



# TRIFACTA

## Trifacta Release Notes

Version: 9.2

Doc Build Date: 07/29/2022

**Copyright © Trifacta Inc. 2022 - All Rights Reserved. CONFIDENTIAL**

These materials (the “Documentation”) are the confidential and proprietary information of Trifacta Inc. and may not be reproduced, modified, or distributed without the prior written permission of Trifacta Inc.

EXCEPT AS OTHERWISE PROVIDED IN AN EXPRESS WRITTEN AGREEMENT, TRIFACTA INC. PROVIDES THIS DOCUMENTATION AS-IS AND WITHOUT WARRANTY AND TRIFACTA INC. DISCLAIMS ALL EXPRESS AND IMPLIED WARRANTIES TO THE EXTENT PERMITTED, INCLUDING WITHOUT LIMITATION THE IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE AND UNDER NO CIRCUMSTANCES WILL TRIFACTA INC. BE LIABLE FOR ANY AMOUNT GREATER THAN ONE HUNDRED DOLLARS (\$100) BASED ON ANY USE OF THE DOCUMENTATION.

For third-party license information, please select **About Trifacta** from the Help menu.

|  |     |
|--|-----|
| 1. <i>Release Notes</i>                        | 4   |
| 1.1 <i>Changes to System Behavior</i>          | 5   |
| 1.1.1 <i>Changes to the Language</i>           | 6   |
| 1.1.2 <i>Changes to the APIs</i>               | 28  |
| 1.1.3 <i>Changes to Configuration</i>          | 50  |
| 1.1.4 <i>Changes to the Object Model</i>       | 59  |
| 1.1.5 <i>Improvements to the Type System</i>   | 64  |
| 1.1.6 <i>Changes to User Management</i>        | 72  |
| 1.2 <i>Release Notes 9.2</i>                   | 77  |
| 1.3 <i>Release Notes 8.7</i>                   | 89  |
| 1.4 <i>Release Notes 8.2</i>                   | 98  |
| 1.5 <i>Release Notes 7.6</i>                   | 116 |
| 1.6 <i>Release Notes 7.1</i>                   | 125 |
| 1.7 <i>Release Notes 6.8</i>                   | 135 |
| 1.8 <i>Release Notes 6.4</i>                   | 143 |
| 1.9 <i>End of Life and Deprecated Features</i> | 149 |

# Release Notes

This section contains release notes for published versions of Trifacta®.

# Changes to System Behavior

The following pages contain information about changes to system features, capabilities, and behaviors in this release.

# Changes to the Language

## Contents:

- Release 8.11
  - *Nest transformation explicitly types transformed column*
- Release 8.10
  - *Split by position no longer requires sorted list of positions*
- Release 8.5
  - *Support for numeric separators in NUMFORMAT function*
  - *New functions*
- Release 8.4
  - *New functions*
- Release 8.3
- Release 8.2
- Release 7.10
  - *New functions*
- Release 7.9
  - *Transform Builder now supports All columns option*
- Release 7.8
  - *Rename transform now supports Upper / Lower and Left / Right options*
- Release 7.5
  - *New Functions*
  - *New Functions*
- Release 7.3
  - *New Functions*
- Release 7.2
  - *New Functions*
- Release 7.1
  - *New Functions*
- Release 6.11
  - *New Functions*
  - *PARSEDATE function now supports four default Datetime format values*
  - *ignore case parameter added to string functions*
  - *Expanded parameters for SUBSTITUTE function*
- Release 6.9
  - *MODE functions return lowest value in evaluated set if there is a tie*
  - *New Functions*
- Release 6.8
  - *New Functions*
- Release 6.6
  - *New Functions*
- Release 6.5
  - *New Functions*
- Release 6.4
  - *Improvements to metadata references*
- Release 6.3
  - *New Functions*
  - *Optional input formats for DateFormat task*
- Release 6.2
  - *New Functions*
  - *ARRAYELEMENTAT function accepts new inputs*
- Release 6.1
- Release 6.0
  - *New Functions*
  - *Changes to LIST\* inputs*
  - *Renamed functions*
  - *FILL Function has new before and after parameters*

- *Release 5.9*
  - *New functions*
- *Release 5.8*
  - *File lineage information using source metadata references*
  - *New math and statistical functions for arrays*
- *Release 5.7*
  - *WEEKNUM function now behaves consistently across running environments*
- *Release 5.6*
  - *URLPARAMS function returns null values*
- *Release 5.1*
  - *Wrangle now supports nested expressions*
  - *SOURCEROWNUMBER function generates null values consistently*
  - *New Functions*
- *Release 5.0.1*
  - *RAND function generates true random numbers*
- *Release 5.0*
  - *Required type parameter*
  - *Deprecated aggregate transform*
  - *New search terms*
  - *Support for <> operator*
  - *ROUND function takes optional number of digits*
  - *New Functions*
- *Release 4.2.1*
- *Release 4.2*
  - *New Filter transform*
  - *New Case transform*
  - *Rename transform now supports multi-column rename*
  - *Delete specified columns or delete the others*
  - *New string comparison functions*
  - *NOW function returns 24-hour time values*
  - *New Transforms*
  - *New Functions*

---

The following changes have been applied to *Wrangle* in this release of Trifacta®.

## Release 8.11

### Nest transformation explicitly types transformed column

In prior releases, when a Nest transformation was applied to a column to nest values into Arrays or Objects, the resulting column was re-inferred by the Trifacta application. This re-inference should not be necessary, since the target column's data type is effectively declared in the transformation definition.

