



CONNIX 13.0

Release Notes

May 26, 2017

Table of Contents

CONNX 13.0 Release Notes.....	4
Overview.....	4
CONNX Architecture on Windows.....	5
CONNX Architecture on UNIX.....	6
CONNX Architecture (ODBC/JDBC/OLEDB/.NET Provider)	7
CONNX Client Engine for Windows	7
CONNX Client Engine for UNIX	7
CONNX Data Dictionary	7
CONNX Server.....	7
CONNX JDBC Thin Client.....	7
CONNX JDBC Server.....	7
CONNX JDBC Router	8
CONNX DataSync	8
CONNX DataSync Transformation Server.....	8
CONNX InstantdbSync (new in CONNX 12)	8
CONNX KPiSync (new in CONNX 12).....	9
CONNX Excel Add-in (new in CONNX 12)	9
CONNX Client PC Functions	11
CONNX Server Functions	11
Requirements for CONNX 13.0.....	12
Data (Host) Server Requirements.....	12
CLIENT PC REQUIREMENTS.....	15
Unix Client System Requirements	16
JDBC Pure Java Client Requirements.....	17
InstantdbSync Requirements	17
Obtaining a current JDK (Java Development Kit)*	18
Compatible Front Ends	18
Installation Instructions	20
64 bit Considerations	20
User Account Control (UAC).....	22
New Platform Support	23
Critical Changes.....	23
Changes/Bug Fixes for CONNX 13.0.....	24
Changes/Bug Fixes for CONNX 12.5.....	28
Changes/Bug Fixes for CONNX 12 SP3	34
Changes/Bug Fixes for CONNX 12 SP2	38
Changes/Bug Fixes for CONNX 12 SP1.....	42
Changes/Bug Fixes for CONNX 12.....	44
Upgrading from prior versions of CONNX/DataSync.....	51
Upgrading from prior versions of InstantdbSync/Open Systems Event Replicator .	51
CONNX .Net Data Provider - Connection Pooling and Pooled Connection Timeout..	51

CONNX 13.0 Release Notes

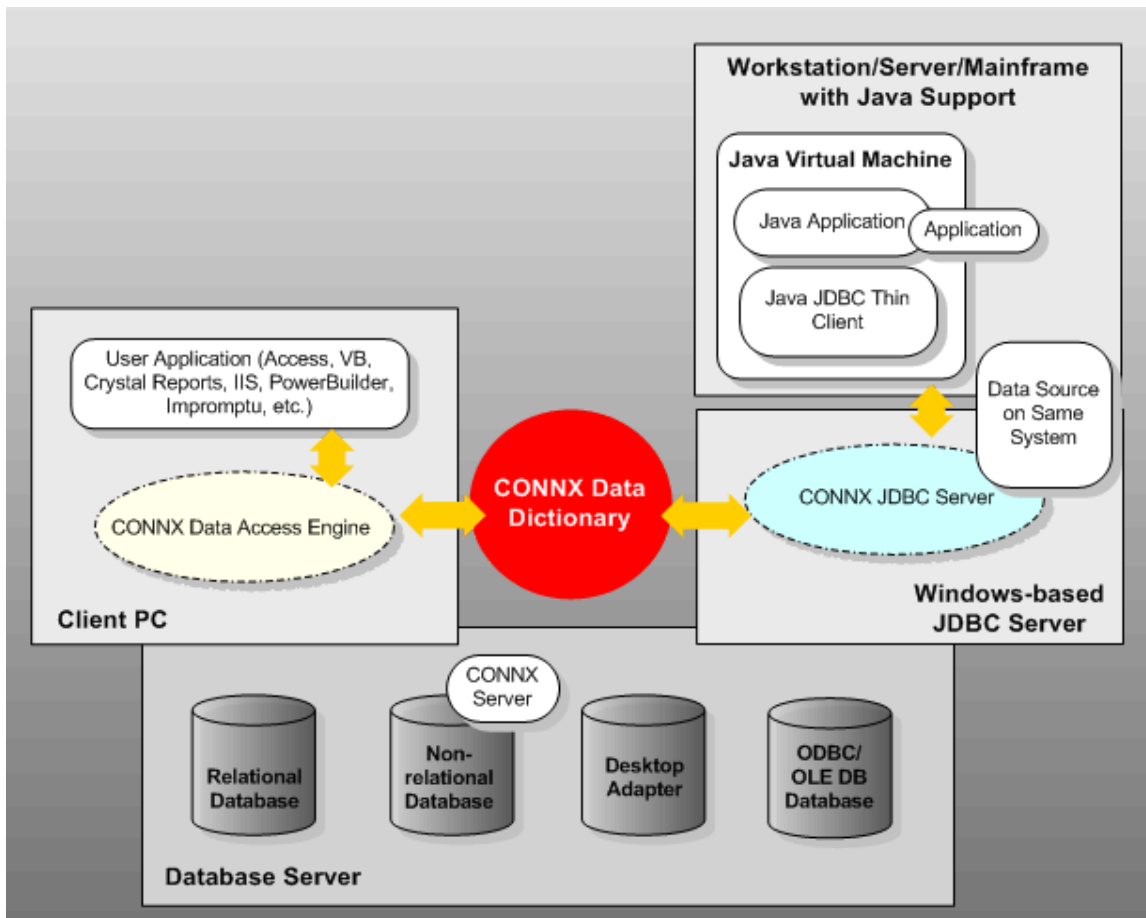
Overview

CONNX provides businesses secure read/write real-time access to all enterprise data from any platform as if all the data existed in one relational database. All data is then accessible using standard SQL and any standards-based application. CONNX acts as a reusable data access framework for projects throughout the enterprise. CONNX supports Adabas, C-ISAM, DB2, DISAM, MicroFocus, VSAM, IMS, Oracle, RMS, Rdb, PostGreSQL, DBMS, DataFlex, POWERflex, SQL Server, Sybase, Informix, and any OLE DB, ODBC, .NET, JDBC, UNIX, or Linux data source.

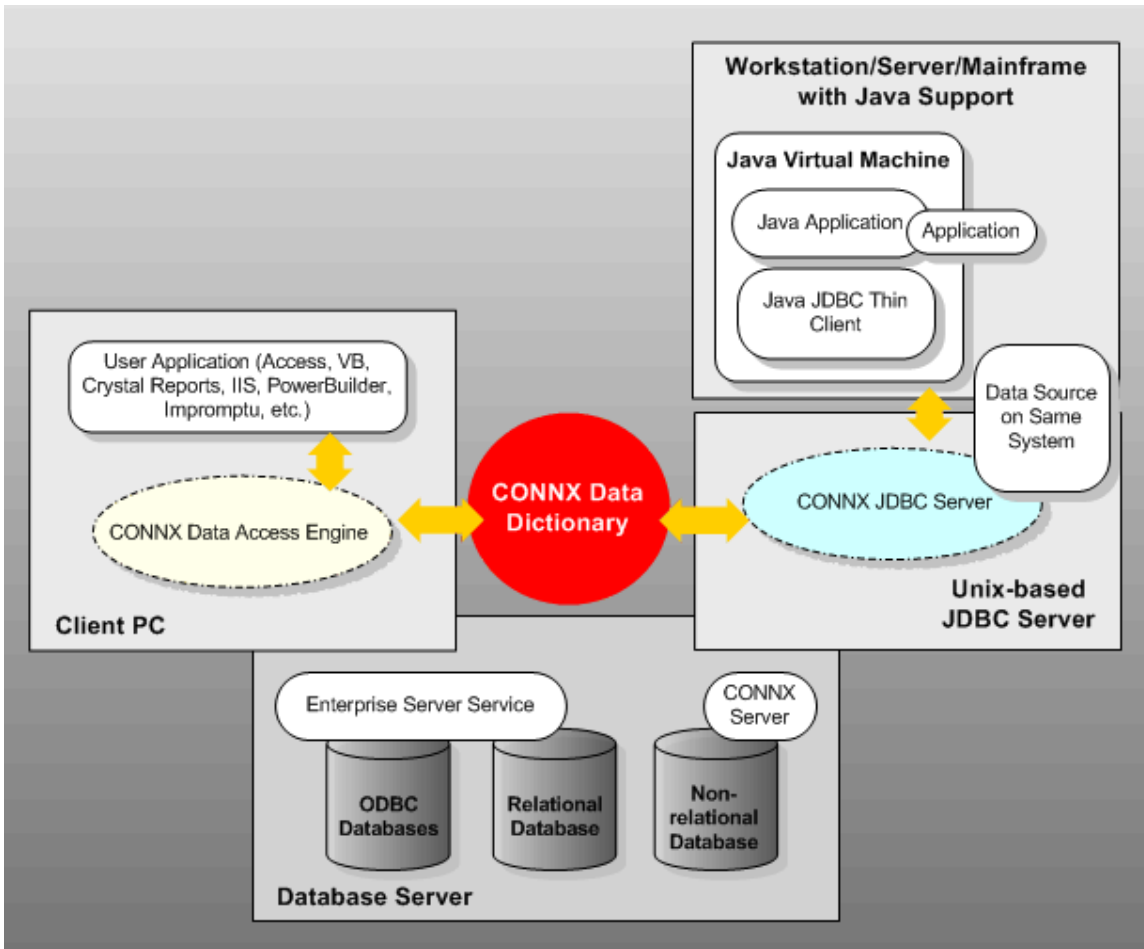
Here are just some of the implementations in which CONNX 13 can be utilized:

