



ConnX 14.0

Release Notes
October 7, 2019

Table of Contents

ConnX 14.0 Release Notes	3
Overview	3
ConnX Architecture on Windows	4
ConnX Architecture on UNIX	5
ConnX Architecture (ODBC/JDBC/OLEDB/.NET Provider)	6
ConnX Client Engine for Windows	6
ConnX Client Engine for UNIX	6
ConnX Data Dictionary	6
ConnX Server	6
ConnX JDBC Thin Client	6
ConnX JDBC Server	6
ConnX JDBC Router	7
ConnX DataSync	7
ConnX DataSync Transformation Server	7
ConnX InstantdbSync	7
ConnX KPiSync	8
ConnX Excel Add-in	9
ConnX Client PC Functions	10
ConnX Server Functions	10
Product Installation	11
Requirements for ConnX 14.0	11
Data (Host) Server Requirements	11
CLIENT PC REQUIREMENTS	14
Unix Client System Requirements	15
JDBC Pure Java Client Requirements	16
InstantdbSync Requirements	16
Obtaining a current JDK (Java Development Kit)*	17
Compatible Front Ends	17
64 bit Considerations	19
User Account Control (UAC)	21
New Platform Support	22
Critical Changes	22
Changes/Bug Fixes for ConnX 13.8	33
Changes/Bug Fixes for ConnX 13.5	38
Changes/Bug Fixes for ConnX 13.0	41
Changes/Bug Fixes for ConnX 12.5	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 14.0 Release Notes

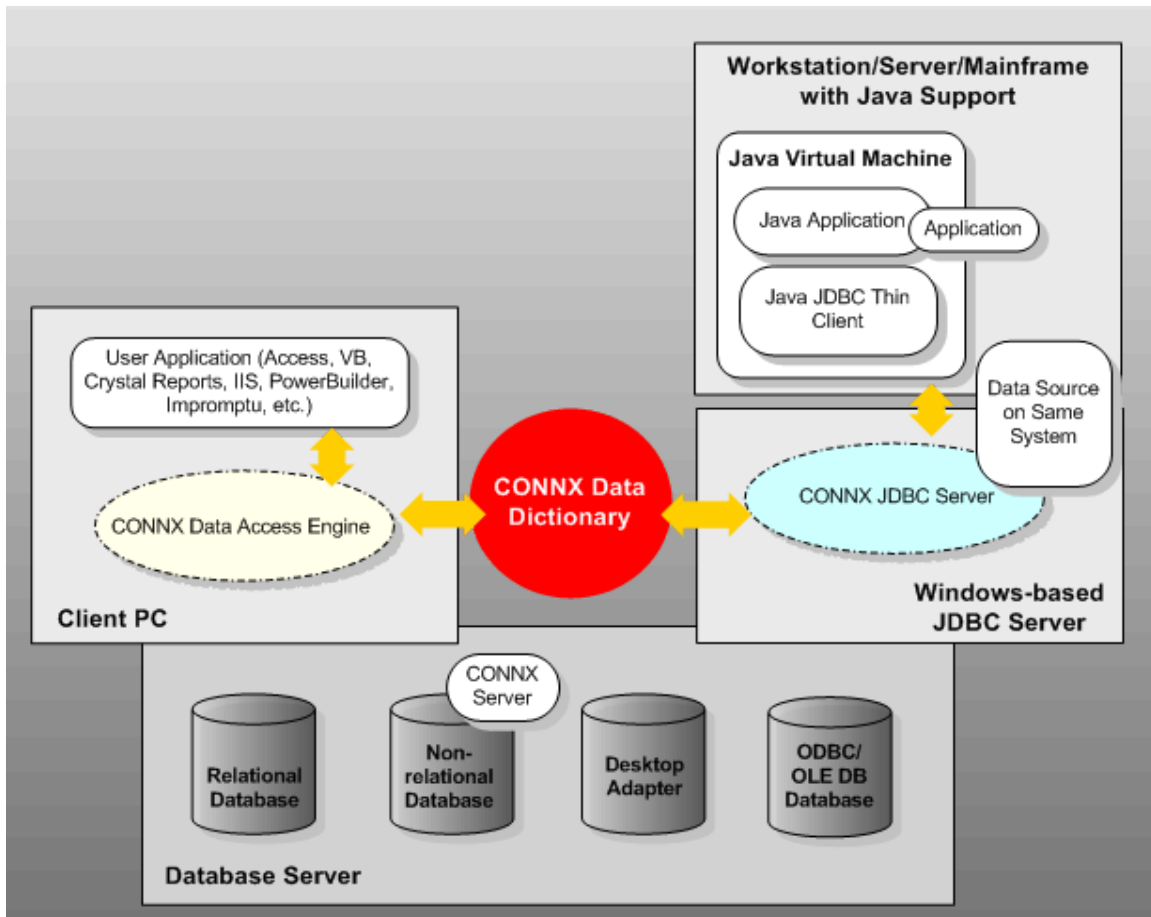
Overview

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.