Beginning in this release, the output column of these Nest transformations is explicitly typed to Array or Object data type, based on the definition of the transformation.

**NOTE:** Existing uses of the Nest transformation are not immediately affected. However, if these transformations are edited, then the changes may cause unexpected results and breakages in downstream transformations. If the recipe was originally designed expecting a different data type, subsequent steps may have been used to clean up the nested data, assuming that it was String values or some other data type. If the output column is now explicitly typed as Array or Object data type, these steps may be broken. You may be able to fix these broken steps by explicitly typing the output column to String after the Nest transformation and before your subsequent steps.

For more information, see *Nest Your Data*.

## Release 8.10

### Split by position no longer requires sorted list of positions

Beginning in this release, when you create a Split by position transformation, the numeric values indicating the positions do not need to be listed in sorted numeric order.

**Tip:** You can now do faster iteration since you can add new positions as needed when previewing the transformation.

For more information, see *Split Column*.

## Release 8.5

### Support for numeric separators in NUMFORMAT function

Beginning in Release 8.5, the NUMFORMAT function supports the following configurable separators for localizing numeric values:

| Option Name        | Description   |
|--------------------|---|
| Decimal Separator  | The string used to separate the integer part of a Decimal value from its fractional part. |
| Grouping Separator | The string used to separate a group of digits.  |

For more information, see *NUMFORMAT Function*.

### New functions

| Function Name            | Description  |
|--------------------------|--|
| <i>NUMVALUE Function</i> | Converts a string formatted as a number into an Integer or Decimal value by parsing out the specified decimal and group separators. A string or a function returning formatted numbers of String type or a column containing formatted numbers of string type can be inputs. |

## Release 8.4

### New functions

Documentation for the following functions is now available.

| Function Name               | Description  |
|-----------------------------|--|
| <i>FINDNTH Function</i>     | Returns the position of the nth occurrence of a letter or pattern in the input string where a specified matching string is located in the provided column. You can search either from left or right.                   |
| <i>PARSESTRING Function</i> | Evaluates an input against the String datatype. If the input matches, the function outputs a String value. Input can be a literal, a column of values, or a function returning values. Values can be of any data type. |
| <i>PARSEARRAY Function</i>  | Evaluates a String input against the Array datatype. If the input matches, the function outputs an Array value. Input can be a literal, a column of values, or a function returning String values.                     |
| <i>PARSEOBJECT Function</i> | Evaluates a String input against the Object datatype. If the input matches, the function outputs an Object value. Input can be a literal, a column of values, or a function returning String values.                   |



## Release 8.3

None.

## Release 8.2

None.

## Release 7.10

### New functions

| Function Name              | Description   |
|----------------------------|---|
| <i>EOMONTH</i><br>Function | Returns the serial date number for the last day of the month before or after a specified number of months from a starting date. |

## Release 7.9

### Transform Builder now supports All columns option

Beginning in Release 7.9, select **All** columns option has been added in the Transform Builder.

| Option Name | Description                        |
|-------------|------------------------------------|
| All         | Selects all columns in the dataset |

### Example:

|                            |                |
|----------------------------|----------------|
| <b>Transformation Name</b> | Rename columns |
| <b>Parameter: Option</b>   | Add suffix     |
| <b>Parameter: Columns</b>  | All            |
| <b>Parameter: Suffix</b>   | _new           |

The following is the list of the transformations that accept the All option for selecting columns:

- Date format
- Delete columns
- Move columns
- Rename column
- Replace
- Replace text or patterns
- Replace cells
- Replace text between delimiters
- Replace by position
- Replace mismatched values
- Replace missing values
- Edit with formula
- Change column type
- Text format
- Unpivot columns

For more information, see *Transform Builder*.

## Release 7.8

### Rename transform now supports Upper / Lower and Left / Right options

Beginning in Release 7.8, the Rename transform supports the following new options:

| Option Name                | Description   |
|----------------------------|---|
| Convert to lowercase       | Converts existing column names to lowercase                                   |
| Convert to UPPERCASE       | Converts existing column names to uppercase                                   |
| Keep from beginning (left) | Specifies the number of characters to keep from the beginning of column names |
| Keep from end (right)      | Specifies the number of characters to keep from the end of column names       |

For more information on rename columns, see *Rename Columns*.

## Release 7.5

### New Functions

#### Approximation statistical functions:

**Tip:** Approximation functions are suitable for larger datasets. As the number of rows increases, accuracy and calculation performance improves for these functions.

| Function Name                         | Description   |
|---------------------------------------|---|
| <i>APPROXIMATEMEDIAN Function</i>     | Computes the approximate median from all row values in a column or group. Input column can be of Integer or Decimal.                        |
| <i>APPROXIMATEPERCENTILE Function</i> | Computes an approximation for a specified percentile across all row values in a column or group. Input column can be of Integer or Decimal. |
| <i>APPROXIMATEQUARTILE Function</i>   | Computes an approximation for a specified quartile across all row values in a column or group. Input column can be of Integer or Decimal.   |

#### base64 encoding functions:

| Function Name                | Description  |
|------------------------------|--|
| <i>BASE64ENCODE Function</i> | Converts an input value to base64 encoding with optional padding with an equals sign (=). Input can be of any type. Output type is String. |
| <i>BASE64DECODE Function</i> | Converts an input base64 value to text. Output type is String.   |

## Release 7.4

### New Functions

| Function Name               | Description  |
|-----------------------------|--|
| <i>WEEKDAYNAME Function</i> | Derives the full name from a Datetime value of the corresponding weekday as a String. Source value can be a reference to a column containing Datetime values or a literal. |