Data Migration	Web Development
Application/ Data Integration	Application Development
Ad Hoc Reporting	Data Warehousing

CONNX Architecture on Windows



CONNX Architecture on UNIX



CONNX Architecture (ODBC/JDBC/OLEDB/.NET Provider)

CONNX initially began as an ODBC Driver to RMS Data sources but since has evolved into a middleware product that has drivers/adapters for many different interfaces and data sources. The client is based on an ODBC driver, which is a dynamic link library that applications call to access data located in remote systems. The CONNX ODBC driver processes the ODBC function calls, submits requests to the appropriate data source, and then returns the results.

CONNX Client Engine for Windows

The CONNX Client Engine is based on an ODBC driver, which is a dynamic link library that applications call to access data located in remote systems. The CONNX ODBC driver processes the ODBC function calls, submits requests to the appropriate data source, and then returns the results.

CONNX Client Engine for UNIX

The CONNX UNIX Client is based on an ODBC driver, which is a shared library that applications call to access data located on remote systems.

CONNX Data Dictionary

The CONNX Data Dictionary (CDD) is a repository of information about the database tables and fields accessed through CONNX, including structure and security. It contains the metadata about the source information and provides a GUI screen for easy and intuitive maintenance of the metadata, stored procedures, security, and views.

CONNX Server

All CONNX servers are full-featured and translate SQL requests into native database requests. The CONNX ODBC driver makes the server transparent to the end user. Server functions for DataFlex, Oracle, and DB2 are resident on the client PC. The third-party driver determines the location of the server components for ODBC and OLE DB data sources.

CONNX JDBC Thin Client

The CONNX JDBC thin client allows read/write access to a CONNX data source from any client machine possessing a JDK (1.3+). JDKs exist for most platforms. The CONNX JDBC thin client is a Type 3 driver that processes the JDBC function calls and submits requests to the CONNX JDBC Server and then returns the results.

CONNX JDBC Server

The CONNX JDBC server handles requests from the CONNX JDBC thin client and accesses the target data sources. The CONNX JDBC server component is available as either a Microsoft Windows or UNIX server component that enables access to multiple heterogeneous data sources.

CONNX JDBC Router

The CONNX JDBC Router component is necessary only if Web applets are served by a non-Windows Web server. The router is a Java application placed on the non-Windows Web server. It is designed to route JDBC requests to the CONNX JDBC server.

The CONNX JDBC Router is required for installation if any of the following conditions apply:

- The Java Applet is served by a Web server that is hosted on a different machine than the JDBC Server that is being called by the applet; for example, when the JDBC Server is on machine PROD1 and the Web Server is on machine PROD2.
- The Java Applet is served by a Web Server that is running in a browser.

CONNX DataSync

The CONNX DataSync component is a stand-alone product available to customers who purchase a CONNX DataSync License. It offers the ability to synchronize tables and views from a source database to a target database. It includes a scheduler to specify at which times the synchronizations can take place. It also offers the ability to incrementally synchronize the target feed when changes take place on the source.

CONNX DataSync Transformation Server

The CONNX DataSync Transformation Server component is a standalone product available to customers who purchase a CONNX DataSync Transformation Server License. It extends the functionality of CONNX DataSync by offering the ability to create data transformations; for example, a source feed that includes multiple tables joined together, aggregate queries, or complex functions. The product offers many features that allow it to be classified as an ETL Tool.

CONNX InstantdbSync (new in CONNX 12)

CONNX InstantdbSync provides Real-time Data Replication/Synchronization, Real-time Change Data Capture and Real-time Data Warehousing.

In today's world of increasing data volumes and Big Data, we are seeing databases that are designed to store and process hundreds of millions of fields, handle large volumes of data, and manage high-speed transactions, all while they house the critical business data needs of today's fast-paced companies. With the databases of today, applications are able to save and retrieve complex data structures in a single physical record and in a single operation. Today's databases have significant transactional performance advantages over traditional databases, which would require several joins between many tables to accomplish the same task.

The InstantdbSync solution gives companies the ability to seamlessly move data from a source database to any number of relational, non-relational, Cloud, desktop, or in memory databases for easier access, data reporting, and analysis.

- Fast to Implement
- Easy to Use
- Scalable
- The Tool for Minimizing Project Risk

New for CONNX InstantdbSync 13.0

- Support for MySQL as a source database on Linux as well as Windows
- Support for ACD (Add, Change and Delete) targets on all supported source databases. ACD replications are a special type of replication that track the operations done against a database rather than replicating the data itself.

CONNX KPiSync (new in CONNX 12)

CONNX KPiSync is a Mobile Application to help users 'Visualize their Key Business Metrics on any device in real time'. The CONNX KPiSync solution was designed with the purpose of providing users with a quick and easy method of delivering Key Performance Indicators from all facets of a business to users via any device, smartphone tablet, or PC.

CONNX KPiSync keeps you in touch with your critical business information (KPi's - key performance indicators) from any corporate application or data store, any time, all the time, real-time, whether you are in the office, on the go, or at home. KPiSync makes it simple to keep track of your critical business metrics simply by defining a key performance indicator which will query your enterprise data. These indicators are then pushed to your desktop, laptop, tablet, or mobile devices, using our secure Microsoft Azure Cloud.

For CONNX Version 12.5 and above we've added a new granular licensing option. The new functionality allows the customer to create quotas for users and usage, so you can regulate the number of pushes, number of users, or number of KPi's by users.

CONNX Excel Add-in (new in CONNX 12)

The Excel Add-In allows for a quicker and easier data access to any data source within Microsoft Excel. The CONNX wizard will allow a streamlined connection and data selection for the over 100 supported data sources. CONNX extends Excel to allows for

connection to multiple data sources and allows joins between tables from different data sources

For CONNX Version 12.5 and above we've added Read/Write capability. In previous versions users only had the ability to import their data into Excel and manipulate it. This feature was implemented for security reasons to uphold the integrity of your data and give a wider scope of users in your organization access to your data without worry. While this feature is still in place, we've added the "write" function which allows higher lever users the ability to make changes to your data, safely and easily using Excel.

CONNX Client PC Functions

Functions for the client PC in the CONNX distributed architecture include the following:

Data Conversion	Sorting
Metadata Retrieval (CONNX CDD)	Grouping
First Pass SQL Optimization	Extended SQL Functions
Partial Joins	CONNX Security

CONNX Server Functions

Functions for the data server in the CONNX distributed architecture include the following*:

Indexed Retrieval	Data Compression (on Request)
Non-Index Retrieval	Remote Procedure Calls (RPCs)
Partial Joins	Database Security Client

*DataFlex, POWERflex, Oracle, and DB2 the server functions are resident on the client PC.

Requirements for CONNX 13.0

Data (Host) Server Requirements

Please see the accompanying documentation for additional System Requirements.