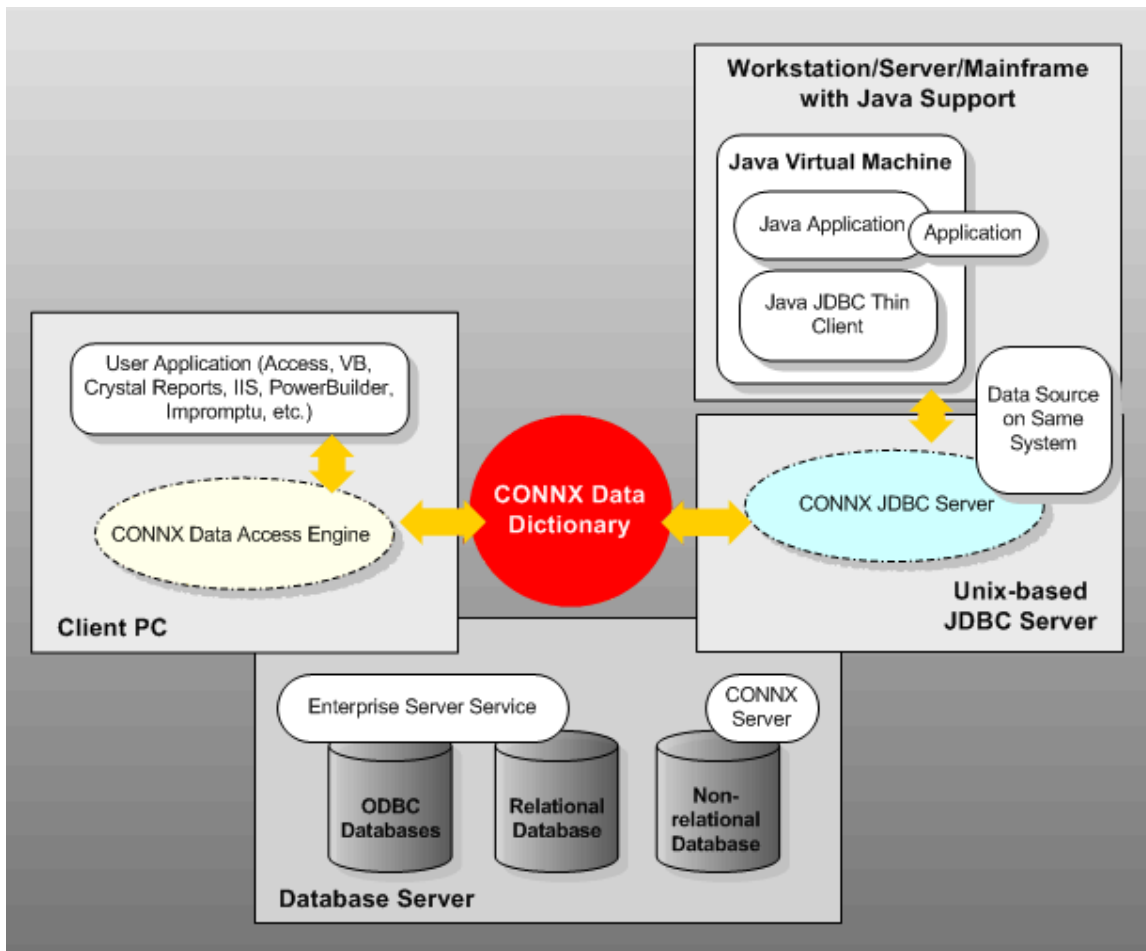
Here are just some of the implementations in which ConnX 14 can be utilized:

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

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/Open Systems Event Replicator 13.5

- Command line support for Pause/Resume replication, start Initial State and Get Status functions
- Performance statistics displayed in Replication Administrator

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

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 KPIs by users.

ConnX Excel Add-in

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.

Product Installation

Important:

The Adabas SQL Gateway product installation is now available in the Software Download Center (SDC) in Empower: <https://empower.softwareag.com> (login required). Please consult the file INSTALL.TXT contained in the download file for further details on installation issues. Installation instructions can in the Installation Guide located on the Empower web site.

