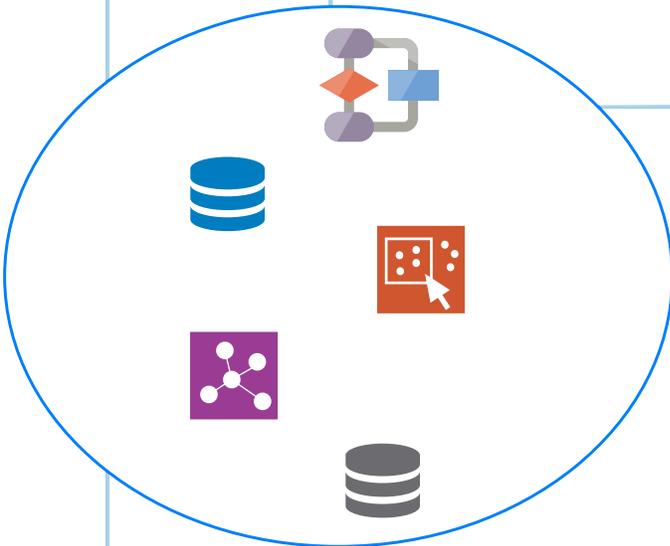
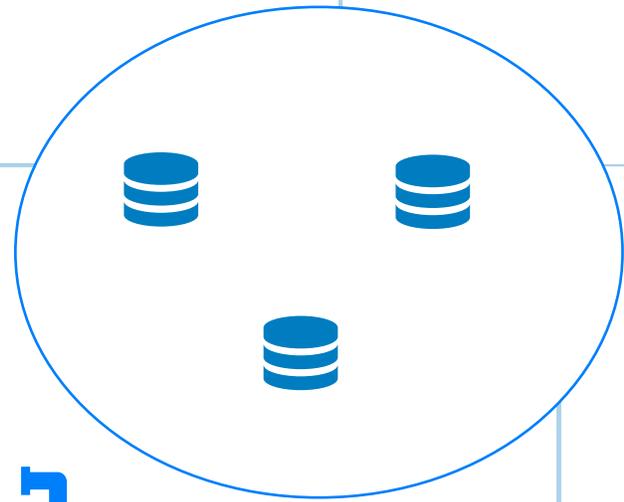


# infogix® **Data3Sixty**®

Release Notes  
version 2.1.2

01



---

## Chapter 1 What's New in 2.1.2?

Analysis Changes .....	5
Group Node .....	5
Repartition Node .....	5
Allow Filter Specification in some Data Store Input types .....	5
Add Support for Accurate Output Counts when Caching is Off	5
Add New Field for Master Instance Type in Dedicated Sizing .	5
Configuration for Value Limit in the Profile Node .....	5
Allow Decryption of Encrypted Values when Exporting to External Data Stores .....	6
Avoid Unnecessary Caching in Analyses .....	6
Automatically Increase Spark Driver Memory when Number of Files Increases .....	6
Security Enhancements.....	6
Support External Admin Groups .....	6
Support for re-encrypting of DEK when KEK changes .....	6
Allow Groups with No Members .....	6
Enforce Infogix S3 Bucket Policies .....	7
Data Store and Data View Improvements .....	7
Support Impala as Data View Store.....	7
Support Hive as Data View Store .....	7
Add Support for Configurable Marker File in HDFS.....	7
Allow user to specify part of Data View Table Name .....	7
Decimal Precision in Search Screen Results for In-line Editing and in View/Edit Screens.....	7
Add Execution Sizing and Execution Properties to Data Views	8
Case Management Improvements .....	8
Decrease Lower Limits of Column Width and Height to Save Space .....	8
Executions History Improvements .....	8
Make Work Id, Process Id, and Stage Name Copyable .....	8
Import/Export Improvements .....	8
Improve Error Message when Importing a File that is not a Definition Export File .....	8
On-premises Installation and Infrastructure Improvements .....	8
Add a standby PostgreSQL Database Instance .....	8
Support for full Vertica Backups.....	9
Allow Customer to add their own Certificates to App Server Key Store .....	9



---

Support Apache Tomcat version 8.0.47 and 9.0.4 .....	9
Add Support for HDFS and Yarn client configs in Installer .....	9
Ability to pass additional JVM arguments to Apache Tomcat ..	9
Improved restore command to fully support Data Migration ....	9
Wait for Process to start or stop rather than using sleep().....	9
Upgrade PostgreSQL from 9.5 to 10.2 .....	9
System Property to turn off Execution of Scheduled Stages .	10