Database	Hardware	Network	Operating System	Memory/ HD requirements
Digital RMS (any version)	Compaq/DEC VAXServer Compaq/DEC AlphaServer	UCX 3.0 or above compatible TCP/IP Software	OpenVMS/VAX OpenVMS/Alpha {AXP] VMS 5.3 and above Itanium 64-bit	12mb VAX 32 mb Alpha Working Memory 20k Blocks HD avail
Oracle Rdb (version 4.1) (version 6.0 and above)	Compaq/DEC VAXServer Compaq/DEC AlphaServer	UCX 3.0 or above Compatible TCP/IP Software	OpenVMS/VAX OpenVMS/Alpha [APX] VMS 5.3 and above	12mb VAX 32 mb Alpha Working Memory 20k Blocks HD avail
Oracle DBMS (version 4.3 and above)	Compaq/DEC VAXServer Compaq/DEC AlphaServer	UCX 3.0 or above Compatible TCP/IP Software	OpenVMS/VAX OpenVMS/Alpha [APX] VMS 5.3 and above	12mb VAX 32 mb Alpha Working Memory 20k Blocks HD avail
Oracle RDBMS (version 7.3 and above)	Compaq/DEC VAXServer Compaq/DEC AlphaServer Personal Computer (Intel) Sun Workstation IBM RS/6000(AIX)	SQLNet 2.x And TCP/IP	OpenVMS/VAX OpenVMS/Alpha VMS 5.3 and above [APX] Microsoft Windows NT UNIX(ANY)	12mb VAX 32 mb Alpha Working Memory 20k Blocks HD avail

Database	Hardware	Network	Operating System	Memory/ HD requirements
C-ISAM	SunSparc RS/6000 Intel HP Server	TCP/IP	SunOS AIX Linux SCO HPUX Windows XP SP2 and above	5 mb of HD space 32mb RAM
DISAM	SunSparc RS/6000 Intel HP Server	TCP/IP	SunOS AIX Linux SCO HPUX Windows XP SP2 and above	5 mb of HD space 32mb RAM
Micro Focus	SunSparc RS/6000 Intel HP Server	TCP/IP	SunOS AIX Linux SCO HPUX Windows XP SP2 and above	5 mb of HD space 32mb RAM
DataFlex & PowerFlex (any version)	Personal Computer Sun Workstation	Any supported protocol under Windows	Windows, UNIX	
Any OLE DB Compliant data source Sybase Informix SQL Server	No requirements except those of the database itself and the third-party driver An ODBC Level 2- compliant driver must exist for the platform and database.	TCP/IP software Requirements of third-party driver	No requirements except those of the database itself and the third-party driver	No requirements except those of the database itself and the third- party driver

DB2 Database	Hardware	Network
DB2/6000; DB2 UDB for AIX	AIX 4.3 and above	TCP/IP and SNA/LU 6.2
DB2/MVS V4R1 and above	MVS	SNA/LU 6.2 only
DB2 UDB for z/OS and OS/390	z/OS and OS/390	TCP/IP and SNA/LU 6.2
DB2/400 V3R1 and above	OS/400	SNA/LU 6.2 only
DB2/400 V4R2 and above; DB2 UDB for iSeries	OS/400 and iSeries	TCP/IP and SNA/LU 6.2
DB2 UDB Enterprise Server Edition	Windows XP SP2 and above	TCP/IP and SNA/LU 6.2
DB2 UDB for Linux Enterprise Server Edition	Linux	TCP/IP

CONNX for VSAM Product	Operating System	Supported File Types	Network Software	CICS Version/Release
CONNX for CICS/VSAM	OS/390 and z/OS	VSAM	TCP/IP V3R2 and above	V4R1 or TS 1.x and above
CONNX for VSAM / QSAM / PDS	OS/390 and z/OS	VSAM / QSAM / PDS	TCP/IP V3R2 and above	N/A
CONNX for CICS/VSAM	VSE 2.3 and below	VSAM	TCP/IP (CSI / IBM), Barnard TCP/IP Stack	V2R3 and below
CONNX for CICS/VSAM	VSE 2.4 and above	VSAM	TCP/IP (CSI / IBM), Barnard TCP/IP Stack	TS 1.1.1 and above

Adabas SQL Gateway (CONNX for Adabas) Product	Operating System	Network Software
Adabas	OS/390, z/OS, VSE, Windows XP SP2 and above, Solaris, HPUX, AIX, VSE, Linux Intel, zLinux	TCP/IP, Barnard TCP/IP Stack (VSE only)

CLIENT PC REQUIREMENTS

	Minimum	Recommended
Available space on hard drive	150 MB	250 MB
OS	Windows XP SP2 and above (32/64bit)	Windows 7 and above (32/64bit)
Network Connectivity	Microsoft TCP/IP OR Oracle OCI Client version 9 and above (Oracle Only) OR SNA/LU.6.2 with TCP/IP or DLC Network Protocol (DB2 Only)	Microsoft TCP/IP OR Oracle OCI Client version 9 and above (Oracle Only)
Access or permission on the appropriate databases	YES	YES

Unix Client System Requirements

PC Linux Client System Requirements	
Hardware	Processor: Intel Pentium class or above Memory: 512 MB
Operating System	Any Linux OS which supports Linux Kernel 2.6.18 or above, for example, Fedora Core Release 6 or above, RedHat Enterprise Linux, version 4 or above, or SUSE Enterprise Linux 11 or above. Please see the documentation for your specific Linux distribution to determine the Linux kernel version.
Free Hard Disk Space	50 MB
Software – ODBC Driver Manager	Any ODBC Driver Manager

Solaris Client System Requirements	
Hardware	Processor: UltraSPARC Memory: 512 MB
Operating System	Sun OS 5.8 or above
Free Hard Disk Space	50 MB
Software – ODBC Driver Manager	Any ODBC Driver Manager

AIX Client System Requirements	
Hardware	Processor: IBM e-Server P-Series or RS/6000 Memory: 512 MB
Operating System	AIX 5.x Operating System: IBM AIX 5L Version 5.1, system maintenance level 2 (64-bit) or Version 5.2
Free Hard Disk Space	50 MB
Software – ODBC Driver Manager	Any ODBC Driver Manager

HP-UX Client System Requirements	
Hardware	Processor: PA-RISC Memory: 512 MB
Operating System	HP-UX 11.0 (64-bit) or HP-UX V11.11i (64-bit)
Free Hard Disk Space	50 MB
Software – ODBC Driver Manager	Any ODBC Driver Manager

JDBC Pure Java Client Requirements

Requirement	Minimum
JDK*	1.3 for JDBC server. 1.7 for JMS server
Hard Drive Space	10 MB Free
Network Connectivity	TCP/IP

InstantdbSync Requirements

Requirement	Minimum
Operating System	64bit Windows Server class operating system
Source Database	SQL Server 2008 and above or MySQL 5.6 and above
Hard Drive Space	20 MB Free
RAM	8 GB

Obtaining a current JDK (Java Development Kit)*

JDKs are available through your platform vendor.

Platform	URL
Windows XP SP2 and above	http://www.oracle.com/technetwork/java/archive-139210.html
Sun Solaris	http://www.oracle.com/us/sun/index.htm
Linux	http://www.oracle.com/technetwork/java/javase/overview/index.html
VMS-Alpha	http://www.compaq.com/java/download/index.html
SCO Unix	http://www.sco.com/developers/java/
IBM (AS/400, OS/390, VM/ESA, AIX, z/OS,VSE)	http://www.ibm.com/developerworks/java/jdk/
SGI	http://www.sgi.com/partners/?/devtools/languages/javafaq.html
HP-UX	http://h18012.www1.hp.com/java/download/

* The platform vendor is usually the best source for platform specific JDKs. JDK required for CONNX client machine using JDBC only. Please follow your vendor’s instructions for installation.

