
Retek[®] Extract Transform and Load[™] 11.1.1

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[®] 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This unpublished work is protected by confidentiality agreement, and by trade secret, copyright, and other laws.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:

©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.retек.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.

RETL 11.1.1 Overview

The RETL 11.1.1 release addresses two critical issues.

- The first addresses an issue 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 issue addresses a logical error in fullouterjoin, where a user might see an ArrayIndexOutOfBounds 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.

RETL 11.1.1 Issues Resolved



We strongly recommend 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" and "Backwards Compatibility Notes" sections as well as the "Known Issues" section of these Release Notes. All certified Retek products using RETL must be regression and volume tested with the RETL 11.1 tool prior to its production deployment at customer sites.

Issues identified in the 11.1 noted below have been resolved in the RETL 11.1.1 patch release:

Issue #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 schemafile 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.

RETL 11.1.1 Compatibility

Following is a listing of RETL 11.1.1 operating system and database compatibilities.

OS	Version	Arch	DB	Version	RETL Executable
IBM AIX	64 bit	5.1	Oracle	9.2.0.2	IBM AIX
HP-UX	64 bit	11i	Oracle	9.2.0.2	HP-UX
HP-UX	64 bit	11i	Oracle	9.0.1.4	HP-UX
HP-UX	32 bit	11i	Oracle	8.1.7.4	HP-UX
HP-UX	32 bit	11i	NCR Teradata	2r4.2	HP-UX
HP-UX	64 bit	11i	None	None	HP-UX
HP-UX	32 bit	11i	None	None	HP-UX
Sun Solaris	64 bit	5.8	Oracle	9.2.0.2	Sun Solaris
Sun Solaris	64 bit	5.8	Oracle	9.0.1.4	Sun Solaris
Sun Solaris	32 bit	5.8	Oracle	9.0.1.4	Sun Solaris
Sun Solaris	32 bit	5.8	Oracle	8.1.7.4	Sun Solaris
Sun Solaris	32 bit	5.8	NCR Teradata	2r4.2	Sun Solaris
Sun Solaris	64 bit	5.8	None	None	Sun Solaris
Sun Solaris	32 bit	5.8	None	None	Sun Solaris



Notes:

- RETL 11.x only requires one binary and installation for all database / platform combination listed above.
- Certification with DB2 v7.2.3 was not performed as part of the RETL certification effort. Currently, no Retek clients are using RETL with DB2 so the determination was made to re-schedule the DB2 certification effort of DB2 v8 for the next generally available release of RETL.
- Solaris 2.8, Solaris 8 and SunOS 5.8 are the same platform.
- If your current operating system or database is not specified above, please contact Retek Customer Support. The RETL is able to process datasets either directly from RDBMS database tables or from Unix flat file based systems. The RETL stores information about each dataset that describes the data structures (the metadata). The means by which the metadata is supplied is dependent on the data interface operators provided by the RETL and may require use of a schema to ensure that datasets are consistent from the source data to the target data.

Backwards-Compatibility to 10.x versions

- RETL 11 versions are stricter when enforcing 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 schemafile. 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 'schemafile' property in the export operator.
- RETL 10.x versions always format float fields with 6 digits of precision. In 11.x versions, the precision is the same as the input data to RETL. This hasn't been shown to cause any issues in certification efforts, but it should be noted as a behavioral difference. This is documented as Bug #441.
- 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 11.1 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 RETL 11.1 Programmer's Guide for more information.
- RETL 11.x versions may require more physical memory for certain flows than 10.x releases.
- Certification with DB2 was not performed as part of the RETL 11.1 certification effort. Currently, no Retek clients are using RETL with DB2 so the determination was made to re-schedule the DB2 certification effort of DB2 v8 for the next generally available release of RETL.

DB2 Certification

Certification with DB2 was not performed as part of the RETL 11.1 certification effort. Currently, no Retek clients are using RETL with DB2 so the determination was made to re-schedule the DB2 certification effort of DB2 v8 for the next generally available release of RETL.

General Issues of RETL (Release 11.1 and Prior)

- 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.

Known Issues Specific to RETL 11.1

- Bug #407 - RETL 11.1 does 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.
- Bug #461 – RETL 11.1 should provide a more specific error messages with more record and field level details.
- 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 #632 - An unintended feature of 10.x allowed input datafiles to contain more fields than are specified in the schemafilename. RETL 11.1 does not allow more fields to exist in the datafile and will throw an exception should this happen.
- Bug #646 - Terawrite performs a 'delete' on a table when terawrite's 'mode' property == 'truncate'. This results in slow truncation on large tables.
- 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.