## **Chapter 2    Technical Preview Features**

Infogix Data3Sixty Govern Integration .....	11
Additional Functions .....	11
Sqoop In Node.....	11
Sqoop Out Node .....	11
Kinesis Node.....	12

## **Chapter 3    Customer Reported Issues**

Unable to Delete Individual Data View Data Snapshots .....	13
Once a Dimension is added to a Data View and Saved, it can't be Deleted .....	13

Infogix, the Infogix checkmark logo, Infogix Assure, Infogix Insight, ACR, ACR/Detail, ACR/Summary, ACR/Workbench, ACR/Connector, ACR/Instream, ACR/File, Infogix ER, Infogix DTM, Infogix Perceive, Infogix Nexix, Infogix Data3Sixty, Infogix Sagacity, the Infogix ART logo, DataMind, ClickMind, HealthMind, RevMind, RouteMind, TrafficMind, Information Integrity, Turning Reform into Revenue, and Information With Confidence are registered trademarks of Infogix, Inc.

Infogix Controls and The Information Integrity Experts are trademarks of Infogix, Inc.

Any other trademarks or registered trademarks are the property of their respective owners.

Copyright 2018 Infogix, Inc. All rights reserved.

Confidential—Limited distribution to authorized persons only, pursuant to the terms of Infogix, Inc. (FKA Unitech Systems, Inc.) license agreement. This document is protected as an unpublished work and constitutes a trade secret of Infogix, Inc.

This document and the information contained herein are the property of Infogix, Inc. Reproduction or use in whole or in part of this document and the information contained herein by anyone without prior written consent of Infogix, Inc. is prohibited.

Publication Date 5/1/2018

# What's New in 2.1.2?

The *Infogix Data3Sixty Release Notes* briefly describes enhancements and modifications included in Infogix Data3Sixty 2.1.2. For more detailed and functional descriptions of newly added features, please consult the Infogix Data3Sixty User Guide.

## Analysis Changes

### Group Node

The newly added Group node allows users to group records by one or multiple fields, in order to produce an associated array. The arrays that are associated to each group can then be processed using other Analysis features, such as array functions or the Javascript node.

### Repartition Node

Repartitioning of data may now be performed within Analyses, using the newly added Repartition node. Along with the Cache Data node, the Repartition node can now be found under the newly added 'System' tab, which has been added to the Analysis Designer palette.

### Allow Filter Specification in some Data Store Input types

A new Filter Records tab allows users to create a filter expression for JSON, ORC, or Parquet file based Data Stores at the Data Store Input node level. Filter expressions can be created using the general expression language used throughout Infogix Data3Sixty or SQL. Filtering may be applied at the record level or the column level.

### Add Support for Accurate Output Counts when Caching is Off

A new check box has been added to Analysis settings, to control whether accurate output counts should be kept for each node in the Analysis' Executions History.

When checked, this check box will cause the system's Execution History to track record counts at each node within an Analysis. When unchecked, the system will not track record counts.

By default, Collect Accurate Record Counts is checked and disabled. To enable the check box, you will need to disable the Cache Output of Nodes check box. Note that Collect Accurate Record Counts must be checked when Cache Output of Nodes is checked and that the UI will enforce this.

### Add New Field for Master Instance Type in Dedicated Sizing

While previously a user could control the instance type used during the execution of an Analysis within Analysis settings, there was no option to control the master instance type. With this release, this functionality has been added.

### Configuration for Value Limit in the Profile Node

Within the Profile node, a new Value Limit parameter can be used to control how many values can be displayed in the Profile field's *values* and *patterns* array fields.

## ■ What's New in 2.1.2?

---

For the *values* field, if more values are found than the Value Limit, no values will be displayed. For the *patterns* field, the Value limit acts as a maximum. That is, no matter how many patterns exist, only up to x patterns will be displayed, where x is the Value Limit. By default, Value Limit is set to 10,000; this is the value that will be used if Value Limit is left blank.

### **Allow Decryption of Encrypted Values when Exporting to External Data Stores**

System Administrators will now have the option to decrypt values that have been encrypted by the system, within certain Analysis nodes. To support this, the following Analysis nodes now feature a *Decrypt Secure Values* check box: Data Store Output (External Store Types only) and Sqoop Out.

### **Avoid Unnecessary Caching in Analyses**

This enhancement allows the system to more intelligently decide which nodes to cache and which ones not to cache, thereby preventing it from caching nodes that do not need to be cached.

