in the other field.

In RETL 11.2.2, a new field is constructed and its value set to the same as the source field.

Bug #1183 – Incorrect error message when FILTER’s filter property syntax is incorrect

In RETL versions 11.0 through 11.2.1, RETL displayed the following error message when the syntax of the filter property of FILTER was incorrect:

No filter logic is specified.

In RETL 11.2.2, if the filter property syntax is incorrect, the following error is displayed:

```
E349: Invalid syntax in "filter" property: filter-value
```

Bug #1211 – 32bit JVM is not properly installed on AIX

In RETL 11.2.1, the tarball for the 32-bit JVM for AIX was corrupt. This bug has been fixed in RETL 11.2.2.

Bug #1222 – NullPointerException in TERAWRITE

In RETL 11.2 and 11.2.1, if TERAWRITE was used to write data to a Teradata table that had more columns than the dataset and in a different order, it is possible for TERAWRITE to throw a NullPointerException because the dataset fields were not correctly mapped to the table columns.

This bug has been corrected in RETL 11.2.2.

Bug #1226 – Incorrect handling of DECIMAL fields in TERAFOREAD

In RETL 11.2 and 11.2.1, TERAFOREAD incorrectly read DECIMAL fields from the database.

This bug has been corrected in RETL 11.2.2.

Bug #1237 – RQST00000061454: orawrite strips trailing spaces breaking changecapture

In RETL versions 11.0 through 11.2.1, spaces at the end of records are not preserved in ORAWRITE because the default behavior of SQL*Loader is to trim the trailing spaces. This bug can cause problems when the data is later retrieved and compared against another dataset using CHANGECAPTURE because the trimmed trailing space causes the record to be flagged as a change.

RETL 11.2.2’s ORAWRITE operator accepts a new property “controlfileloadoptions” that will be written to the SQL*Loader control file. Setting “controlfileloadoptions” to “PRESERVE BLANKS” will preserve trailing spaces. This action must be coupled with a delimited load (specified by setting the “outputdelimiter” property) in order to work well with changecapture.

Changes since RETL 11.2

See above for enhancements in the 11.2.1 release.

- Fixes Bug#434 – Allow logging to log actual times spent per operator.
- Fixes Bug#544 – Missing CONVERTFUNCTION results in no-op instead of error.
- Fixes Bug#860 – 4 RETL instances much quicker than 4 ORAREADS.
- Fixes Bug#905 – Progress indicators.
- Fixes Bug#906 – Performance statistics show operator initialization time under framework.
- Fixes Bug#934 – dblookup does not partition.
- Fixes Bug#936 – Partitioning error when funneling multiple partition groups with different numpartitions.
- Fixes Bug#942 – Incorrect schema resolution in generator after inline funnel.
- Fixes Bug#945 – Negative dfloats can break fixed-length output schema.
- Fixes Bug#946 – Occasional NullPointerException in funnel.
- Fixes Bug#953 – Join key datatypes must match.
- Fixes Bug#955 – Enhanced logging
- Fixes Bug#957 – Assertions should be enabled when RFX_DEBUG=1.
- Fixes Bug#961 – Partitioned lookup is incorrect when lookup table is non-keyed partitioned.
- Fixes Bug#963 – Mismatched numpartition settings cause assertion error.
- Fixes Bug#964 – Date comparison failure in parser.
- Fixes Bug#966 – Should be able to specify multiple outputs in parser.
- Fixes Bug#967 – IndexOutOfBoundsException in orawrite.
- Fixes Bug#969 – Groupby errors on null value for min/max.
- Fixes Bug#970 – Double formatting with magnitude < 0.
- Fixes Bug#971 – first field nulls not handled properly in comparison and assignments.
- Fixes Bug#972 – Enhance graphing with performance measurements.
- Fixes Bug#973 – Verify graphing performance measurement accuracy.
- Fixes Bug#974 – Add and adjust dataset monitoring capabilities.
- Fixes Bug#975 – Improve the tests surrounding the new parser functionality.
- Fixes Bug#981 – Partitioned lookup may not scale.
- Fixes Bug#982 – Delimited orawrite should specify length for CHAR.
- Fixes Bug#984 – DEFAULTS in rfx.conf do not override operator properties.
- Fixes Bug#985 – NullPointerException from StreamRedirector.
- Fixes Bug#986 – Comparison between different field types gives incorrect error.

- Fixes Bug#990 – dblookup sets incorrect SQL to look up in database.
- Fixes Bug#994 – -g must be last option on command line.
- Fixes Bug#995 – Convert can invalidate sort order on presorted data.