Compatible Front Ends

OLE DB	ODBC	JDBC	Application
	√	√	Any JDBC-compliant application
√	√		Any ODBC- or OLE DB-compliant application
		√	Apache Web Server
√	√		Borland C++
√	√	√	Borland Delphi
	√	√	Borland JBuilder
	√		Cognos Impromptu
√	√	√	Crystal Reports
	√		Dharma ODBC Integrator
	√		GIS (Geographical Information Software)
√	√	√	Internet Information Server (IIS)

OLE DB	ODBC	JDBC	Application
	√		JetForms
√	√		Microsoft Access
	√		Microsoft Excel (MSQuery)
√	√		Microsoft SQL Server (linked server technology)
	√		Microsoft Transaction Server (MTS)
√	√		Microsoft Visual Basic
√	√		Microsoft Visual Basic for Applications (VBA)
√	√		Microsoft Visual C++, Microsoft Visual Studio
√	√		Microsoft Visual Studio .NET
	√	√	Netscape (iPlanet) Enterprise Server
	√		Oracle Developer/Designer 2000
	√		Oracle Discover
	√		Oracle Heterogeneous Services
√	√		PowerBuilder
√	√		Paradox for Windows
	√		Sagent
	√	√	Star Office
		√	Sun Forte
		√	Sun Netbeans
		√	Sun Netra Web Server
	√		Visual FoxPro for Windows

Installation Instructions

Installation instructions for CONNX are available from <http://www.connx.com/products/helpdesk.html>

64 bit Considerations


When CONNX 13 is installed on a 64 bit Windows operating system, both the 32 bit and the 64 bit components of CONNX are installed. By default, the 32 bit components are installed in C:\CONNX32 and the 64 bit components are installed in C:\Program Files\CONNX. Both 32 bit and 64 bit executables are accessible from the Start Menu. Note: When CONNX is installed on a 32 bit Windows operating system, only the 32 bit components are installed.

Accessing 32 bit only data sources from 64 bit applications

It is possible to access a 32 bit only data source, such as Dataflex on Windows, C-ISAM/D-ISAM on Windows, etc. from a 64 bit application using the CONNX Enterprise Server Service (ESS). Using the ESS, a 64 bit application such as MS SQL Server can load the 64 bit CONNX client. The CONNX Solutions CDD can then be configured to access the 32 bit data source via the 32 bit Enterprise Server Service. This configuration allows the 64 bit client to call into the 32 bit ESS via TCPIP which, in turn, is able to load the 32 bit only DLLs used to access the data. The opposite is also true: if you have a 64 bit only data source that you need to access from a 32 bit application, you can use the 64 bit ESS to access the data and pass it to the 32 bit CONNX client.


Configuring 32 bit and 64 bit components

CONNX is configured with the CONNX Configuration Manager. With CONNX 12, the CONNX Configuration Manager can be used for managing both the 32 bit and 64 bit components.



There is a radio button to select which registry setting to configure. Usually, the registry settings will be the same for both the 32 bit and 64 bit components. Checking the “Apply changes to both 32bit and 64bit registry” checkbox will cause a setting made for one component to be made for the other, as well. There are some settings, however, where it may be necessary to maintain different values for the two components; for example the port the ESS listens on. In these cases, this check box should be unchecked when changing the value.

On the InfoNaut tab and the User DSNs tab, the settings are not differentiated between 32 bit and 64 bit. The selection radio buttons are not displayed, and any settings that are made automatically apply to both.



Default Ports

The following is a list of 32 and 64 bit components and the default ports they listen on:

Component Name	Default Port
32 bit Enterprise Server Service	6500
64 bit Enterprise Server Service	6502
32 bit JDBC Server	7500
64 bit JDBC Server	7502
License Server	7501
InstantdbSync Message Queue	9200
InstantdbSync Controller	9205
JMS Server	7600

If CONNX is installed in an environment where a firewall is present, these ports need to be opened.

User Account Control (UAC)

CONNX 12 fully supports environments running the Microsoft Windows User Account Control. Because the CONNX Configuration Manager requires read/write access to the registry, the following notes apply:

- **CONNX Configuration Manager**
 - **With the UAC on:**

The CONNX Configuration Manager requires administrator level permissions to have read/write access to all registry settings except for those located on the InfoNaut and User DSN tabs. If the user has administrator rights on the system, the CONNX Configuration Manager will request promotion to an administrator level. Depending on UAC settings, the user may or may not be prompted. If the user does not have administrator rights, the CONNX Configuration Manager will allow the user to view settings in a read-only mode. The InfoNaut and User DSN tabs will function in a read/write mode even if the user does not have administrator rights on the system.
 - **With the UAC off:**

No special permissions are required. The CONNX Configuration Manager will function without prompting as long as the user has authority to execute programs and has access to the CONNX directory.

New Platform Support

- added linux support for dataflex huffman compression table
- added support for Teradata
- added support for MySQL as an InstantdbSync source database

Critical Changes

Please take special note of the following corrections and changes (listed below as well) as they may cause a change in expected behavior.

- We have discovered an incompatibility with the Oracle 10.1 Instant Client. For users who wish to use the Oracle 10 instant client, they must upgrade to 10.2.
- Oracle 8 is no longer supported. This applies to both the Oracle client as well as the Oracle 8 database.
- When upgrading from previous versions of InstantdbSync or Open Systems Event Replicator, the message queue must be drained prior to performing the upgrade. To drain the message queue, stop all activity on the source database(s) and wait for the queue length to go to 0 on the status screen. Note: This only applies to customers that are replicating to relational targets, ACD targets or JMS targets. This does not apply to Open Systems Event Replication customers doing Adabas to Adabas replication.
- Changed default value of CNXNOPREAUTHORIZE from 0 to 1 for IMS server.
- Changed default GROUP name for CICS server installations to be CNXGROUP instead of CNXvrr
- All shipping executables are now digitally signed
- Fixed default value for CONNECTIONPOOLING. The correct default is 2.
- Made default data type on oracle table creation for longvarchar, longvarbinary and nlongvarchar as - CLOB, BLOB and NCLOB respectively.
- The underlying DataSync database CONNXStore has been upgraded to PostgreSQL version 9.3.4. The installer will run a conversion process. The database will be backed up prior to conversion.
- Update the data types shown in the Query Builder to show ANSI 92 data types instead of the .net datatypes.
- REGION parameter in JCL for mainframe servers changed to 0M.
- Mainframe servers - changed default for ALLOWMIXEDPWD from 0 to 1. This will allow mixed case passwords on the mainframe data servers. If mixed case passwords are not enabled on the mainframe, this setting should be set to 0 in CNXPAMS.
- DECNet is no longer supported on OpenVMS systems. This affects RMS, RDB and DBMS server components.

- Replaced CONNX Solutions program folder on the Windows Start menu with separate folders for each CONNX product. This change makes us consistent with the new Windows 10 menu structure.

Changes/Bug Fixes for CONNX 13.0

New Features

- **InstantdbSync:**
 - Added support for MySQL as a source on Linux (Intel)
 - Added ACD support for MySQL and Adabas as source
- **Added zOS passphrase support**
- **Added support for Adabas SSX security on LUW – to activate, set configuration ADA_SECURITY to 1**
- **Added Adabas timestamp with timezone support**
- **Added support to Excel Add-In to allow editing of filtered data sets (data sets generated with SQL Statements containing where criteria)**
- **Added support for building Adabas sample tables**
- **ES-17649 - allows user to bypass cap by specifying a session name of "BYPASSCAP" in JDBC server**

These are the major issues that were resolved and features that were added in CONNX 13.0:

CONNX client/server

- Corrected errors in MySQL CLOB/BLOB support
- Fix error where MAXROWS was returning an error with no text
- Fixed crash in debug tracing with large error messages
- Implemented a new function in CONNX called CNXRawConvert.
The Syntax is as follow:
CNXRawConvert(<datatype name>, <expression to convert>, <offset of bytes into expression>, <length after offset to convert>, <precision if necessary>, <scale if necessary>, <codepage if necessary>)
This function provides the full power of all of the data conversion routines in the CDD from a SQL function which allows a section of a string to be easily pulled out and converted as required.
- Added Auto Bulk Mode feature
- Fixed errors with MySQL Decimal data type
- Fixed problem with extended characters in key data not working properly since in PostgrSQL
- Fix for Drop index when going through the ODBC interface
- Fixed error when doing an update to a BINARY or VARBINARY field with a where clause in the SQL statement
- Fix for usage of bit datatype in where clause for PostgreSQL
- Fixed situation where error messages for MAXROWSCOMPARED and MAXROWSFETCHED were the same. They are different now.
- Fixed problem with IMS Insert errors not being diagnosed properly
- Fixed problem where IMS rollback was always returning an error

JDBC Server

- Fix problem with JDBC Threads not being Colnitialized

CDD Manager

- Fixed error where scale was missing in RMS import of a Dibol file
- Fix for security on information_schema

InstantdbSync/Open Systems Event Replicator

- Fixed MYSQL producer error where it would not package key buffer correctly if key on source table was NOT the first column.
- Fixed error in SQL Server EP where the EP was not replicating records that contained date/time values that were formatted with non-US formatting
- Fixed problem where there where two license seats being taken by the Rep Admin
- Fixed iTrac issue ARN-74: "Initial state error on columns with position after decimal point"

InfoNaut

- Auto size columns for very small result sets - less than 60 columns, less than 500 rows

License Server

- Fix to MSU license error message – display the server name of the connection instead of the logical database name

DataSync

- Fixed error in transformation server when the source did not have an index
- Fix to enable transform support for tables that have CNXROWID selected
- Fixed problem where the Custom Index name on the Transform properties window could not be cleared once it was set

Excel Add-In

- Corrected layout errors
- Fixed error when opening an odc for edit in a worksheet, if the process of setting up the worksheet fails in getting the table schema, it leaves the worksheet in a state where another odc cannot be selected.