Requirements for ConnX 14.0

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)	TCP/IP	OpenVMS/VAX OpenVMS/Alpha VMS 5.3 and above [APX] Microsoft Windows Server 2012 and above 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 HPUX Windows Server 2012 and above	5 mb of HD space 32mb RAM
DISAM	SunSparc RS/6000 Intel HP Server	TCP/IP	SunOS AIX Linux HPUX Windows Server 2012 and above	5 mb of HD space 32mb RAM
Micro Focus	SunSparc RS/6000 Intel HP Server	TCP/IP	SunOS AIX Linux HPUX Windows Server 2012 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 Server 2012 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 Server 2012 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 Server 2012 and above (32/64bit)	Windows Server 2012 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 or Itanium 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	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
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

64 bit Considerations

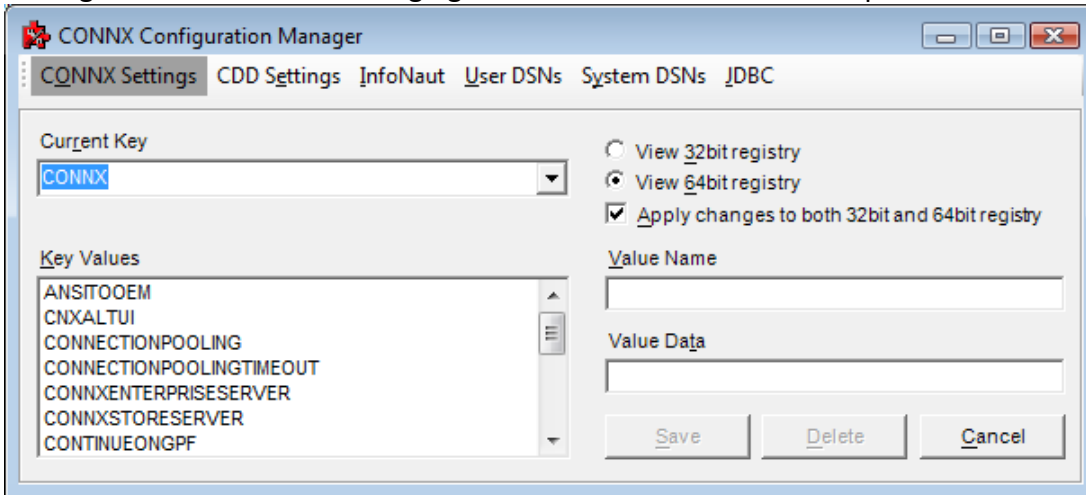
When ConnX is installed on a 64 bit Windows operating system, both the 32 bit and the 64 bit components of ConnX are installed. By default, the 32 bit components are installed in C:\CONNX32 and the 64 bit components are installed in C:\Program Files\CONNX. Both 32 bit and 64 bit executables are accessible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

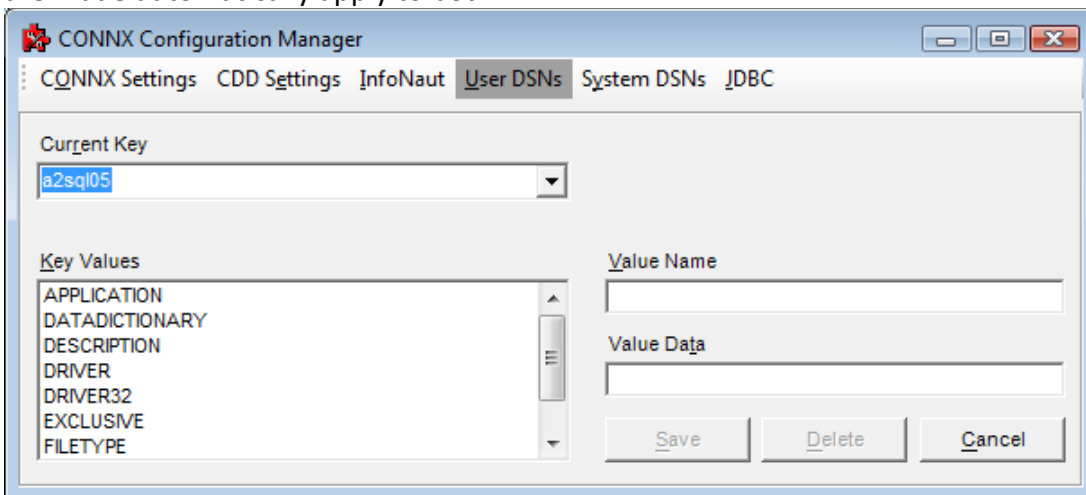
Configuring 32 bit and 64 bit components

ConnX is configured with the ConnX Configuration Manager. 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 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 support for Oracle as an InstantdbSync source database
- Added support for Docker Containers
- added Linux support for dataflex huffman compression table
- added support for Teradata
- added support for MySQL as an InstantdbSync source database
- added support for Redshift

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.