## Release 7.3

### New Functions

| Function Name                   | Description   |
|---------------------------------|---|
| <i>ROLLINGMAXDATE Function</i>  | Computes the rolling maximum of date values forward or backward of the current row within the specified column. Inputs must be of Datetime type.              |
| <i>ROLLINGMINDATE Function</i>  | Computes the rolling minimum of Date values forward or backward of the current row within the specified column. Inputs must be of Datetime type.              |
| <i>ROLLINGMODEDATE Function</i> | Computes the rolling mode (most common value) forward or backward of the current row within the specified column. Input values must be of Datetime data type. |

## Release 7.2

### New Functions

| Function Name                          | Description  |
|--|--|
| <i>KTHLARGESTDATE Function</i>         | Extracts the ranked Datetime value from the values in a column, where $k=1$ returns the maximum value. The value for $k$ must be between 1 and 1000, inclusive. Inputs must be valid Datetime values.                              |
| <i>KTHLARGESTUNIQUEDATE Function</i>   | Extracts the ranked unique Datetime value from the values in a column, where $k=1$ returns the maximum value. The value for $k$ must be between 1 and 1000, inclusive. Inputs must be Datetime.                                    |
| <i>KTHLARGESTDATEIF Function</i>       | Extracts the ranked Datetime value from the values in a column, where $k=1$ returns the maximum value, when a specified condition is met. The value for $k$ must be between 1 and 1000, inclusive. Inputs must be Datetime.        |
| <i>KTHLARGESTUNIQUEDATEIF Function</i> | Extracts the ranked unique Datetime value from the values in a column, where $k=1$ returns the maximum value, when a specified condition is met. The value for $k$ must be between 1 and 1000, inclusive. Inputs must be Datetime. |
| <i>MINDATEIF Function</i>              | Returns the minimum Datetime value of rows in each group that meet a specific condition. Set of values must valid Datetime values.   |
| <i>MAXDATEIF Function</i>              | Returns the maximum Datetime value of rows in each group that meet a specific condition. Set of values must valid Datetime values.   |
| <i>MODEDATEIF Function</i>             | Returns the most common Datetime value of rows in each group that meet a specific condition. Set of values must valid Datetime values.   |

## Release 7.1

### New Functions

This release introduces the following functions to calculate the difference between two valid dates.

#### Date calculation functions:

| Function Name            | Description   |
|--------------------------|---|
| <i>MINDATE Function</i>  | Computes the minimum value found in all row values in a Datetime column.              |
| <i>MAXDATE Function</i>  | Computes the maximum value found in all row values in a Datetime column.              |
| <i>MODEDATE Function</i> | Computes the most frequent (mode) value found in all row values in a Datetime column. |

#### Work day functions:

| Function Name                    | Description   |
|----------------------------------|---|
| <i>NETWORKDAYS Function</i>      | Calculates the number of working days between two specified dates, assuming Monday - Friday workweek. Optional list of holidays can be specified.   |
| <i>NETWORKDAYS INTL Function</i> | Calculates the number of working days between two specified dates. Optionally, you can specify which days of the week are working days as an input parameter. Optional list of holidays can be specified. |
| <i>WORKDAY Function</i>          | Calculates the work date that is before or after a start date, as specified by a number of days. A set of holiday dates can be optionally specified.  |
| <i>WORKDAYINTL Function</i>      | Calculates the work date that is before or after a start date, as specified by a number of days. You can also specify which days of the week are working days and a list of holidays via parameters.      |

### Time zone conversion functions:

| Function Name                    | Description  |
|----------------------------------|--|
| <i>CONVERTFROM UTC Function</i>  | Converts Datetime value to corresponding value of the specified time zone. Input can be a column of Datetime values, a literal Datetime value, or a function returning Datetime values.                        |
| <i>CONVERTTOUT C Function</i>    | Converts Datetime value in specified time zone to corresponding value in UTC time zone. Input can be a column of Datetime values, a literal Datetime value, or a function returning Datetime values.           |
| <i>CONVERTTIME ZONE Function</i> | Converts Datetime value in specified time zone to corresponding value second specified time zone. Input can be a column of Datetime values, a literal Datetime value, or a function returning Datetime values. |

## Release 6.11

### New Functions

This release introduces the following functions to evaluate String values against individual data types.

| Function Name              | Description  |
|----------------------------|--|
| <i>PARSEBOOL Function</i>  | Evaluates a String input against the Boolean datatype. If the input matches, the function outputs a Boolean value. Input can be a literal, a column of values, or a function returning String values.  |
| <i>PARSEFLOAT Function</i> | Evaluates a String input against the Decimal datatype. If the input matches, the function outputs a Decimal value. Input can be a literal, a column of values, or a function returning String values.  |
| <i>PARSEINT Function</i>   | Evaluates a String input against the Integer datatype. If the input matches, the function outputs an Integer value. Input can be a literal, a column of values, or a function returning String values. |

### PARSEDATE function now supports four default Datetime format values

Prior to Release 6.11, the PARSEDATE function required that you submit an array of Datetime formats as the second parameter of the function.

In Release 6.11, the PARSEDATE function supports the following default Datetime values:

```
'yyyy-MM-dd HH:mm:ss'
'yyyy/MM/dd HH:mm:ss'
'yyyy-MM-dd'
'yyyy/MM/dd'
```

**NOTE:** These defaults are used only if the function reference does not contain a second parameter of an array of valid Datetime formats.

For more information, see *PARSEDATE Function*.

### ignore case parameter added to string functions

In Release 6.11, the following functions now support an ignore case parameter. It is not required, and the default value is `false`. By default, matches are case-sensitive.