Changes since RETL 11.1

- Fixes Bug#407 – In framework we treat all int and uint type as long in Java version.
- Fixes Bug#421 – Rejected records get printed to the screen which can produce too much output.
- Fixes Bug#441 – Support for properly-formatted precision in dfloat/sfloat.
- Fixes Bug#445 – Parallel file reads.
- Fixes Bug#461 – Include specific error in import rejects.
- Fixes Bug#464 – Determine if we really need hostname/port for oraread/write.
- Fixes Bug#484 – Need build information option.
- Fixes Bug#488 – Better exception handling in Operator base class.
- Fixes Bug#493 – Record key hash non-unique with multiple key fields.
- Fixes Bug#499 – Import does not use nullvalue property to detect null values.
- Fixes Bug#503 – Fixed width of zero not allowed.
- Fixes Bug#507 – Schema.parseKeyFields(...) should be more robust.
- Fixes Bug#508 – Number format is not standardized on export.
- Fixes Bug#510 – userid/password properties do not exist for db2read/teraferead operators.
- Fixes Bug#515 – jdbcwrite should set proper type in constructor.
- Fixes Bug#516 – Implement native load mechanisms for orawrite in 11.x versions.
- Fixes Bug#517 – Implement native load mechanisms for terawrite in 11.x versions.
- Fixes Bug#522 – Generator does not implement partnum_offset or partcount_incr.
- Fixes Bug#532 – final_delimiter not implemented.
- Fixes Bug#534 – Missing len from fixed schema file causes delimited export.
- Fixes Bug#541 – Incorrect handling of too few fields in import.
- Fixes Bug#543 – Multiple use of same dataset as input throws null pointer exception.
- Fixes Bug#547 – JDBCWrite should check schema owner.
- Fixes Bug#548 – Port is not configurable for Teradata (terawrite).
- Fixes Bug#553 – Checks if there is sqlldr.
- Fixes Bug#555 – Strip schema attribute not handled.
- Fixes Bug#557 – fieldmod does not check if fieldnames actually exist in incoming dataset.
- Fixes Bug#558 – createtablemode should only allow recreate/create.

- Fixes Bug#561 – RETL does not abort gracefully if dbwrite password not specified.
- Fixes Bug#564 – Should modify maxsize if database column maxsize < field maxsize.
- Fixes Bug#572 – Merge jdbcread methods with DatabaseUtils and ColumnInfo.
- Fixes Bug#574 – logical error in changecapture.
- Fixes Bug#579 – Fatal error trapping when using SQL-loader.
- Fixes Bug#606 – Handling environment problems with RETL executable/library.
- Fixes Bug#610 – In the fixed import, we should check if the total len of the record equals the sum of the len of fields.
- Fixes Bug#632 – Incorrect handling of too many fields in import.
- Fixes Bug#635 – Generator does not create null values same as 10.x.
- Fixes Bug#641 – Generator sequence does not respect destination field datatype.
- Fixes Bug#643 – Convert cannot handle duplicate source fields in one CONVERTSPEC.
- Fixes Bug#647 – Reject file incorrect format.
- Fixes Bug#648 – Incorrect contents of reject file.
- Fixes Bug#675 – Sort of generated string fails on re-import.
- Fixes Bug#678 – Multibyte fixed width export incorrect/slow.
- Fixes Bug#685 – Funnel should not be waiting when there is no data in any one of inputs.
- Fixes Bug#686 – Dataset should check with deadlock listener always even though the dataset has been expanded.
- Fixes Bug#694 – Non-UTF encodings cause exception in import.
- Fixes Bug#696 – 11.x can break the output schemafilename definition.
- Fixes Bug#698 – sortfunnel outputs wrong data – logical problem.
- Fixes Bug#700 – Generator sequence does not work with doubles.
- Fixes Bug#701 – Binop incorrect when destination field is a source field.
- Fixes Bug#703 – Improved performance when handling errors.
- Fixes Bug#705 – Performance bottlenecks in orawrite.getRecordStringBuffer(...).
- Fixes Bug#708 – Performance improvement in jdbcread.
- Fixes Bug#709 – Performance degradation in jdbcread due to string comparison.
- Fixes Bug#710 – Sort does not handle non-existent 'key' gracefully.
- Fixes Bug#713 – changecapture codes incorrect in documentation.
- Fixes Bug#714 – Error message bad when null value is in non-null field.
- Fixes Bug#716 – Delimiter property of sort not working correctly.
- Fixes Bug#718 – NullPointerException in lookup.
- Fixes Bug#719 – rightouterjoin invalid schema.