14.0:

- The ConnX function CastAsCONNXType has been deprecated. Please use CNXRawConvert instead. CastAsCONNXType does not support expressions and makes assumptions about having the raw pre-converted column data available. CNXRawConvert does not have these limitations.
- Turned on the CLIENT_FOUND_ROWS flag 100% for the MySQL driver - this will have the effect of always returning the number of rows "matched" during an update - vs the number of rows actually changed. In other words, if I update a row to a value that matched the original values, MySQL by default says 0 rows updated because the data did not change, even though 1 or more rows matched. This behavior causes problems with InstantDBSync which relies on an accurate reporting of the number of rows updated in order to detect potentially serious error conditions.
- The following CDD data type names:
 - "JMS Integer Text"
 - "JMS Big Integer Text"
 - "JMS Double"
 - "JMS Numeric"
 - "JMS DateTime (yyyy-mm-ddThh:mm:ss.fffffffff)"
 - "JMS Time from SQL Server Time2 (hh:mm:ss.fffffffff)"
 - "JMS Time from Timestamp (hh:mm:ss.fffffffff)"
 - "JMS Byte"
 - "JMS Byte Bit (1/0)"
 - "JMS Date (YYYY-MM-DD)"
 - "JMS Binary"Have had their CDD data type names changed to:
 - "XML Integer Text"
 - "XML Big Integer Text"
 - "XML Double"

- "XML Numeric"
- "XML DateTime (yyyy-mm-ddThh:mm:ss.fffffff)"
- "XML Time from SQL Server Time2 (hh:mm:ss.fffffff)"
- "XML Time from Timestamp (hh:mm:ss.fffffff)"
- "XML Byte"
- "XML Byte Bit (1/0)"
- "XML Date (YYYY-MM-DD)"
- "XML Binary"
- Performance statistics are turned on by default for replication

13.8 and earlier:

- Windows XP, Windows 7 and Windows Server 2008 are no longer supported for ConnX 13.8 and above.
- SCO Unix is no longer a supported platform
- RM Cobol is no longer supported
- 32bit support for DISAM, CISAM and Microfocus have been dropped. 64bit versions of these data servers are still supported.
- When connecting to VSE with the Barnard TCP/IP stack, you need **BSTTENVR.OBJ** from **Build 257 pre32** or newer. Please contact Barnard Software, Inc. for a copy of this file.
- ConnX is no longer supported on Solaris 5.7. Solaris 5.8 and above is required to install and run ConnX 13.5 and above.
- InstantdbSync/Open Systems Event replicator now uses sqlregistry64 rather than sqlregistry32 on Linux/Unix systems. The installation process will automatically migrate settings from the 32bit registry database to the 64bit registry database during the upgrade. When adding or updating replication settings, please use sqlregistry64. This change effects version 13.5 and above.
- 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 versions of InstantdbSync or Open Systems Event Replicator previous to 12.5, 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.

Software AG

- 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 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.
- DECNet (Pathworks) is no longer supported when connecting to 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 14.0

New Features

- Added support to LUW Replication for webMethods Integration Server as a target
- Added Docker support to LUW Replication
- Added Ability to use JDBC Updatable rowsets in conjunction with "select .. for update" clause
- Added CNXROWNUMBER support for CONNX for DBMS
- Added read lock support to CONNX for RMS
- Added support of setObject and getObject JDBC statement interface
- Added support to the database connection button in the CDD Manager to have the ability to connect to TLS/SSL data servers
- Import CDD from command line using TLS/SSL
- Added ability to create log file when using CDD administrator in command line mode
- Added auditing capabilities
- Added ability for Flushopenfilecache to work on a specific RMS table file
- Added new Data Type called "Timestamp19 (YYYYMMDDTHHMMSSFFFZ)" which is stored as a text string, but becomes an ODBC timestamp when we read it.
- Added REGEX_LIKE feature
- Added Snowflake support to CONNX.