Install

- Added support for building Adabas sample tables

Changes/Bug Fixes for CONNX 12.5

New Features

- **InstantdbSync:**
 - Added support for MySQL as a source database (Windows only)
 - Added support for ACD targets
- **Added support for KPiSync usage quotas**
- **Added support Read/Write capability to Excel Add-in**
- **Added native MySQL driver for Windows**
- **Added feature to KPiSync to allow SMTP Mail to be sent when KPI's are pushed, or there is an error.**
- **Enhanced support for deploying .NET applications to the cloud.**

These are the major issues that were resolved and features that were added in CONNX 12.5:

CONNX client/server

- Oracle 8 client and database is no longer supported
- An incompatibility exists between CONNX and the Oracle 10.1 Instant client. Users wishing to use the Oracle 10 Instant client must upgrade to the Oracle 10 Instant client version 10.2.
- Fixed bug in temp table join & lookup logic where the join field is some type of integer.
- Optimized Adabas server to now use a single S1 to get the count for queries like "select count(*) from table where col = ? and col = ?" when those columns are in a single adabas descriptor.
- Adabas server has been optimized to no longer issue an L1 following an S1 if the ISN Quantity is 1.
- Fix for bug in dataflex where we used the global variable isrecnum for the current file record number.
- Fixed issue with the UNIQUEID not being thread safe for CISAM, DISAM and Microfocus.
- Added ISALIVE check for Adabas.
- Added support for LIMIT in CONNX Views.

- Fixed Oracle crash when connection fails and connection is reattempted
- Fixed error with Oracle blob insert/select where blob field is NULL
- Added support to .NET data provider which enables .NET applications to be deployed in the cloud without installing any CONNX client code other than the .NET driver itself.
- CONNX will now return an error if the IN list is empty - for example:
select ... from table .. where col in ()
- Improved the Adabas multi fetch logic to always fetch the maximum when using Sx calls.
- Added Client based time zone selection support for Adabas via a new configuration setting CONNX.TIMEZONE - this will get passed to the Adabas database on the OpenDatabase call in order to support the timestamp with timezone data types.
- Corrected IMS import logic to treat a return code of 4 as a warning instead of an error during a FABMMAIN based import.
- Fixed problem with 1145 code page translation from 1252.
- Fixed an IMS error caused by the server building an SSA with a column name that came from a COBOL copybook but was not defined in the DBD.
- Fixed problem where the IMS DLI server would crash in certain circumstances.
- Modified IMS server to send the actual IMS error and reason codes back to the client instead of internal CONNX error codes.
- Added support for IMS secondary indexes.
- Performance optimizations for IMS inserts.
- Improved client side error messages for IMS errors.
- Fixed IMS ODBA server error where SETS (commit) was failing when running in ODBM.
- Fix for grandchild IMS segments not always returning all rows.
- Fixed IMS DLI JCL templates to allow calls into DLI without having to APF authorize PSB/DBD libs.
- Changed default value of CNXNOPREAUTHORIZE from 0 to 1 for IMS server.

JDBC Server

- Corrected an issue where imbedded Nulls from JAVA did not translate correctly in JDBC Server.
- Added pre-fetch logic into the java client.
- Implemented support for JDBC IsValid function.
- Fixed memory overwrite that could cause the JDBC server to hang.

RCI Client

- Fixed a problem with statements not being released when the query results in an error on prepare.
- Added a statement Cache to the RCI Layer to help improve performance for queries issued in a loop thousands of times with different parameters.
- Optimization so exec & fetch with RCI happens in a single TCP/IP RPC.
- Removed unnecessary GetNextResult call from RCI.
- Fixed issue where updatable cursors were cached in the RCI.
- Fixed problem with low values not being generated properly when S1MODE=1 was enabled.

CDD Manager

- Added support to Adabas sysobjh import to use the length in the sysobjh file for LA fields instead of the max.
- Fixed IMS import problems where some segments were being skipped.
- Fixed an error when importing using an index text specification file. If the key field in the file had the same name as the field in the copybook, the CDD manager would crash.
- Fixed an error when doing a copybook overlay over a FABMMAIN import and the key field had the same name as the one in the copybook, the offsets for the remaining copybook fields were off by the length of the key field.

InstantdbSync/Open Systems Event Replicator

- Added support for MySQL as a source database on Windows
- Added support for ACD (add, change and delete) replications. These replications allow the tracking of the changes rather than replicating data.
- Fixed problem where SQL Server EP did not recognize any SQL Server Native Client with a version number higher than 10.
- Added support for ACD (add, change and delete) targets. (SQL Server as source)
- Fixed error with SQL Server updates when the primary key on the source was modified.
- Added support for SQL Server 2016 as a source database.
- Fixed problem with A2A replication using Adabas 6.4 and above where a status change for a group of replications would fail for the entire group if any of the individual replications failed to change status.
- Fixed memory leak in the Administrator's Status screen.
- Fixed Adabas to Adabas initial state failure if more than 150 tables were involved in the initial state.

KPiSync

- Corrected layout problems on mobile devices
- Added context menus to the main grids of the admin
- Allow license to limit users/kpi's/number of pushes per month. This provides flexibility in pricing.
- Added feature to allow SMTP Mail to be sent when KPI's are pushed, or there is an error.
- Modified the logging mechanism to no longer use the Windows Event viewer. Log messages are now written to the file "KPiSyncAdmin.log" in the KPiSync folder.

InfoNaut

- Fixed a problem with CONNX Security where a table was not visible in the table list of infonaut when you revoke select access to one or more columns of the table.
- Fixed minor GUI issues.
- Fixed issue that was not allowing passthrough queries to return a result set.

License Server

- Added support for KPISync user/kpi/push quotas.
- Fixed problem with serial number lengths and OEM licenses.

DataSync

- Fix for bug in Datasync where the oracle NUmber data type had 20 bytes of "garbage" after it - causing incremental syncs to detect changes where there were none
- Added support to allow target metadata to be forced all upper case or all lower case.
- Fix the appearance of the transform properties form.
- Fix for DataSync bug where using a source table with a space in the source table of a transform caused errors when trying to run a sync on that transform.

Excel Add-In

- Added Read/Write capability

Install

- Suppress the error about starting the CONNX DataSync Schedule Service during the install.
- Fixed the z/OS installer to display an error if the FTP command line fails.
- Fixed the z/OS installer to allow spaces in the folder name where CONNX is installed.
- Improved logging for z/OS and z/VSE installers.
- Fixed issue on zOS installer where the install did not stop if there was a timeout error.
- Added Active and Passive FTP mode options for the z/OS installer.
- Fixed error with registering the KPiSync scheduling service.