### **Automatically Increase Spark Driver Memory when Number of Files Increases**

As the number of files for a Data Store starts to increase, Spark requires that the amount of memory allocated to the driver be increased to void OOM errors. This enhancement allows the system to automatically perform this increase as part of the automatic sizing algorithm, which is in effect when the user does not override sizing in an Analysis.

## **Security Enhancements**

### **Support External Admin Groups**

Additional functionality has been added to the Mapping Enterprise SSO feature for User Groups. This new functionality allows external groups to be mapped to groups within the product that will cause all Users within the Group to be treated as system wide administrators.

### **Support for re-encrypting of DEK when KEK changes**

In order to support rotation of Key Encrypting Keys (KEK), a command line tool to re-encrypt all Data Encryption Keys (DEK) with a new KEK is now provided.

### **Allow Groups with No Members**

A Group with 0 Users may now be created and saved. Doing so can facilitate subsequent external group mapping for Enterprise SSO users. Previously, a Group needed at least 1 member to be saved.

### **Enforce Infogix S3 Bucket Policies**

Changes have been made to increase security measures that prevent unauthorized access to S3 buckets associated with other tenants.

Specifically:

- Within a deployment, External Data Stores may not access data within the storage used for Internal Data Stores to store the data for all customers.
- External Data Stores may not access S3 buckets that end in *sagacity.infogix.com*, unless the bucket starts with the Deployment Id of the tenant. For example, tenant ABC can access bucket *abc123.sagacity.infogix.com* but cannot access bucket *xyz.sagacity.infogix.com*.

### **Data Store and Data View Improvements**

#### **Support Impala as Data View Store**

On Enterprise Installations that use a Cloudera Hadoop cluster, Apache Impala is now supported for Regular Performance Data Views.

#### **Support Hive as Data View Store**

On Enterprise Installations that use a Hortonworks Hadoop cluster, Apache Hive is now supported for Regular Performance Data Views.

#### **Add Support for Configurable Marker File in HDFS**

This setting has been supported for the Cloud version of the product for some time, but with this release it has been added to the Enterprise version. Within Data Stores that use the HDFS file system, it can be used to prevent writing to files that should not be written to unless a Marker is found.

#### **Allow user to specify part of Data View Table Name**

Within Data Views, a new, optional parameter called Table Prefix Name allows users to create a prefix of up to 30 characters that will be added to the Data View's underlying, unique name. Doing so can then allow users to search for the Data View by name, when querying the database that retains Data View information, with an external tool. Without a Table Name Prefix, a Data View's name will instead just be a non-meaningful, alphanumeric string. Adding a Table Name Prefix can make searching for the Data View easier.

#### **Decimal Precision in Search Screen Results for In-line Editing and in View/Edit Screens**

The Decimal Precision that is displayed when performing in-line editing on numeric fields within Case Management or when using a View/Edit screen has been improved, as follows:

## ■ What's New in 2.1.2?

---

- Decimal fields now respect the Scale setting specified for the field or will have a decimal precision of 4 if no Scale is specified.
- Currency fields now respect the Scale setting specified for the field or will have a decimal precision of 2 if no Scale is specified.
- Floating point fields will have a decimal precision of 4.

### **Add Execution Sizing and Execution Properties to Data Views**

A new tab has been added to Data Views, available only to system wide administrators and users with Administer privileges, that allows users to define the Data View's execution sizing parameters and execution properties. The parameters found in this tab are similar to those found in Analysis Settings.

## **Case Management Improvements**

### **Decrease Lower Limits of Column Width and Height to Save Space**

It was reported by a customer that a minimum height of 70 in column definitions for Case Management screens was too large and resulting in wasted space. In response, the minimum height for columns was reduced to 40 and the minimum width was reduced to 200.

## **Executions History Improvements**

### **Make Work Id, Process Id, and Stage Name Copyable**

This enhancement makes the Stage Name, Work Id, and Process Id fields shown in the Executions History screen and the Work Id shown in the View Content screen copyable. This capability was added because users often need these field's values for use in other parts of the product.

## **Import/Export Improvements**

### **Improve Error Message when Importing a File that is not a Definition Export File**

This enhancement has changed the error message that is shown in this scenario from "Unexpected Error" to "The file you uploaded was not a valid definition export file."

## **On-premises Installation and Infrastructure Improvements**

### **Add a standby PostgreSQL Database Instance**

On premises installations may now be configured to have a warm standby database instance that will take over operations if the primary database instance fails.