Retek Extract Transform and Load

- Fixes Bug#721 – Import of multi-byte fixed-width files fails.
- Fixes Bug#722 – ConvertByteToBytes incorrect.
- Fixes Bug#723 – LONGVARCHAR only partially supported.
- Fixes Bug#724 – Cannot handle "select NULL something from test_jdbc".
- Fixes Bug#725 – Generator backwards compatibility issue with field length.
- Fixes Bug#726 – ArrayIndexOutOfBoundsException when missing field type in schema.
- Fixes Bug#727 – NullPointerException when numloader > 1 and RFX_SHOW_SQL=1.
- Fixes Bug#728 – I18N and fixed width incorrect.
- Fixes Bug#729 – The dataset hang when generator-->noop with independent thread with very small buffersize and huge data.
- Fixes Bug#734 – NullPointerException when paging.
- Fixes Bug#736 – ArrayOutOfBounds in export.
- Fixes Bug#737 – Rejects file not flushed when exiting due to too many rejects.
- Fixes Bug#738 – Parallel orawrite does not alternate TEMPDIRs.
- Fixes Bug#739 – tempdir property should override setting in rfx.conf.
- Fixes Bug#740 – intra-operator parallelism properties should have standard names.
- Fixes Bug#743 – Temporary directories should be checked.
- Fixes Bug#745 – Schema assembler should ensure field delimiter != record delimiter.
- Fixes Bug#746 – orawrite should ensure output delimiter != record delimiter.
- Fixes Bug#747 – orawrite cannot handle preload and postload for stored procedure.
- Fixes Bug#748 – Convolutd error output when setting properties.
- Fixes Bug#749 – Issue warning instead of error for unrecognized properties.
- Fixes Bug#751 – Combine DataSet, SDataSet, DataSet_Simple, and DataSet_RawBytes.
- Fixes Bug#752 – Add --graphviz-files
- Fixes Bug#753 – sp_prequery should be executed BEFORE the query.
- Fixes Bug#754 – Heterogeneous output schemas.
- Fixes Bug#756 – Add 64bit JVM support in RETL 11.
- Fixes Bug#757 – Convert does not handle TYPEPROPERTY properly.
- Fixes Bug#758 – Int32/Int64 Field type rollover.
- Fixes Bug#760 – Binop odd division rounding error.
- Fixes Bug#761 – Lookup returns false matches.
- Fixes Bug#762 – Unsigned integer issues.
- Fixes Bug#764 – NumberUtils.ConvertIntToBytes poor algorithm.

- Fixes Bug#769 – Sort failure on date without century.
- Fixes Bug#771 – Sort is not removing temporary files.
- Fixes Bug#772 – RETL 11.x does not support partitioning.
- Fixes Bug#774 – Lookup ArrayIndexOutOfBoundsException with multiple lookup tables.
- Fixes Bug#776 – Out of memory exception should abort flow.
- Fixes Bug#779 – Fieldmod loses SortKeys.
- Fixes Bug#780 – Schema assembler sets sort keys incorrectly.
- Fixes Bug#782 – Default value for properties not set for join.
- Fixes Bug#784 – graphviz cleanup.
- Fixes Bug#786 – Should error when record len != sum(field lens) in schemafile.
- Fixes Bug#788 – Forwards compatibility for operator properties.
- Fixes Bug#790 – Join with database operator.
- Fixes Bug#791 – Groupby requires reduce property before Min, Max, ... Fields.
- Fixes Bug#792 – Verify retl.jar during RETL startup.
- Fixes Bug#793 – Verify retl.jar during RETL startup.
- Fixes Bug#796 – Enhanced performance logging.
- Fixes Bug#799 – Deadlock detection not working for pipelines with more than two outputs.
- Fixes Bug#800 – Occasional errors when paging.
- Fixes Bug#801 – Paging may result in too many files open.
- Fixes Bug#802 – DataSet paging files may not be created.
- Fixes Bug#803 – Unneeded check for duplicate tables in existence.
- Fixes Bug#811 – leftouterjoin does not support spaces in nullvalue property.
- Fixes Bug#812 – fieldmod blows away sorted order.
- Fixes Bug#813 – rightouterjoin invalid sort order.
- Fixes Bug#816 – Record.Join uses incorrect calculation to get # of fields.
- Fixes Bug#818 – Poor message displayed when output record breaks schema.
- Fixes Bug#821 – Delimited import strips trailing whitespace in 11 but not in 10.
- Fixes Bug#822 – Package debug/non-optimized version of retl.jar
- Fixes Bug#823 – The RETL Programmer's Guide description of ifnotfound property is incorrect.
- Fixes Bug#824 – Improve logging information with flow name.
- Fixes Bug#825 – Paths in RETL script create problems on Windows.
- Fixes Bug#826 – Add 'loader_options' parameter to orawrite.