- *STARTSWITH Function*
- *ENDSWITH Function*
- *EXACT Function*
- *MATCHES Function*
- *STRINGGREATERTHAN Function*
- *STRINGGREATERTHANEQUAL Function*
- *STRINGLESSTHAN Function*
- *STRINGLESSTHANEQUAL Function*
- *SUBSTITUTE Function* (additional details below)

### Expanded parameters for SUBSTITUTE function

In Release 6.11, the accepted parameter inputs have been expanded.

| Parameter Name     | pre-Release 6.11 inputs                                     | Release 6.11 or later inputs   | Notes   |
|--------------------|---|--|---|
| string_source      | String literal, column, or function returning String values | No change  |   |
| string_pattern     | String literal, pattern, or regex                           | String literal, pattern, or regex or a column or function returning String values    |   |
| string_replacement | String literal  | String literal or a column or function returning String values                       |   |
| ignore_case        | n/a   | If <code>true</code> , matching is case-insensitive. Default is <code>false</code> . | New parameter   |
| pattern_before     | String literal, pattern, or regex                           | No change  | Not permitted when <code>string_pattern</code> or <code>string_replacement</code> is of column data type. |
| pattern_after      | String literal, pattern, or regex                           | No change  | Not permitted when <code>string_pattern</code> or <code>string_replacement</code> is of column data type. |

For more information, see *SUBSTITUTE Function*.

## Release 6.9

### MODE functions return lowest value in evaluated set if there is a tie

Suppose you have the following set of values:

| Price | Quantity |
|-------|----------|
| 5.23  | 3        |
| 3     | 7        |
| 7.88  | 3        |
| -3.12 | -1       |
| 0     | 6        |

|      |    |
|------|----|
| 5.23 | 0  |
| 8.37 | 38 |

You apply the following transformation:

|                                   |                             |
|-----------------------------------|-----------------------------|
| <b>Transformation Name</b>        | New formula                 |
| <b>Parameter: Formula type</b>    | Single row formula          |
| <b>Parameter: Formula</b>         | MODEIF(Quantity, Price < 5) |
| <b>Parameter: New column name</b> | 'modeif-test'               |

For the rows where price < 5, there is no most commonly occurring value in Quantity.

In Release 6.8 and earlier, the returned value was `null`.

In Release 6.9 and later, the returned value is the lowest value among the evaluated set. Among the qualifying rows, the lowest value in the Quantity column is -1.

- See *MODE Function*.
- See *MODEIF Function*.
- See *ROLLINGMODE Function*.
- See *LISTMODE Function*.

## New Functions

This release introduces the following statistical functions.

| Function Name              | Description   |
|----------------------------|---|
| <i>MEDIAN Function</i>     | Computes the median from all row values in a column or group. Input column can be of Integer or Decimal.                  |
| <i>PERCENTILE Function</i> | Computes a specified percentile across all row values in a column or group. Input column can be of Integer or Decimal.    |
| <i>QUARTILE Function</i>   | Computes a specified quartile across all row values in a column or group. Input column can be of Integer or Decimal.      |
| <i>CORREL Function</i>     | Computes the correlation coefficient between two columns. Source values can be of Integer or Decimal type.                |
| <i>COVAR Function</i>      | Computes the covariance between two columns using the population method. Source values can be of Integer or Decimal type. |
| <i>COVARSAAMP Function</i> | Computes the covariance between two columns using the sample method. Source values can be of Integer or Decimal type.     |

## Release 6.8

## New Functions

This release introduces the sampling method of calculating statistical functions. The following are now available:

| Function Name             | Description   |
|---------------------------|---|
| <i>STDEVSAMP Function</i> | Computes the standard deviation across column values of Integer or Decimal type using the sample statistical method.  |
| <i>VARSAAMP Function</i>  | Computes the variance among all values in a column using the sample statistical method. Input column can be of Integer or Decimal. If no numeric values are detected in the input column, the function returns 0. |

|                                   |   |
|-----------------------------------|---|
| <i>STDEVSAMPIF Function</i>       | Generates the standard deviation of values by group in a column that meet a specific condition using the sample statistical method.                       |
| <i>VARSAMPIF Function</i>         | Generates the variance of values by group in a column that meet a specific condition using the sample statistical method.                                 |
| <i>ROLLINGSTDEV SAMP Function</i> | Computes the rolling standard deviation of values forward or backward of the current row within the specified column using the sample statistical method. |
| <i>ROLLINGVARSA MP Function</i>   | Computes the rolling variance of values forward or backward of the current row within the specified column using the sample statistical method.           |

## Release 6.6

### New Functions

| Function Name         | Description   |
|-----------------------|---|
| <i>SINH Function</i>  | Computes the hyperbolic sine of an input value for a hyperbolic angle measured in radians. The value can be a Decimal or Integer literal or a reference to a column containing numeric values.    |
| <i>COSH Function</i>  | Computes the hyperbolic cosine of an input value for a hyperbolic angle measured in radians. The value can be a Decimal or Integer literal or a reference to a column containing numeric values.  |
| <i>TANH Function</i>  | Computes the hyperbolic tangent of an input value for a hyperbolic angle measured in radians. The value can be a Decimal or Integer literal or a reference to a column containing numeric values. |
| <i>ASINH Function</i> | Computes the arcsine of an input value for a hyperbolic angle measured in radians. The value can be a Decimal or Integer literal or a reference to a column containing numeric values.            |
| <i>ACOSH Function</i> | Computes the arccosine of an input value for a hyperbolic angle measured in radians. The value can be a Decimal or Integer literal or a reference to a column containing numeric values.          |
| <i>ATANH Function</i> | Computes the arctangent of an input value for a hyperbolic angle measured in radians. The value can be a Decimal or Integer literal or a reference to a column containing numeric values.         |