Changes/Bug Fixes for CONNX 12 SP3

New Features

- **InstantdbSync:**
 - Added support for selectively undeploying replications from the Deployed Replications Tab
- **Added support for Teradata**
- **Added support for Adabas mixed case FDTs**
- **Added in support for getting the proper schema timestamp when importing Adabas tables**

These are the major issues that were resolved and features that were added in CONNX 12 SP3:

CONNX client/server

- Fixed fractional error for timestamps in temporary tables
- Fixed error with unsupported Oracle time functions – now process by CONNX
- Fixed problem with NVARCHAR data not being returned correctly in Oracle
- Fixed error where data server on VSE did not always shut down properly when using the Barnard TCP/IP stack
- Added support for Teradata
- Fixed Group By error on column alias that is **also** a column name
- Fix for problem with "ODBC Only" where an error returned by SQLFetch was not diagnosed properly
- Fixed DISAM error message on DROP TABLE when in same connection after insert
- Fixed problem where a trailing statement terminator semicolon caused a case statement query to fail
- Fixed errors with 8 byte integers (BIGINT)
- Fix for dataflex descending/ascending indexes
- Fixed RDB problem on Itanium systems where REAL data type not working correctly
- Added in support for getting the proper schema timestamp when importing Adabas tables

- Added support for mixed case Adabas FDTs
- Fixed problem with keyword names and name quoting
- Added setting that will always allow NULLS to be inserted - useful for IMS variable length segments
- Fixed problem with ROLLBACK not working for IMS with the DLI (batch) interface
- IMS DLI - Add configurable CNXPARMS parameters for the DFSRRC00 parameters. They are in the form of IMSxxxx where xxxx is the IBM name of the parameter
- Fixed intermittent Oracle timestamp bug
- Fix for error reporting with ODBC interface

JDBC Server

- Fixed error where we were not issuing rollback on transactions placed into the pool

CDD Manager

- Fixed problem where when turning security on for one column of a table, the table is no longer visible to Infonaut
- Fixed problem with left over data when a user is deleted
- Added powerhouse import to DISAM
- Fixed errors with copybook imports for IMS
- Fix for IMS Import with occurs clause not setting code page properly
- IMS - fixed FABMMAIN output parsing error. If there were more than 9 PCBs and the 10th or greater PCB referenced a different DBD, it was not being processed

InstantdbSync

- Fixed error in EC engine where it failed to initialize the record loaders after coming up from a bad shutdown
- Fixed problem where the EC was logging initial state replay errors when debug was turned off
- Added the ability to selectively undeploy a replication from the deployed replications tab
- Fix to allow the Enterprise and Desktop adapters to create tables with Unicode wide characters
- Fixed license failure in controller
- Corrected logging of table name on startup
- Fixed a problem where if the user changes the CONNX User password in the cdd, but does not change it in the Rep Admin config servers screen before they deploy again it leaves the controller in a wait state

KPiSync

- Fixed notification errors

InfoNaut

- Fixed issues when doing updates in the grid

License Server

- Fixed problem where the license server crashed if the license database was unreadable. It now issues an error message in the log

DataSync

- Fixed error where transformation server was unable to create Oracle index when multiple indexes use the same column
- Fixed some cryptic error messages
- Prevent CONNXStore from logging to the Windows Event Viewer
- Solve problem where the hosts file contains an entry like <machine-address> localhost which effectively removes localhost from being used as a machine name with psql
- Corrected problem with inserts into RDB with NULL values

Install

- Fixed problem where temp files were not being written to tempinst directory
- Prevent installation from continuing if the upgrade from previous version of CONNXStore fails.
- Fixed error where installer would sometimes reboot without a prompt before the installation was finished
- Fixed error that occurred if the tempinst folder did not exist
- Fixed FTP errors when installing to VSE and OpenVMS
- Fixed JCL errors when installing IMS to zOS with SMS enabled
- Changed IMS Utility HLQ to fully qualified path

Changes/Bug Fixes for CONNX 12 SP2

These are the major issues that were resolved and features that were added in CONNX 12 SP2:

CONNX client/server

- DECNet is no longer supported on OpenVMS systems. This affects RMS, RDB and DBMS server components.
- Added new Wizards to simplify the creation and deployment of CONNX SSIS packages in SQL Server.
- Fixed crash on some Unix systems when gethostbyname was called and there was no entry in the hosts file.
- Added SESSION support for Adabas so that the RCI Session name appears in show sessions
- Added support for Redshift
- Added support for Netezza
- Added create table support for HIVE
- Added {noclientfilter} option so that all client filtering will be bypassed
- Fixed crash when stopping Adabas data server on mainframe.
- Fixed a bug found from the incorrect column size being returned on an nvarchar column in the schema tables
- Fix problem where OLEDBProvider would not properly handle multiple errors - last one was always overwritten
- Fix bugs in ODBC implementation of Create Table - and added redshift support to automatically lowercase the tablenamees
- Fixed Abend s913 on mainframe data servers. This error occurred if the user that the data server was running as did not have permission to open a spool file
- Fixed crash when querying schema table VIEW_TABLE_USAGE and there are views in the CDD that reference non-existent columns
- Fix problem where subqueries in UNION statements would incorrectly attempt to resolve columns from the "top" union statement instead of its parent.
- Fix so that a call to get next rowset does not destroy the context for prepared passthrough queries
- Fix problem with order in which we were releasing OLEDB Interfaces

- Fix for problem with mixed desc/asc indexes with Dataflex
- Added a guard to prevent a crash when a query is issued on a connection that has not been established.
- Fixed problem where we generated the wrong key length in the format buffer for MUs within a PE for Adabas
- Added the ability to create a TIMESTAMP column for Oracle
- Fixed a memory overwrite in the code that processes Adabas blobs
- Fixed IMS import errors
- Fixed memory leak in IMS DLI server
- Fixed bug where IMS DLI server was not issuing rollbacks
- Added a new setting FORCESHARECONNECTIONS that will force connection sharing on read/write connections (needed for IMS Batch - where we can only have 1 data server instance)
- Fix bug where {maxrows x} with order by was not returning correct data.

JDBC Server

- Added support for updatable cursors
- Change default JDBC Packet size from 8k to 64k
- Adjusted JDBC logging so we redisplay the SQL statement on re-executes of prepared statements.
- Fix to set the Adabas session ID (even if not set) unconditionally, so old session IDs are not displayed if connection pooling is enabled
- Added USER name logging
- Fix case problem with connxJdbc.jar (should be connxjdbc.jar)

CDD Manager

- Updates to the CDD Manager GUI. This updates the menu, toolbar, icons in tree view and the grid controls. Size and position of frames are now sticky.

InstantdbSync

- Enhanced processing of updates to use an in-memory cache. This cache greatly improves the performance of update processing.
- Fixed timeout error when deploying a large number of producers
- Fixed error with mixed mode Adabase to Adabase (A2A) and Adabas to Relational (A2R). If an A2A and A2R replication were deployed from the same EP and then one of them was removed on a redeploy leaving only the other, the admin showed the replication as being removed, but it was still replicating in the EP.
- Fixed bug in the SQL Server EP recovery code so that if the connection to the source database is lost, instead of shutting down and restarting, the EP will stay up and go into a retry loop. This allows it to stay responsive to status requests and work messages.
- Fixed erroneous message stating possible skipped transaction.
- Fixed intermittent crash in engine when it was starting after a bad shutdown

KPiSync

- Added subscription support. The KPiSync admin can now specify which KPi's or groups of KPi's will be seen by which users.

InfoNaut

- Added missing hotkeys to main window
- Fixed problem where infonaut was not using the correct regional date format for dbms date values in non-US regions.

License Server

- Added support for offline revoke

DataSync

- Fixed errors while upgrading from CONNXStore version 7 to version 9.
- Corrected CONNXStore errors when installed in a directory that contains spaces in the name.
- Fixed bug where Compact Database issued an error saying it couldn't find a role name
- Fix for DataSync CDC Transform bug, if multiple unique indexes were present, DataSync was just picking the best one. However now it saves the index that the user selects so the sync uses that one. Only for CDC transforms.

Install

- Modernized server installer and Unix/Linux installer GUI's to be consistent with industry standard look and feel.
- Enhanced Unix/Linux installers to allow SFTP and SCP transfers in addition to the previous FTP and manual modes.
- Enhanced Mainframe installer to allow RMODE=SPLIT.
- Enhanced z/VSE installer to allow selection between active and passive FTP.
- Replaced CONNX Solutions program folder on the Windows Start menu with separate folders for each CONNX product. This change makes us consistent with the new Windows 10 menu structure.
Note: the shortcuts that were previously in CONNX Solutions | CONNX are now in CONNX Driver.