Retek Extract Transform and Load

- Fixes Bug#829 – Maxlength of 0 creates problems for orawrite.
- Fixes Bug#830 – slsildmdm runs faster in 10.x than in 11.x.
- Fixes Bug#832 – Incorrect null in Fullouterjoin.
- Fixes Bug#833 – Error in double to byte.
- Fixes Bug#834 – Dtty files removed with -g option.
- Fixes Bug#835 – Ned specific error message in oraread query.
- Fixes Bug#836 – Missed record in multi-threaded import.
- Fixes Bug#838 – Stack overflow in generator if limit < init.
- Fixes Bug#839 – Copy does not perform deep enough copy.
- Fixes Bug#841 – dfloat fields can break fixed output schema.
- Fixes Bug#842 – Lookup table loading is dependant on data record production.
- Fixes Bug#843 – Assertion is fired in allocation size_profile script.
- Fixes Bug#845 – Assertion in Parser_ByteSequence.readFixedField in Allocation size_profile.
- Fixes Bug#847 – Pipeline initialization and schema resolution is serial.
- Fixes Bug#848 – ArrayIndexOutOfBoundsException when 1 INPUT provided to lookup.
- Fixes Bug#852 – Operator jdbcwrite needs parallel property.
- Fixes Bug#853 – Check for WHERE in dblookup select_sql property incorrect.
- Fixes Bug#854 – Partitioned Orawrite runs preload and postload for each partition.
- Fixes Bug#855 – db2write deletes log files even after a fatal error.
- Fixes Bug#859 – Monitor moethod of funnel does not work for empty datasets.
- Fixes Bug#860 – 4 RETL instances much quicker than 4 Orareads.
- Fixes Bug#862 – Adding batched funnel algorithms causes dataset._pblock to be null.
- Fixes Bug#864 – Should print a warning if orawrite cannot be partitioned.
- Fixes Bug#867 – verboserejects should be a visible property in import.
- Fixes Bug#868 – RETL does not print any warning if rejects + verboserejects=false.
- Fixes Bug#869 – Errors not caught in DatabaseUtils.Connect().
- Fixes Bug#871 – ImportRejectSupporter throws an error on processing rejected records.
- Fixes Bug#872 – Multiple-query JDBCRead is slower than multiple JDBCReads.
- Fixes Bug#873 – jdbcread needs to partition.
- Fixes Bug#874 – Import partitioning does not calculate number of partitions correctly.
- Fixes Bug#875 – Import partitioning does not need to end thread.
- Fixes Bug#877 – verify_retl chokes on space between -d and database type.

- Fixes Bug#881 – Schema compatibility checking in export is too strict.
- Fixes Bug#883 – Partitioning fails in aipt_future_delivery.
- Fixes Bug#887 – Execute jdbcread subthreads in parallel.
- Fixes Bug#888 – Dataset-less funnel.
- Fixes Bug#889 – Check for retl.jar in RETL script does not check RETL_DIST_LIB variable.
- Fixes Bug#890 – Funnel should have 11.1 algorithm in 11.2.
- Fixes Bug#891 – Generated HTML for graphs contains incorrect message about dotted file.
- Fixes Bug#892 – Funnel should drop sort order.
- Fixes Bug#895 – Add create_options property to orawrite.
- Fixes Bug#901 – Include 1.4.1 AIX JRE in package rather than 1.4.0.
- Fixes Bug#903 – Inlining lookup may improve performance.
- Fixes Bug#904 – Splitter needs to end thread when partitioning.
- Fixes Bug#907 – Non-deterministic output if debugging is set on.
- Fixes Bug#910 – Incorrect schema resolution for fullouterjoin.
- Fixes Bug#911 – Binop produces incorrect field type when multiplying int by dfloat.
- Fixes Bug#912 – NullPointerException in binop when left or constleft not specified.
- Fixes Bug#913 – Conflicting versions of rfx.conf num partitions and -n option do not seem to resolve properly.
- Fixes Bug#919 – Import partitioning does not work with multiple files.
- Fixes Bug#920 – Poor error message when nullable missing in import schemafile.
- Fixes Bug#921 – Hash partitioning after splitter uses incorrect dimensions.
- Fixes Bug#922 – Import - multi-threaded imports are losing records.
- Fixes Bug#923 – Export should partition.
- Fixes Bug#924 – FileUtils concatenation() errors when output file does not exist yet.
- Fixes Bug#925 – Misleading error message "record was too long" in import.
- Fixes Bug#927 – Changelog algorithm does not allow for partitioned imports.
- Fixes Bug#929 – install_jre fails on some servers.
- Fixes Bug#930 – Cannot override default logger file from config file.
- Fixes Bug#931 – Log file should be in temp directory if directory not specified.
- Fixes Bug#932 – Unable to write to log file.
- Fixes Bug#935 – changecapturelookup interferes with lookup's tablekeys property.