## Release 6.5

### New Functions

| Function Name        | Description   |
|----------------------|---|
| <i>SIN Function</i>  | Computes the sine of an input value for an angle measured in radians. The value can be a Decimal or Integer literal or a reference to a column containing numeric values.   |
| <i>COS Function</i>  | Computes the cosine of an input value for an angle measured in radians. The value can be a Decimal or Integer literal or a reference to a column containing numeric values.   |
| <i>TAN Function</i>  | Computes the tangent of an input value for an angle measured in radians. The value can be a Decimal or Integer literal or a reference to a column containing numeric values.  |
| Cotangent Function   | See <i>TAN Function</i> .   |
| Secant Function      | See <i>COS Function</i> .   |
| Cosecant Function    | See <i>SIN Function</i> .   |
| <i>ASIN Function</i> | For input values between -1 and 1 inclusive, this function returns the angle in radians whose sine value is the input. This function is the inverse of the sine function. The value can be a Decimal or Integer literal or a reference to a column containing numeric values. |

|                       |   |
|-----------------------|---|
| <i>ACOS Function</i>  | For input values between -1 and 1 inclusive, this function returns the angle in radians whose cosine value is the input. This function is the inverse of the cosine function. The value can be a Decimal or Integer literal or a reference to a column containing numeric values.   |
| <i>ATAN Function</i>  | For input values between -1 and 1 inclusive, this function returns the angle in radians whose tangent value is the input. This function is the inverse of the tangent function. The value can be a Decimal or Integer literal or a reference to a column containing numeric values. |
| Arccotangent Function | See <i>ATAN Function</i> .  |
| Arcsecant Function    | See <i>ACOS Function</i> .  |
| Arccosecant Function  | See <i>ASIN Function</i> .  |

## Release 6.4

### Improvements to metadata references

Broader support for metadata references: For Excel files, `$filepath` references now return the location of the source Excel file. Sheet names are appended to the end of the reference. See *Source Metadata References*.

## Release 6.3

### New Functions

| Function Name              | Description   |
|----------------------------|---|
| <i>PARSEDAT E Function</i> | Evaluates an input against the default input formats or (if specified) an array of Datetime format strings in their listed order. If the input matches one of the formats, the function outputs a Datetime value. |

### Optional input formats for DateFormat task

The DateFormat task now supports a new parameter: Input Formats. This parameter specifies the date format to use when attempting to parse the input column.

- If the parameter is specified, then the value of the parameter is used to parse the inputs.
- (default) if the parameter is not specified, then the following common formats are used for parsing the input:



```
'M/d/yy',
'MM/dd/yy',
'MM-dd-yy',
'M-d-yy',
'MMM d, yyyy',
'MMMM d, yyyy',
'EEEE, MMMM d, yyyy',
'MMM d yyyy',
'MMMM d yyyy',
'MM-dd-yyyy',
'M-d-yyyy',
'yyyy-MM-ddXXX',
'dd/MM/yyyy',
'd/M/yyyy',
'MM/dd/yyyy',
'M/d/yyyy',
'yyyy/M/d',
'M/d/yy h:mm a',
'MM/dd/yy h:mm a',
'MM-dd-yy h:mm a',
'MMM dd yyyy HH.MM.SS xxx',
'M-d-yy h:mm a',
'MMM d, yyyy h:mm:ss a',
'EEEE, MMMM d, yyyy h:mm:ss a X',
'EEE MMM dd HH:mm:ss X yyyy',
'EEE, d MMM yyyy HH:mm:ss X',
'd MMM yyyy HH:mm:ss X',
'MM-dd-yyyy h:mm:ss a',
'M-d-yyyy h:mm:ss a',
'yyyy-MM-dd h:mm:ss a',
'yyyy-M-d h:mm:ss a',
'yyyy-MM-dd HH:mm:ss.S',
'dd/MM/yyyy h:mm:ss a',
'd/M/yyyy h:mm:ss a',
'MM/dd/yyyy h:mm:ss a',
'M/d/yyyy h:mm:ss a',
'MM/dd/yy h:mm:ss a',
'MM/dd/yy H:mm:ss',
'M/d/yy H:mm:ss',
'dd/MM/yyyy h:mm a',
'd/M/yyyy h:mm a',
'MM/dd/yyyy h:mm a',
'M/d/yyyy h:mm a',
'MM-dd-yy h:mm:ss a',
'M-d-yy h:mm:ss a',
'MM-dd-yyyy h:mm a',
'M-d-yyyy h:mm a',
'yyyy-MM-dd h:mm a',
'yyyy-M-d h:mm a',
'MMM.dd.yyyy',
'd/MMM/yyyy H:mm:ss X',
'dd/MMM/yy h:mm a',
```

These formats are a subset of the date formatting strings supported by the product. For more information, see *Datetime Data Type*.

## Release 6.2

### New Functions

| Function Name        | Description   |
|----------------------|---|
| <i>RANK Function</i> | Computes the rank of an ordered set of value within groups. Tie values are assigned the same rank, and the next ranking is incremented by the number of tie values. |

|                           |  |
|---------------------------|--|
| <i>DENSERANK Function</i> | Computes the rank of an ordered set of value within groups. Tie values are assigned the same rank, and the next ranking is incremented by 1. |
|---------------------------|--|

### ARRAYELEMENTAT function accepts new inputs

In previous releases, the ARRAYELEMENTAT function accepted a second input parameter to specify the index value of the element to retrieve. This "at" parameter had to be an Integer literal.