These are the major issues that were resolved and features that were added in ConnX 14.0:

ConnX client/server

- Fix for memory overwrite when processing an unknown structure area in an RMS file.
- Fixed .NET Data Provider exception when issuing an SQL Statement with a LIMIT
- Turned on the CLIENT_FOUND_ROWS flag 100% for the MySQL driver - this will have the effect of always returning the number of rows "matched" during an update - vs the number of rows actually changed. In other words, if I update a row to a value that matched the original values, MySQL by default says 0 rows updated because the data did not change, even though 1 or more rows matched. This behavior causes problems with InstantDBSync which relies on an accurate reporting of the number of rows updated in order to detect potentially serious error conditions.
- Added setting CONNX\ERRORONTRUNCATE – When doing an INSERT SELECT, if a string column is too wide to fit in the target column the string is truncated to fit and then moved to the target. When ERRORONTRUNCATE is set to 1 and there is a truncation of data detected, it will throw an error and stop the entire Insert Select. When the value is set to 0, the data is truncated and processing continues. The default value is 0 which is the behavior of previous versions.
- Fixed parsing error with Oracle floating values larger than or equal to 1e125 and smaller than or equal 1e-125
- Changed CNXStop so that it logs every step to the system console
- Fixed crash in GROUP BY logic
- Fixed crash in MIN/MAX functions
- Fix for ACE1745 - RSP 201 AdaSCR password was not being set properly after recovery from Transaction timeout (nucleus error 9)
- Created a new aggregate function call GROUP_CONTACT which will concatenate into a single string all of the items in the group - up to 1024 length string

- Added a new optional parameter to {fn setfilename} called refresh, which will refresh the indexes for the newly attached file - this is only done in memory, and the CDD is not affected.
- Fixed memory overwrite when DB2 tracing is enabled and the DB2 table name is 128 characters (max length)
- Added new feature that will only allow connections from CDDs where the CDD password matches this new setting - CDDPASSWORD - this is a data server setting.
- Changed behavior of Regex to match that of Oracle - partial matches are allowed now.
- CONNX now returns an error with a SQL statement that contains a mismatched starting comment with no ending comment (/* with no */)
- Fixed memory overwrite in the aggregate function logic
- Fixed iTrac ACE-1760 "SELECT COUNT(*) returns invalid results"
- Fix for RMS empty file with blockio
- Fixed iTrac CXA-59 PostgreSQL - COUNT function returns error when used in PostgreSQL 10.1 or newer
- Replaced VMS command LOGOFF (which is not the correct command) with LOGOUT
- Fixed iTrac CXA-47 - Rounding problem with DECIMAL data type on negative numbers
- Fixed problem where creating a table with a period "." in the name resulted in an error
- Fixed iTrac CXA-56 Problem with some aggregate functions with unicode data type caused the length to be cut in half with Access 2013
- Fix for Support Incident: 5364280 Query uses NU superdescriptor to look for zero.
- Fix for Support Incident: 5369451 Duplicate rows in query result in complex superde.
- Fix for problem where CONNX did not recognize Solaris hashed passwords with a non-default number of rounds
- Fixed issue so query plans (where the priority count is zero in all tables) do not differ per platform
- Fix problem with index information not being calculated properly for some MU/PE tables in adabas.

- Added Snowflake support to CONNX.
- Added support for DEBUG logical for VMS
- Fixed crash where MU or PE count > 192 would result in a crash if the record buffer was too small
- Fix for Acucobol - SQLUnbind causes future fetches to fail.
- Fixed truncate command when there is a SQL view clause on the table
- Fixed crash huge sql statement - put in a guard to prevent stack overflow if nesting depth exceeds 5000
- Fixed iTrac ACE1791: Audit logging resulted in nucleus response code 47 when selecting MU subtable.
- Fixed problem with TableCache logic when {fn setfilename } was used - tables were never reused from the cache
- Fix usage of ineligible complex superdescriptor
- Added row number support for Codasyl dbms

InfoNaut

- Fixed iTrac ACE-1754 - In InfoNaut when your SQL ends with a comment marked with double hyphen, clicking the Execution Plan has no effect. The query gets performed and not just analyzed.
- Add save results in 3 new formats - Excel 2007 and newer, Semicolon separated values and custom separated values.
- Added feature so that when a query is saved to an infonaut query file (.ciq) the formatting is kept when they are opened.