### **Support for full Vertica Backups**

Previously, the product only supported incremental Vertica backups. With this release, full backups are now supported and customers may schedule them.

### **Allow Customer to add their own Certificates to App Server Key Store**

Customers may now load their own Security Certificates into the App Server's Key Store. Doing so can allow the customer to make an SSL connection via an Analysis Plugin node, if such a connection is required.

### **Support Apache Tomcat version 8.0.47 and 9.0.4**

In order to be able to deploy the application to the Ireland AWS data center, to support European clients, the Cloud Installer code base has been updated to support Apache Tomcat version 8.0.47. Additionally, the Enterprise Installer code base has been updated to support Apache Tomcat version 9.0.4.

### **Add Support for HDFS and Yarn client configs in Installer**

To support a High Availability Hadoop configuration, the product installer can now merge a Yarn config and an HDFS config into the client config jar.

### **Ability to pass additional JVM arguments to Apache Tomcat**

In order to support Cloudera integration requirements, the ability to pass additional JVM arguments to Apache Tomcat has been added to the Enterprise version of the product. Passage of these arguments is performed through use of a new property in the *common.properties* file.

### **Improved restore command to fully support Data Migration**

In this release, the *restore* command has been improved to support data migration from one database to another. To do so, the *restore* command can now detect if a database schema already exists; if it does, the user will now have the option to recreate it.

### **Wait for Process to start or stop rather than using sleep()**

In this release, the *startApplicationDb* and *startStandbyApplicationDb* tasks have been improved by the removal of *sleep()* statements.

### **Upgrade PostgreSQL from 9.5 to 10.2**

To support the ability to add a standby instance of PostgreSQL, the product's version of PostgreSQL has been upgraded from 9.5 to 10.2. Additionally, to support this upgrade, the system will now validate that the correct version of PostgreSQL data (10.2, not 9.5) is being used.

## ■ What's New in 2.1.2?

---

### **System Property to turn off Execution of Scheduled Stages**

A system property to support turning off execution of scheduled stages for a given installation has been added. This capability allows users to turn off scheduled executions in lower level environments but at the same time be able to define schedules in the lower level environments, so that definitions don't need to be edited in higher level environments.

# Technical Preview Features

The following features are available within Infogix Data3Sixty as “Tech Previews”. A feature classified as Tech Preview is not guaranteed to be 100% functional but has been included in this release to enable customer exploration. Infogix encourages you to explore these “Tech Preview” features and provide feedback on your experiences. All “Tech Preview” functionality is marked as such within the Infogix Data3Sixty UI.

## Infogix Data3Sixty Govern Integration

Infogix Data3Sixty can be configured per Environment to push metadata about its Data Stages to Infogix Data3Sixty Govern. The settings that make this possible are available within the *Admin->Environments* option in the product’s Main Menu. Within Environments, the configuration takes place under the *Govern* tab.

Once a connection is established, the product will then push metadata to Govern whenever a Data Stage is created, updated, or deleted. A Metadata Synchronization feature may also be used on demand or scheduled.

## Additional Functions

The following functions have been added for use in the Analysis Designer. Detailed descriptions of each function can be found within the product itself.

- ISMATCHREGEX
- ISBLANK
- ISEMPY
- ISNUMBER
- ISDATE
- IDENTIFYDATATYPE
- CONTAINSREGEX
- INDEXREGEX
- PHONEMATCH
- ZIPCODEMATCH
- NAMEMATCH
- COMPANYNAMEMATCH
- ADDRESSMATCH

## Sqoop In Node

The Sqoop In node can be used to efficiently ingest data from an external, relational database. This feature utilizes Apache Sqoop.

## Sqoop Out Node

The Sqoop Out node can be used to efficiently output data to an external, relational database. This feature utilizes Apache Sqoop.

## ■ Technical Preview Features

---

### **Kinesis Node**

The Kinesis Node can be used to pull in live data streams, for usage in an Analysis.

# Customer Reported Issues

This section of the release notes lists Customer Reported Issues that were resolved in Infogix Data3Sixty 2.1.2.

## **Unable to Delete Individual Data View Data Snapshots**

It was reported that there was an issue with the feature that allows the user to delete individual runs from a Data View. When users tried to delete a record associated with a particular run, the record would not be removed from the View Content UI. This issue has been fixed.

## **Once a Dimension is added to a Data View and Saved, it can't be Deleted**

It was reported that there was an issue with Data View Dimension deletion, that prevented users from deleting Dimensions after the Data View had been saved. This issue has been fixed.