Beginning in this release, the function also accepts for this second "at" parameter:

- Names of columns containing Integer values
- Functions that return Integer values

For more information, see *ARRAYELEMENTAT Function*.

### Release 6.1

None.

### Release 6.0

#### New Functions

| Function Name                      | Description  |
|------------------------------------|--|
| <i>ARRAYINDEXOF Function</i>       | Computes the index at which a specified element is first found within an array. Indexing is left to right.   |
| <i>ARRAYRIGHTINDEXOF Function</i>  | Computes the index at which a specified element is first found within an array, when searching right to left. Returned value is based on left-to-right indexing. |
| <i>ARRAYSLICE Function</i>         | Returns an array containing a slice of the input array, as determined by starting and ending index parameters.   |
| <i>ARRAYMERGEELEMENTS Function</i> | Merges the elements of an array in left to right order into a string. Values are optionally delimited by a provided delimiter.                                   |

### Changes to LIST\* inputs

The following LIST-based functions have been changed to narrow the accepted input data types. In previous releases, any data type was accepted for input, which was not valid for most data types.

In Release 6.0 and later, these functions accept only Array inputs. Inputs can be Array literals, a column of Arrays, or a function returning Arrays.

**NOTE:** You should references to these functions in your recipes.

| LIST* Functions             |
|-----------------------------|
| <i>LISTAVERAGE Function</i> |
| <i>LISTMIN Function</i>     |
| <i>LISTMAX Function</i>     |
| <i>LISTMODE Function</i>    |
| <i>LISTSTDEV Function</i>   |
| <i>LISTSUM Function</i>     |

## Renamed functions

The following functions have been renamed in Release 6.0.

| Release 5.9 and earlier | Release 6.0 and later |
|-------------------------|-----------------------|
| LISTUNIQUE Function     | UNIQUE Function       |

## FILL Function has new before and after parameters

Prior to Release 6.0, the FILL function replaced empty cells with the most recent non-empty value.

In Release 6.0, *before* and *after* function parameters have been added. These parameters define the window of rows before and after the row being tested to search for non-empty values. Within this window, the most recent non-empty value is used.

The default values for these parameters are `-1` and `0` respectively, which performs a search of an unlimited number of preceding rows for a non-empty value.

**NOTE:** Upon upgrade, the FILL function retains its preceding behavior, as the default values for the new parameters perform the same unlimited row search for non-empty values.

For more information, see *FILL Function*.

## Release 5.9

### New functions

The following functions can now be applied directly to arrays to derive meaningful statistics about them.

| Function               | Description  |
|------------------------|--|
| ARRAYSORT Function     | Sorts array values in the specified column, array literal, or function that returns an array in ascending or descending order.                 |
| TRANSLITERATE Function | Transliterates Asian script characters from one script form to another. The string can be specified as a column reference or a string literal. |

## Release 5.8

### File lineage information using source metadata references

Beginning in Release 5.8, you can insert the following references into the formulas of your transformations. These **source metadata references** enable you to continue to track file lineage information from within your datasets as part of your wrangling project.

**NOTE:** These references apply only to file-based sources. Some additional limitations may apply.

| reference               | Description  |
|-------------------------|--|
| <code>\$filepath</code> | Returns the full path and filename of the source of the dataset. |

|                   |  |
|-------------------|--|
| \$sourcerownumber | Returns the row number for the current row from the original source of the dataset. <div> <b>NOTE:</b> This reference is equivalent to the SOURCEROWNUMBER function, which is likely to be deprecated in a future release. You should begin using this reference in your recipes. </div> |
|-------------------|--|

For more information, see *Source Metadata References*.

## New math and statistical functions for arrays

The following functions can now be applied directly to arrays to derive meaningful statistics about them.

| Function                    | Description   |
|-----------------------------|---|
| <i>LISTSUM Function</i>     | Computes the sum of all numeric values found in input array. Input can be an array literal, a column of arrays, or a function returning an array. Input values must be of Integer or Decimal type.                |
| <i>LISTMAX Function</i>     | Computes the maximum of all numeric values found in input array. Input can be an array literal, a column of arrays, or a function returning an array. Input values must be of Integer or Decimal type.            |
| <i>LISTMIN Function</i>     | Computes the minimum of all numeric values found in input array. Input can be an array literal, a column of arrays, or a function returning an array. Input values must be of Integer or Decimal type.            |
| <i>LISTAVERAGE Function</i> | Computes the average of all numeric values found in input array. Input can be an array literal, a column of arrays, or a function returning an array. Input values must be of Integer or Decimal type.            |
| <i>LISTVAR Function</i>     | Computes the variance of all numeric values found in input array. Input can be an array literal, a column of arrays, or a function returning an array. Input values must be of Integer or Decimal type.           |
| <i>LISTSTDEV Function</i>   | Computes the standard deviation of all numeric values found in input array. Input can be an array literal, a column of arrays, or a function returning an array. Input values must be of Integer or Decimal type. |
| <i>LISTMODE Function</i>    | Computes the most common value of all numeric values found in input array. Input can be an array literal, a column of arrays, or a function returning an array. Input values must be of Integer or Decimal type.  |

## Release 5.7

### WEEKNUM function now behaves consistently across running environments

In Release 5.6 and earlier, the WEEKNUM function treated the first week of the year differently between the Trifacta Photon and Spark running environments:

- **Trifacta Photon week 1 of the year:** The week that contains January 1.
- **Spark week 1 of the year:** The week that contains at least four days in the specified year.