CDD Manager

- Fixed bug in import of Excel files in to the CDD - skips over import of stored procedures since these are not supported.
- Fix for CDDPassword length 10 or longer not working properly
- CDDPASSWORD can now be a comma separated list with support for up to 10 passwords
- Fixed issue where CONNX did not import Oracle indexes properly if the index contained more than one column and one of the columns was descending.

CDD Comparison Tool

- Added ability to check consistency of Adabas rotated tables (MUs and PEs)
- Fixed iTrac Issue CXA68: Octet length error when comparing 2 CDD's

InstantdbSync/Open Systems Event Replicator

- Fixed crash in Oracle EP when Oracle unique index was greater than 1700 bytes
- Changed the Replication Admin so that instead of opening the Event_Replication license and holding it open while the Replication Admin was active, the license is checked and then released.
- Fixed reliability problems when starting the SQL Server EP
- Fixed issue where sending back-to-back on-demand initial states could cause the replication to stay in an “initial state pending” state forever.
- Fixed several errors in Oracle connect/reconnect logic
- Fixed problem where A2A re-deploy loses initial state start and end timestamps.
- Fixed issue where only the first error was being reported when there were multiple errors
- Fixed issue where TRUNCATE was not being replicated when MySQL was the source
- Fixed issue where recmd was not displaying error messages on the console.
- Enhanced repcmd status output to reflect all the information displayed in the replication admin status screen
- Performance improvement when using view clause support with replication
- Fixed performance issue in replication - statements were being re-prepared every transaction
- Fixed message queue performance problems when running on HPUX
- Fixed problem where the source column expression syntax check did not use double quotes around the full table name
- Fixed missing double quotes around the full table name in the create table statement

- Fixed iTrac CXB-76 Duplicate replications with no target table added after dropping target tables.
- Fixed iTrac CXB-78 A2A replications are unchecked in the Design screen when they are undeployed.
- Fixed problem where the ACD replications for Adabas MUPE tables did not pass validation even though they had a correct unique index.
- Fixed iTrac ARN-77: Replication admin does not prevent conflicting primary key mapping between parent file and MU/PE subtables
- Fixed problem where some replications in the Replication Admin were displaying a status of "Replicating" when they should have displayed as "Replication Paused."
- Fixed issue when target table is Oracle and has multiple CLOBS - empty clob values would generate an error.
- Fixed issue where a reimport of a source Adabas flat table with fewer columns resulted in the SQL statements not being correct after being opened in the rep admin.
- Fixed replication error message when it fails diagnostic check.
- Fixed iTrac ARN-103: Invalid column mapping on replication.
- Added the ability for the Replication Admin to create many target tables from one source table - target database combination.
- Performance statistics are turned on by default for replication
-

DataSync

- Fixed issue where CONNXStore column was sometimes truncated when synchronizing to a destination database
- Fix for Support Incident 5368935: While extracting CLOBS via CONNX multiple duplicate records found. The problem was that the CONNX Driver loaded DataSync configuration settings into global memory on DLL startup - and when the setting was changed in the DataSync GUI, the CONNX Core driver still had the old setting
- Fixed problems with Oracle CLOBs
- Added better error message in situations where incremental syncs fail

Excel Add-In

- Fixed iTrac Issue ACE1773: Excel plug-in validation error occurred. The Excel add in was blocking some select statements that were valid.
- Fixed iTrac ACE-1776: Excel Addin does not write back to sample Vehicle and Employees Adabas files when the file is imported with SUPERDESCRIPTORASFIELD turned on.
- Fixed a bug in the alignment of the text of the data task pane buttons.
- Fixed issue where if certain 3rd party excel adds were installed with the CONNX excel add in all buttons for both ribbons would show in the same ribbon
- Changed Write Changes button to say Write Changes to File Only to avoid confusion that this button saves the changes to the database also.

JDBC Server

- Fixed JDBC Logging - added critical section to make it thread safe
- Increased JAVA RPC size from 4k to 100k for performance
- Changed the JDBC Client TCP/IP Send logic so it sends the entire buffer all at once instead of breaking it up into 4k chunks.
- Fix the Java client so that we properly detect the first word of the SQL statement - skipping past any comments or white space
- Fixed iTrac CXA-6: Unix/Linux machines with JDBC server with bigint data type lost a last negative digit for big numbers
- Fix memory overwrite/crash when logging is enabled and SQL statement is greater than 65535.
- Implemented PreparedStatement.setObject(-, -)

