

---

# Retek<sup>®</sup> Extract Transform and Load<sup>™</sup> 11.1.2

## 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.

Retek<sup>®</sup> Extract Transform and Load<sup>™</sup> 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This unpublished work is protected by confidentiality agreement, and by trade secret, copyright, and other laws.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:

©2004 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](http://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
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.

---

# Release Notes

The RETL 11.1.2 patch is a refresh of the fixes in RETL 11.1.1, and it fixes two additional issues.

The first addresses a regression bug where floating point numbers were prematurely rolling into scientific notation, which caused a critical loss of precision in significant digits.

The second fix addresses an algorithm problem in funnel. There is now a new algorithm in funnel that will significantly speed up performance in flows that printed warnings about paging to disk.

The third fix addresses a problem with integer overflows. Incorrect comparisons as a result of the integer overflow caused outerjoin to incorrectly insert null values for fields that shouldn't have had them.

The issues fixed in 11.1.2 rolled in from 11.1.1 are as follows:

- The first addresses an ominous problem in the sfloat/dfloat datatypes that can cause imprecision in insignificant decimal digits and can cause RETL to break output schemas. Because there is still the possibility of output schemas being broken, application groups should always use a 'len' or 'maxlength' of 25 as the length for fixed-length dfloat fields. sfloat/dfloat values must follow the format of 25 digits maximum (including positive/negative signs, decimal point, and all digits of scientific notation, e.g. -134.2342550425, 1.98847e34 are valid while 1.23456789012345678901234567890 is not).
- The second fix addresses a logical error in fullouterjoin, where a user might see an `ArrayIndexOutOfBoundsException` exception in certain flows that use fullouterjoin.

It is recommended that all products that use RETL 11 upgrade to this latest patch version. See below changes since RETL 11.1 for more information on this bug. Please refer to the release notes of RETL 11.1 for additional information about RETL 11 product certifications.



**Note:** It is strongly recommended that before any installation or development work is performed, that the RETL 11.1 Programmer's Guide, provided with the release, is read cover to cover. In particular, read and understand Chapter 2 and the 'Upgrade From A Previous Versions Of RETL'/'Backwards Compatibility Notes' section, and the 'Known Issues' section of the Release Notes. Additionally, product groups should take care to extensively complete all certifications and regression/volume testing with RETL 11 versions before deployment to customer sites.

# Compatibility Matrix

The following represents OS/Database combinations that are supported and have been certified to work properly :

OS Version	Arch	DB	Version
AIX 5.1	64	Oracle	9.2
HP-UX 11i	64	None	None
HP-UX 11i	32	None	None
HP-UX 11i	64	Oracle	9.2
HP-UX 11i	64	Oracle	9.0.1
HP-UX 11i	32	Oracle	8.1.7
HP-UX 11i	32	Tera	2r4.2
Solaris 8	32	None	None
Solaris 8	64	None	None
Solaris 8	64	Oracle	9.2
Solaris 8	64	Oracle	9.0.1
Solaris 8	32	Oracle	9.0.1
Solaris 8	32	Oracle	8.1.7
Solaris 8	32	Tera	2r4.2



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

# Installation

You may download from "<http://insideretek/development/retl.asp>". This release package contains everything needed for all platforms/databases.

You can check the contents of the package and verify that it is valid by executing the following command:

```
gunzip -c <package.tar.gz> | tar -tf -
```

To extract the package use the following command:

```
gunzip -c <package.tar.gz> | tar -xf -
```

This will create the directory and the package contents (see below).

## Package Contents

The following is a brief description of the contents of retl.11.1.2.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
rff	Symbolic link to retl provided for backwards compatibility
verify_retl	The RETL verification tool
README.verify_retl	The README for verify_retl



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

## Retek Extract Transform and Load

---

lib/

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

File	Description
retl.jar	RETL runtime library.
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

docs/

This directory contains the documentation associated with RETL.

File	Description
ReleaseNotes.txt	This file.

etc/

This directory contains the documentation associated with RETL.