This issue was caused by Spark following an ISO-8601 standard and relying on the joda datetimeformatter.

Beginning in Release 5.7, the WEEKNUM function behaves consistently for both Trifacta Photon and Spark:

- **Week 1 of the year:** The week that contains January 1.

For more information, see *WEEKNUM Function*.

## Release 5.6

### URLPARAMS function returns null values

In Release 5.1 and earlier, the URLPARAMS function returned empty Objects when no answer was computed for the function.

In Release 5.6 and later, this function returns null values in the above case.

See *URLPARAMS Function*.

## Release 5.1

### Wrangle now supports nested expressions

Beginning in Release 5.1, all Wrangle functions now supported nested expressions, which can be arithmetic calculations, column references, or other function calls.

**NOTE:** This feature is enabled by default, as this change does not break any steps created in previous versions of the product. It can be disabled if needed. See *Miscellaneous Configuration*.

**NOTE:** This capability represents a powerful enhancement to the language, as you can now use dynamic inputs for all functions.

The following expression is a valid transform in Wrangle. It locates the substring in `myString` that begins with the `@` sign until the end of the string, inclusive:

```
derive value: substring(myString, find(myString, '@', true, 0), length(myString))
```

#### Nested arithmetic expressions:

Suppose you wanted just the value after the `@` sign until the end of the string. Prior to Release 5.1, the following generated a validation error:

```
derive value: substring(myString, find(myString, '@', true, 0) + 1, length(myString))
```

In the above, the addition of `+1` to the second parameter is a nested expression and was not supported. Instead, you had to use multiple steps to generate the string value.

Beginning in Release 5.1, the above single-step transform is supported.

#### Nested column references:

In addition to arithmetic expressions, you can nested column references. In the following example, the previous step has been modified to replace the static `+1` with a reference to a column containing the appropriate value (`at_sign_offset`):

```
derive value: substring(myString, find(myString, '@', true, 0) + at_sign_offset, length(myString))
```

#### Nested function references:

Now, you can combine multiple function references into a single computation. The following computes the total volume of a cube of length `side` and then multiplies that volume by the number of cubes (`cube_count`) to compute the total `cube_volume`

```
derive type: single value: MULTIPLY(POW(cube_side,3),cube_count) as: 'cube_volume'
```

For more information, see *Wrangle Language*.

## SOURCEROWNUMBER function generates null values consistently

The SOURCEROWNUMBER function returns the row number of the row as it appears in the original dataset. After some operations, such as unions, joins, and aggregations, this row information is no longer available.

In Release 5.0.1 and earlier, the results were confusing. When source row information was not available, the function was simply not available for use.

In Release 5.1 and later, the behavior of the SOURCEROWNUMBER function is more consistent:

- If the source row information is available, it is returned.
- If it is not available:
  - The function can still be used.
  - The function returns null values in all cases.

For more information, see *SOURCEROWNUMBER Function*.

## New Functions

| Function Name                         | Description  |
|---------------------------------------|--|
| <i>ARRAYELEMENTAT Function</i>        | Returns element value of input array for the provided index value.   |
| <i>DOUBLEMETAPHONE Function</i>       | Returns primary and secondary phonetic spellings of an input string using the Double Metaphone algorithm.                                |
| <i>DOUBLEMETAPHONEEQUALS Function</i> | Returns <code>true</code> if two strings match phonetic spellings using Double Metaphone algorithm. Tolerance threshold can be adjusted. |
| <i>UNIQUE Function</i>                | Generates a new column containing an array of the unique values from a source column.  |

## Release 5.0.1

### RAND function generates true random numbers

In Release 5.0 and earlier, the RAND function produced the same set of random numbers within the browser, after browser refresh, and over subsequent runs of a job.

- During job execution, a default **seed value** was inserted as the basis for the function during the execution of the job.
- In some cases, this behavior is desired.

In Release 5.0.1 and later, the RAND function accepts an optional integer as a parameter. When this new seed value is inserted, the function generates deterministic, pseudo-random values.

- This version matches the behavior of the old function.

**NOTE:** On all upgraded instances of the platform, references to the RAND function have been converted to use a default seed value, so that previous behavior is maintained in the upgraded version.

- If no seed value is inserted as a parameter, the RAND function generates true random values within the browser, after browser refresh, and over subsequent job runs.

**NOTE:** Be aware that modifying your dataset based on the generated values of `RAND( )` may have unpredictable effects later in your recipe and downstream of it.

For more information, see *RAND Function*.

## Release 5.0

### Required type parameter

Prior to Release 5.0, the following was a valid *Wrangle* step:

```
derive value:colA + colB as:'colC'
```

Beginning in Release 5.0, the `type` parameter is required. This parameter defines whether the transform is a single or multi-row formula. In the Transform Builder, this value must be specified.

The following is valid in Release 5.0:

```
derive type:single value:colA + colB as:'colC'
```

See *Derive Transform*.

See *Transform Builder*.

### Deprecated aggregate transform

In Release 4.2.1 and earlier, the aggregate transform could be used to aggregate your datasets using aggregation functions and groupings.

In Release 5.0 and later, this transform has been merged into the pivot transform. The aggregate transform has been deprecated and is no longer available.

**NOTE:** During upgrade to Release 5.0 and later, recipes that had previously used the aggregate transform are automatically migrated to use the pivot equivalent.

#### Example 1

Release 4.2.1 and earlier Aggregate:

```
aggregate value:AVERAGE(Scores)
```

Release 5.0 and later Pivot:

```
pivot value: AVERAGE(Score) limit: 1
```

The `limit` parameter defines the maximum number of columns that can be generated by the pivot.