License Server/Licensing

- Changed configuration so that license server picks up configuration settings from the same location as the JDBC server.
- Fixed problem with Revoke service - it now revokes by serial number instead of license code, since the license code may have changed.
- Improved error message when there are no supported databases listed in the registry. The error "You are not licensed for any databases." has been changed to "You are not configured for any databases in the CONNX Configuration Manager under the key CONNX/DATABASES." The Linux/Unix version is worded to reflect the use of sqlregistry.
- Corrected LICENSEPORT to LICENSESERVERPORT in help messages in the installclient script as well as error messages in the license server

Install

- Updated to latest SSH client library to resolve authentication issues.
- Fixed iTrac CXB-73 ADAREX value incorrect after install
- Corrected LICENSEPORT to LICENSESERVERPORT in help messages in the installclient script as well as error messages in the license server
- Provide users with a choice between su and sudo on re-install for Linux/Unix
- Fixed file path naming convention for Excel Add-In installer
- Fixed KeyboardInteractiveAuthentication error for scp and sftp

Known Issues in ConnX 14.0

- When creating new LUW Replications with JMS as a target, the index is not created correctly. This issue only affects JMS targets. If you are doing LUW replication to JMS, please contact technical support for build 19277 or higher.

Changes/Bug Fixes for ConnX 13.8

New Features

- Added support Oracle as an InstantdbSync source.
- Added support for the JDBC Server, License Server, Adabas data server, CISAM data server and DISAM data server to run in Docker Containers
- Added support for SSL/TLS on all TCP/IP connections

These are the major issues that were resolved and features that were added in ConnX 13.8:

ConnX client/server

- Added support for the JDBC Server, License Server, Adabas data server, CISAM data server and DISAM data server to run in Docker Containers
- Added support for SSL/TLS on all TCP/IP connections
- SCO Unix is no longer supported
- Added support for Oracle Timestamp with Timezone data type (previous version did not support timezone)
- Fixed issue where we did not properly detect we were communicating to the mainframe when WCP(Entire network) was used.
- Added logic so if listen socket fails then we reissue the listen.
- Fix crash in Select
cnxclientcodepage(),cnxclienticucodepage(),cnxclientdgcpage(),cnxclientdefaultcodepage() when code page was not supported or zero
- Implemented a specialized bulk update feature for Adabas
- Memory leak fixes
- Fixed problem where pooled connections that became invalidated were being held in the pool until the timeout period expired. Invalid connections are now removed immediately
- Add RMS Wait support to RMS RFA

- Added a new ODBC/JDBC optional connection parameter called SERVERPORT which will enable the user to override the port of the data server from a JDBC/ODBC connection
- Added code to put a pointer to the user id in the Adabas control block for user exits
- Added logic so that string manipulation collation adjusts based on code page for windows (this already happened on Linux/Unix based on LANG= setting)
- Fixed intermittent memory overwrites when running VSE data servers (Adabas or VSAM)
- Changed SYSCNXCOLUMNS to expose the physical column length

InfoNaut

- Fixed error in Show Tables function where a double click on the white area caused an error message

CDD Manager

- Prevent message popups when the CDD Manager is invoked using the command line interface.
- Fixed problem in Adabas DDM import code that caused a crash when there were more than 926 fields defined for a file.
- Fixed problem with vms browse function not working when installed to a location with a space in the name
- Fixed problem with Adabas SYSOBJH imports where the import skipped items when the group level was greater than 6
- Fixed crash if Delete Restriction is attempted when no restrictions are present (Table Security tab)

InstantdbSync/Open Systems Event Replicator

- Added support for Oracle as an InstantdbSync source.
- Corrected problems arising from truncating the source table
- Corrected crashes when running under heavy load with SQL Server as the source.
- Corrected problems with updates not being replicated when running under heavy load with SQL Server as the source
- Resolved memory overwrite errors in SQL Server EP.
- Corrected memory leak which caused a crash during initial state when replicating to an Oracle target and the source table had several billion records.

DataSync

- Fixed intermittent problem where syncs would report a blank table name
- Fixed problem where Incremental syncs failed if there were changes to a table and the source table name had a space in it
- Fixed problem in transform wizard where it was rewriting sql and making joins in the from clause invalid
- Fixed slow performance when running scheduled syncs
- Fixed intermittent problem where scheduled syncs could lock up
- Fixed a problem with license counts
- Fixed issue where error condition still triggered on success task
- Fixed issue where the target schema was not being created in a transform.
- Fixed problem where a sync would fail when column names contained single quotes.
- Fixed problem where if a scheduled sync failed because of a crash, it did not properly detect it and trigger the failure task.