Changes/Bug Fixes for CONNX 12 SP1

These are the major issues that were resolved and features that were added in CONNX 12 SP1:

CONNX client/server

- Fixed problem where SQLCancel would sometimes hang
- Added SSIS Adapter support for SQL Server 2014
- Fixed error with Excel AddIn-unable to create an ODC in Excel when using either the QueryBuilder or typing in your own query
- Fixed problem where trailing spaces were being inserted on passthrough updates when the source was CHAR and the target was VARCHAR.
- Fixed a crash that was due to EXISTS being optimized in such a way that there were two references to the same subsql
- Fixed a problem where long running queries (multiple days) would display incorrect time with {show sessions}
- REGION parameter in JCL for mainframe servers changed to OM.
- Mainframe servers - changed default for ALLOWMIXEDPWD from 0 to 1. This will allow mixed case passwords on the mainframe data servers. If mixed case passwords are not enabled on the mainframe, this setting should be set to 0 in CNXPARMs.
- Fix bug where GroupBy for CHAR data type did not honor the case sensitivity setting
- Fix to ensure that we do not fetch more elements that actually exist in an Adabas file (for mainframe)
- Fix for Adabas when MU within PE returning NULL for occurrences > 180 on mainframe when the MUPEX flag was not set on the Adabas file.
- Implemented DECODE/SWITCH for RDB
- Fixed intermittent crash when using {showsessions}
- Added CLIENTID to {showsessions} output for Adabas

JDBC Server

- Fixed problem where JDBC server would hang if there were more than 1024 connections

CDD Manager

- Fixed rotated array assistant - was not indenting properly
- Fixed import problems with long column and table names
- Fix for VMSBrowse - so that we will try logicals both with and without [000000]

InstantdbSync

- Fixed problem where updates were being replicated to the target when the primary key was a compound key (two or more columns in the key)
- Fixed error when the combined length of column names exceeded 1K on the source
- Fixed erroneous message in the log when a large rollback occurred
- Fixed errors when handling SQL Server compensation records in the transaction log.
- The following limitations exist in this release and will be addressed in future service packs:
 - BLOB/CLOB data types are not supported. This includes TEXT, NTEXT, IMAGE, VARCHAR(MAX), etc.
 - Computed columns have limited support.
 - The index used for replication cannot be a computed column.
 - The source table can have computed columns but the target table cannot.
 - If a computed column is part of the replication, the data in the computed column will not be replicated. In this case, the target column that the computed column is mapped to will receive nulls.

InfoNaut

- Update the data types shown in the Query Builder to show ANSI 92 datatypes instead of the .net datatypes.
- Fix for bug where the formatting was incorrect when a query with datetime fields was run before an SQL script.

DataSync

- Fixed bug where if "Create Transaction Log On Target" was checked, then CRC transformations did not work correctly.
- Fixed deadlock problem

Changes/Bug Fixes for CONNX 12

These are the major issues that were resolved and features that were added in CONNX 12:

CONNX client/server

- Fixed bug where like did not work with a literal ^ when properly escaped as ^^
- Added support for IBM DNS Name Resolver
- Fixed hang when reading large VSAM files on VSE with Barnard TCPIP stack.
- Fixed BLOB/CLOB errors in Oracle
- Fixed a problem on the mainframe data server where an existing listener would shut down if another instance attempted to start on the same port.
- Eliminated unneeded RC calls in Adabas data server.
- Fixed intermittent crash when doing select count(*) type queries with Adabas.
- Fixed lock when {maxrows} is used.
- Fixed problem where SHA-512 passwords were not working on HPUX.
- Fixed crash in schema table logic
- Fix for Oracle LONG, LONG RAW,CLOB & BLOB memory leaks and broken functionality
- Fixed crash in Oracle when doing an INSERT/SELECT with a CLOB/BLOB field
- Fixed connection problems with newer versions of DB2/400
- Fix problem with system tables being intermittently locked
- Added OpenVMS logical called CNXTRUSTEDRPC that will allow DCL type RPCs to be executed with a captive account. (default is turned off)
- Fixed problem with optimizer when using the ansi join syntax
- Added new feature that allows the user to completely control the physical file name on create table for RMS, CISAM, DISAM. Example: create table test(a integer) {physicalnameoverride 'CNXDIR:MYNEWTABLE.DAT'}
- Fix performance issue with accessing catalog tables
- Fixed error in the data server where it was printing incorrect values for some parameters to the log file.
- Fix for Adabas problem where high value was not being generated correctly for unpacked & packed values for a superdescriptor of type 'A'
- Added ODBC function escape {fn ...} support for INSERT(), TRUNCATE() and DATABASE()
- Fixed Bulk load support for SQL Server native client 10
- Fixed problem with intermittent password rejections on Linux
- Fix query plan show to show additional adabas optimizations that occur on cross table joins
- Added SUBQUERYROWESTIMATE to allow more accurate query plan reporting with some types of queries. See the CONNX User's Guide for more information

- Fixed a crash when a query attempts to use parent field from correlated in order by or group by – now returns an error
- Change UPDATE so we always use bind parameters even for null values, so oracle can get better caching performance
- Added DDL support for the IBM CLOCK data type for create table description
- Fix crash found in Select Tables when table type filter is not %
- Added in new function cnxphysicalname so the physical table name can be retrieved from the CDD using SQL
- Fixed problem where bulk mode was not working if the user imported a table with the NUMBER data type with no precision or scale specified
- Added Adabas support for Create index <name> - and drop index <name> (instead of using the special table_index_# index names)
- Added Adabas support for Alter Table Add Column
- Added ability to pass the actual CONNX user name to Adabas in addition to the existing unique ID
- Added diagnostics to catch Adabas ADANUC errors
- Fixed error reporting so we get the proper error message for SendMessage failures
- Fixed oracle truncation error with LONG
- Added support for uniqueidentifier data type (GUID) in SQL Server
- Added a new function - datechar, that can be used to convert a date to char in 4 different formats - usa, iso, jis, eur - this matches the DB2 CHAR function capabilities
- Fix problem where using a subquery as a temp table with DB2 would generate a syntax error
- Added CONNX.ORACLE.SHOWBULKMODE configuration setting so a text file will be created to explain why bulk mode for oracle is not used
- Fixed errors with Oracle Bulk Mode
- Added feature to force bulk mode for Oracle
- Added support for Adabas Blobs on mainframe
- Made default data type on oracle table creation for longvarchar, longvarbinary and nlongvarchar as - CLOB, BLOB and NCLOB respectively
- Fixed errors in Unicode CLOB support in Oracle
- Fixed errors with Oracle Big Double data type
- Implemented ODBC based distributed XA transactions that work with the Microsoft Distributed Transaction Coordinator
- Fix Now, current_date, current_time function for informix – corrected to use TODAY instead of SYSDATE
- Added NULL_SUPPRESSION column to the Adabas schema table COLUMNS
- Added RMS configuration option CNXRMSIGNOREERROR which will allow CONNX to skip corrupted RMS records

- Added support for password protected access to microfocus fileshare files - we now automatically pass the CONNX userid & password to microfocus for authentication
- Fix crash in 64bit ODBC Driver when using SQL Server linked server
- Fix for MySQL Driver crash when we attempt to reposition the cursor after we reach end of cursor
- Fix for Adabas response code 48 when making lots of simultaneous connections
- Added logic so that we propagate the SQL server isolation level to CONNX data server for distributed transactions

JDBC Server

- Fix crash/hang with JDBCConnectionPooling when a connection in the pool becomes "dead"
- Fix bug where GetNextResult does not return NULL if there are no more results to return
- Fixed compatibility problems with DBVisualizer
- Fix issue with GetMaxRows and GetFetchSize returning an error before statement execution
- Fixed problem where JDBC server would hang if there were more than 1024 connections

License Server

- Fix to license server logic - ensure that we re-establish the license server lock if we lose connection to the license server.
- Fix for problem where N-Tier would still open and maintain a socket connection to the license server for the non-ntier databases

CDD Manager

- Fix problem where COBOLFD import would actually work - but report "no tables found"
- Fixed bug where CDD Manager crashes when checking syntax of a view and max CDD security is turned on
- Added Big Endian equivalents of the Long date data types.
- Fixed a crash when importing Adabas tables that contain more than 30,000 columns.
- Fixed import bug where tables would not go into existing Database container when trying to add Dataflex tables to the CDD
- Allow SCT import offset to be overridden
- Fixed bug in Adabas DDL import where we were throwing an error if the DDL specified the name of SEQNO(0) as "ISN_<tablename>"
- Added the ability to not import normalized tables when doing Adabas FDT imports. (only import the "_FLAT" table)
- Fix problem where precision was not being imported properly with quadword decimal with RDB