File	Description
rfl.conf	The default retl 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\_ret

---

# Known Issues

## General Known Issues

- The RETL does not handle arbitrary precision math.
- In the RETL 11.1 or prior releases, carefully check all property names to ensure they exactly match the property name and value spellings noted in the Programmer's Guide. RETL may error on misspelled or unknown properties.
- The export operator will default to use '|' delimited fields if there isn't a 'schemafilename' property specified. This 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.
- The orawrite operator will cause RETL to abend with a return code != 0 only in the case that sql\*loader fails. Should sql\*loader exit prematurely after having reached its maximum number of allowed rejected records, RETL will not abort. Adding this functionality is an enhancement and logged as #579
- 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 10.x and 11.x versions of RETL

- RETL 11 versions are more strict 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.
- An unintended feature of 10.x allowed input datafiles to contain more fields than are specified in the schemafilename. RETL 11 does not allow more fields to exist in the datafile and will throw an exception should this happen. This 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 don't specify a 'schemafilename' 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 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 is because 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 programmer's guide for more info.
- RETL 11.x versions may require more physical memory for certain flows than 10.x releases.
- DB2 has not yet been certified in RETL 11. This certification will take place in the next release of RETL.

- In RETL 11.1.2 and beyond, floating-point numbers in the dfloat datatype will roll over into scientific notation on a number bigger than 10000000000000000 (1E16). 10.x versions never rolled into scientific notation, but this itself could cause imprecision in insignificant digits

### Backwards-Compatibility Issues with 10.x versions

- It is not recommended to upgrade to RETL 11 unless the product using RETL has been certified on RETL 11.

### Known bugs in RETL 11.1.2

- Bug#461 - Should give more specific error messages with source of the error pointing all the way down to the record/field. This will be corrected in a near future version of RETL.
- Bug#407 - RETL 11.1 versions will not reject numeric fields based on maximum values for each particular RETL datatype. For example, the value '100' will not be rejected for an 'int8' datatype. This is a bug that will be fixed in the next release of RETL.
- Bug#632 - Incorrect Handling of Too Many Fields in Import. An unintended feature of 10.x allowed input datafiles to contain more fields than are specified in the schemafilename. RETL 11 does not allow more fields to exist in the datafile and will throw an exception should this happen. This will be fixed in a future version of RETL 11.
- Bug#541 - Incorrect Handling of Too Few Fields in Import. RETL 10.x will reject incoming records with too few fields. 11.x versions will treat the missing fields as null values. This will be fixed to reject records in future versions of RETL.
- Bug#652 - NOT Can Only be Used Once in Filter Expression. The workaround is to reverse all other operations to obtain the same effect. For example, "not STATUS\_CDE eq 'A' and not STATUS\_CDE eq 'C' and not STATUS\_CDE eq 'X' becomes "STATUS\_CDE ne 'A' and STATUS\_CDE ne 'C' and STATUS\_CDE ne 'X'".
- Bug#678 - Multibyte fixed-width exports are slower than single-byte fixed-width exports. This will be corrected in the next version of RETL.
- Bug#646 - Terawrite performs a 'delete' on a table when terawrite's 'mode' property == 'truncate'. This results in slow truncation on large tables. This will be fixed in the next version of RETL.

---

## Changes since RETL 11.1.1

- Fixes regression on Bug#441 - RETL 11.1.1 rolled over into scientific notation if a number > 10 million was specified. This caused problems when doing a UNION ALL in oracle as this will remove scale and precision in integers, effectively transforming them into floating-point numbers that would be subject to premature conversion to scientific notation. This is now fixed as RETL 11.1.2 won't roll over into scientific notation until the dfloat exceeds 10 quadrillion (10000 trillion).
- Fixes Bug#685 - funnel shouldn't be waiting when there is not data in any one of its inputs. Funnel will now yield and switch inputs in a better round-robin fashion instead of waiting on one dataset to finish before moving to the next. This will stop RETL from paging to disk and effectively improves performance for flows that print warnings about disk paging.
- Fixes Bug#832 - An invalid comparison was being performed in int64 fields. The difference between the two values was being cast to an int and returned. If the difference between the two values was larger than the maximum integer, then the returned int would overflow to a negative value. With this fix, RETL will now correctly cast to a long and use that as a comparison.

## Changes since RETL 11.1

- Fixes Bug#441 - RETL 11.1 printed all known digits in floating-point numbers (sfloat and dfloat) even though the last digit could be inaccurate, as is the case for floating-point fields in general. This poses a problem when doing calculations and expecting certain results. Additionally, printing all digits would at times break output schemas. The fix involves only printing 6 digits past the decimal place for floating-point fields, which is consistent with how 10.x versions of RETL work. The exception to this is if an input datafile contains more precision and there aren't any operations on the sfloat/dfloat fields, then the output schemafilename will contain the same precision as the input data.
- Fixes Bug#816 - A bug existed in join code that caused some flows that use fullouterjoin to throw an ArrayIndexOutOfBoundsException. The problem was that the join code was using an incorrect calculation that, at times, would cause the record size to be incorrect. This bug is now fixed and the record is sized correctly.