Excel Add-In

- Corrected update error when a column name had spaces in it
- Corrected problem where the add-in would no longer access data if any changes were made in the connection
- Corrected display problems when running Excel on Windows 10 with certain laptop displays
- Corrected a problem when editing a row in an Adabas table that contains null column. The commit caused an incorrect “no unique row” error even though the row was unique

RCI

- Fix for issue where stored procedures that return result sets were not always reporting the correct column count in the RCI interface.

JDBC Server

- Fix for java executeQuery - if the execute failed the whole statement handle was being closed
- Show warning message in the DSN Registry tool when a new DSN exceeds the 50-character limit and provide an opportunity to correct the >50 characters error

License Server/Licensing

- Adding Hyper threading compensation for SQL Server core checks
- Fix for core count logic if there are multiple connections in the CDD pointing to the same server.

Configuration Manager

- Fixed issue where it was reading/writing from 32bit when 64 was selected and 64bit when 32bit was selected. This affected the CONNX key only

Install

- Fixed missing or non-working shortcut keys in UNIX Client Setup, ConnX Server Setup, License Admin, and Adabas SQL Gtwy Emb SQL Setup
- Fixed issue where references uninstalled items were being left in the menu
- Fixed issue in the Windows logon validation logic which failed if the password contained '&' or '#'
- Corrected problem where the file association for .CDD was incorrect if ConnX was installed in a location other than the default location

Known Issues in ConnX 13.8

- The ConnX client on zLinux does not include the Oracle OCI library. If upgrading from a previous version. Please either contact technical support to obtain an updated build that contains the OCI library or download the Oracle Instant Client directly from the Oracle web site.
- When using MySQL as a target database for InstantdbSync, updates will fail if the target table contains a Unicode string that is longer than 500 bytes. Please contact technical support for an updated build if you encounter this problem.
- The documentation at the time of publication states the minimum supported Windows version to be Windows 7. ConnX version 13.8 has dropped support for Windows 7 and Windows Server 2008. The minimum supported Windows version is Windows Server 2012.

Changes/Bug Fixes for ConnX 13.5

New Features

- Support for Redshift
- Enhanced Replication Administrator to display performance statistics on status screen
- Added command line options to replication controller
- SSL support added for connections between client and server code

These are the major issues that were resolved and features that were added in ConnX 13.5:

ConnX client/server

- Added SSL support
- Added support for Redshift
- Fixed problem where mixed case FDTs on mainframe Adabas were being read as upper case only
- Security fixes
- Fix to problem where connections to the database were not closed properly when license check failed
- Fixed NULLABLE attribute for all IMS tables
- Fix for givesocket/takesocket support on VSE
- Fix for being unable to drop an index using index name on some databases
- Corrected user fault alignments with ConnX for RMS on Itanium
- Fixed linked server problem when table names were not fully qualified

CDD Manager

- Added ability to import VSAM copybooks from local PC
- Fixed Adabas import issue where MU/PE occurrences were being prompted for when ADA_TABLENAME option was set to 1
- Fixed problem where the toolbar was not displayed properly during application startup

InstantdbSync/Open Systems Event Replicator

- Fixed issue where replication admin did not prevent conflicting primary key mapping between parent file and MU/PE subtables.
- Fixed memory leak with Adabas to Adabas replication
- Fixed issues with truncate command when source is SQL Server
- All components are now 64bit
- Added command line functions for Pause/Resume replication, Start Initial State and Get Status
- Added performance statistics to replication admin status screen

InfoNaut

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

DataSync

- Fixed "Unspecified Error" problem with Datasync. The correct diagnostics are now returned
- Fixed GUI issues with tab order and default buttons
- Corrected errors that occurred when Postgres was a database in the DataSync CDD

Install

- Added manual copy option to z/OS server installer
- 32bit compatibility libraries are no longer necessary when installing on a 64bit Linux system
- Added a check to ensure that the main ConnX installer had been run prior to installing the Unix client

Known Issues

- At the time of release, InstantdbSync 13.5 has known issues when replicating tables that have clustered indexes in high volume environments. Please contact technical support for an updated build to correct this issue.
- The default value for the interval at which the System Status screen in the Replication Administrator for InstantdbSync refreshes (STATUSREFRESHINTERVAL) is incorrectly being set to 1024 seconds. To return this value to the correct default value, please set STATUSREFRESHINTERVAL to 30 in the ConnX Configuration Manager under the CONNX\Replication key. This issue is corrected in later patch builds.

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 PostgreSQL
- 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.

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 and above 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 or above from a version prior to 12.5, 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.