Data Dictionary Viewer

- Added in display of Adabas short name of indexes in data dictionary viewer

CDD Comparison Tool

- Fixed memory leak
- Fixed performance problem when comparing very large tables
- Improved GUI performance when comparing very large CDDs.
- Fix problem with CDD compare tool - form was still disabled after the compare finished

InfoNaut

- Added a warning message and made the ODBC Trace button yellow when selected to make it obvious when ODBC tracing is on
- Added digital signature

DataSync

- Improved performance when creating groups with large numbers of tables
- Added support for creating in separate schema for Oracle
- Fixed problem where non-cdc transforms were showing up in the cdc import list in the cdd manager
- Added ability for DataSync to detect column names with invalid characters and automatically quote those names
- Fixed errors with some incremental syncs on tables with timestamps
- Fixed problem with incremental syncs not working when the last column in a table is not mapped (transformation server only).
- Moved sendmail logic from schedule code to command line code - so emails get sent if the command line interface is used
- Fix problem where the clustered index for SQL server was always the first index, not the unique index
- Added logic so force meta upper or lower applies to table name as well as column name
- Fixed error when performing a Savepoint sync on a cdc transform and the "Store Transaction log on target." checkbox is checked
- The underlying CONNXStore database has been upgraded to PostgreSQL version 9.3.4
- Fix problem with insert select from oracle LONG to SQL Server CLOB not working properly
- Added logic so if there is a general network error while we are testing the presence of the target table, we won't switch to full - we will just fail so we can attempt an incremental again in the future
- Fix to print message when a virus scanner is interfering with initdb.exe
- Fix for problem syncing CLOBs with DataSync for MySQL and Visual FoxPro
- Changed DataSync to make all target Date, time, timestamp fields as NULLABLE to deal with situations where different databases have different valid date ranges for dates and some treat values like 0000-00-00 as a real date

Install

- Fixed bug in installconnx UNIX script where we were doing an rm on the connx directory asynchronously in batch mode. This caused the script to check if the file was deleted before the delete was finished which resulted in an error
- Fixed bug where Open Systems Event Replicator Unix installation script always returned with 0 even when there was an error
- Fixed install errors with newer versions of RedHat Linux
- Fixed default value for CONNECTIONPOOLING. The correct default is 2.
- Added a “manual install” option on all Unix/Linux installers to allow an easier installation path for users who do not want to use the built-in ftp transfer method
- Added support to the RCI mainframe installer to allow the user to specify the target code page at install time.
- Changed default GROUP name for CICS server installations to be CNXGROUP instead of CNXvrr.
- Fixed bug in Installconnx unix/linux script where the script failed to run when installing with enhanced security and the base installation directory was different then the location of the install script

Please take special note of the following changes to the CONNX License Administrator.

There are some significant changes to the licensing starting with CONNX 12.

- Licenses have an embedded version – meaning that a CONNX 12 client will require a CONNX 12 (or higher) license. A CONNX license will work with a lower version client. For example a CONNX 12 license will work with a CONNX 11.x or CONNX 10.x client. The version of the CONNX license needed for a given version of the client is based on the major version number. For example, a CONNX 12 license will work with any minor version of CONNX 12 (i.e. CONNX 12.x, etc.)
- CONNX checks the CPU count of the data server against the CPU count generated in the license. On UNIX, Windows and VMS platforms, the CPU count is the number of cores the processor or processors have. On an IBM Mainframe platform, the MSU count is used for licensing purposes. In the License Administrator, the column “CPU/MSU limit” specifies the number of CPUs or MSUs the license is valid for.
- A license may be activated only once and is associated with the license server it was activated on.
- The License Administrator has a “revoke” feature which will remove a license from a server, this will allow a license to be moved from one server to another.

Upgrading from prior versions of CONNX/DataSync

- CONNXStore has been upgraded and is now based on PostgreSQL version 9.3.4 rather than version 7.
- When upgrading to CONNX 12 from an earlier version, the installer will upgrade the underlying CONNXStore database to the new version. A backup of the previous version database will be created. No update to the CDD is required.
- CDD files created in version CONNX 9.0 SP1 and earlier, must be opened in the current version of the CONNX Data Dictionary Admin tool and re-saved so that the CDD is saved in the correct format.
- The CONNXStore database will automatically be upgraded from prior versions of DataSync during installation. If there are any problems, a message should appear, and any problems should be listed in the datasync.log file.
- When upgrading from DataSync 9.0 SP2 and earlier, the first synchronization performed will be a Full Reload sync because the hash method for incremental synchronizations has changed. Note: this only applies when upgrading from 9.0 SP2 and earlier. If upgrading from any version after 9.0 SP2 to CONNX 12 or later, a Full Reload will not be required.
- If prior versions of CONNX were not uninstalled, the CONNX installation may prompt for a reboot so that the new components are properly registered and saved. This prompt to reboot should not be ignored.

Upgrading from prior versions of InstantdbSync/Open Systems Event Replicator

- When upgrading to version 12.5 from a previous version, the message queue must be drained prior to starting the upgrade. To drain the message queue, stop all activity on the source database and wait until the queue length field on the Server Status window of the Administrator shows zero.
- The requirement to drain the message queue only applies when the target is a relational database, JMS queue or ACD table. This requirement does not apply to Adabas to Adabas installations.

CONNX .Net Data Provider - Connection Pooling and Pooled Connection Timeout

This covers the corrections to the pooling and timeout of connections in the CONNX .Net data provider along with some inconsistencies with Microsoft's generic .Net data provider implementation. The following is the correct way to turn on or off the CONNX .Net Data provider connection pooling, and how to set the pooled connection timeout.

Connection Pooling

Turning connection pooling on/off; by default connection pooling is enabled. If connection pooling is enabled the .Net Data provider will hold a connection open for a specified amount of time after the `CNXConnection.Close()` function is called and use it the next time a connection is opened. Since the connection to the server was never closed, the opening of the new connection will be faster if a pooled connection is used.

- The first way to control connection pooling is through the `CNXConnection.PoolConnection` property.
- `CNXConnection.PoolConnection = true`; enables connection pooling in the provider,
- `CNXConnection.PoolConnection = false`; disables connection pooling in the provider.
- This property can be set before or after the connection has been opened, but must be assigned before the `CNXConnection.Close()` function is called.
- A second way to control connection pooling is through the connection string input to the `CNXConnection` object; Add "Pooling=true" to enable connection pooling, "Pooling=false" to disable connection pooling in the provider.
- Ex: "Persist Security Info=True;DD=c:\Test.cdd;UID=test;PWD=test;Mode=ReadWrite;Pooling=False;"

Pooled Connection Timeout

Only used when connection pooling is enabled, this setting controls how long a connection will remain in the pool while not in use. The input is in seconds, so setting it to 20 would mean the connection will remain in the pool for 20 seconds before it is closed. The default setting is 60 seconds, an input value of 0 means there is no timeout. This setting is only used when connection pooling is enabled.

- The first way to set this is with the `CNXConnection.ConnectionPoolTimeout` property.
- `CNXConnection.ConnectionPoolTimeout=10`; connections will last in the pool for 10 seconds after the `CNXConnection.Close()` function is called, before the connection to the server is closed.
- This property can be set before or after the connection has been opened, but must be assigned before the `CNXConnection.Close()` function is called.
- The second way to set this property is through the connection string input to the `CNXConnection` object; Add "Connection Lifetime=25" to set the time in seconds the unused connection will last in the pool.
- Ex: "DD=c:\Test.cdd;UID=test;PWD=test;Pooling=true; Connection Lifetime=25;"

CNXConnection.ConnectionTimeout property change

This property has been changed for clarity and consistency with the Microsoft generic .Net data provider implementation. The description that appears with this function in Visual Studio has been changed to “(Read Only) The time (in seconds) to wait for a connection to open. This is not controlled by the CONNX .Net Data Provider”. Also, since this property should have been read only, it has been changed to read only in the CONNX .Net Data Provider.