#### Example 2

Aggregate:

```
aggregate value:AVERAGE(Scores) group:studentId
```

Pivot:

```
pivot group: StudentId value: AVERAGE(Score) limit: 1
```

For more information, see *Pivot Transform*.

## New search terms

In the new Search panel, you can search for terms that can be used to select transformations for quick population of parameters. In the following table, you can see how terminology has changed in Release 5.0 for some common transforms from earlier release.

**Tip:** You can paste the Release 5.0 terms in the Search panel to locate the same transformations used in earlier releases.

| Release 4.2.1 and earlier transforms | Release 5.0 and later search terms |
|--------------------------------------|------------------------------------|
| aggregate                            | pivot                              |
| keep                                 | filter                             |
| delete                               | filter                             |
| extract on:                          | extractpatterns                    |
| extract at:                          | extractpositions                   |
| extract before:                      | extractbetweenlimits               |
| extract after:                       | extractbetweenlimits               |
| replace on:                          | replacepatterns                    |
| replace at:                          | replacepositions                   |
| replace before:                      | replacebetweenpatterns             |
| replace after:                       | replacebetweenpatterns             |
| replace from:                        | replacebetweenpatterns             |
| replace to:                          | replacebetweenpatterns             |
| split on:                            | splitpatterns                      |
| split delimiters:                    | splitpositions                     |
| split every:                         | splitpositions                     |
| split positions:                     | splitpositions                     |
| split after:                         | splitpatterns                      |
| split before:                        | splitpatterns                      |
| split from:                          | splitpatterns                      |
| split to:                            | splitpatterns                      |

## Support for <> operator

Prior to Release 5.0, the following operator was used to test "not equal" comparisons:

!=

Beginning in Release 5.0, the following operators is also supported:

<>

Example:



```
derive value:IF ((col1 <> col2), 'different','equal') as:'testNotEqual'
```

**Tip:** Both of the above operators are supported, although the <> operator is preferred.

For more information, see *Comparison Operators*.

## ROUND function takes optional number of digits

The ROUND function now supports rounding to a specified number of digits. By default, values are rounded to the nearest integer, as before. See *ROUND Function*.

## New Functions

| Function Name                           | Description   |
|---|---|
| <i>DEGREES Function</i>                 | Generates the value in degrees for an input radians value.  |
| <i>EXACT Function</i>                   | Compares two strings to see if they are exact matches.  |
| <i>FILTEROBJECT Function</i>            | Filters the keys and values from an Object based on specified keys.   |
| <i>HOST Function</i>                    | Returns the host value from a URL.  |
| <i>ISEVEN Function</i>                  | Returns <code>true</code> if an Integer, function returning an Integer, or a column contains an even value. |
| <i>ISODD Function</i>                   | Returns <code>true</code> if an Integer, function returning an Integer, or a column contains an odd value.  |
| <i>KTHLARGESTUNIQUE Function</i>        | Computes the <i>kth</i> -ranked unique value in a set of values.  |
| <i>LCM Function</i>                     | Returns the least common multiple between two input values.   |
| <i>MODE Function</i>                    | Computes the mode (most common) value for a set of values.  |
| <i>MODEIF Function</i>                  | Computes the mode based on a conditional test.  |
| <i>PAD Function</i>                     | Pads the left or right side of a value with a specified character string.                                   |
| <i>PI Function</i>                      | Generates the value for pi to 15 decimal places.  |
| <i>RADIANS Function</i>                 | Generates the value in radians for an input degrees value.  |
| <i>RANDBETWEEN Function</i>             | Generates a random Integer in a range between two specified values.   |
| <i>RIGHTFIND Function</i>               | Locates a substring by searching from the right side of an input value.                                     |
| <i>ROLLINGCOUNTA Function</i>           | Computes count of non-null values across a rolling window within a column.                                  |
| <i>ROLLINGKTHLARGEST Function</i>       | Computes the <i>kth</i> largest value across a rolling window within a column.                              |
| <i>ROLLINGKTHLARGESTUNIQUE Function</i> | Computes the <i>kth</i> largest unique value across a rolling window within a column.                       |
| <i>ROLLINGLIST Function</i>             | Computes list of all values across a rolling window within a column.  |
| <i>ROLLINGMAX Function</i>              | Computes maximum value across a rolling window within a column.   |
| <i>ROLLINGMIN Function</i>              | Computes minimum value across a rolling window within a column.   |
| <i>ROLLINGMODE Function</i>             | Computes mode (most common) value across a rolling window within a column.                                  |
| <i>ROLLINGSTDEV Function</i>            | Computes standard deviation across a rolling window within a column.  |
| <i>ROLLINGVAR Function</i>              | Computes variance across a rolling window within a column.  |
| <i>SIGN Function</i>                    | Computes the positive or negative sign of an input value.   |
| <i>TRUNC Function</i>                   | Truncates a value to the nearest integer or a specified number of digits.                                   |
| <i>URLPARAMS Function</i>               | Extracts any query parameters from a URL into an Object.  |

|                         |   |
|-------------------------|---|
| <i>WEEKNUM Function</i> | Calculates the week that the date appears during the year (1-52). |
|-------------------------|---|

## Release 4.2.1

None.

## Release 4.2

### New Filter transform

Perform a variety of predefined row filtrations using the new `filter` transform, or apply your own custom formula to keep or delete rows from your dataset.

- See *Remove Data*.
- See *Filter Transform*.

### New Case transform

Beginning in Release 4.2, you can use the Transform Builder to simplify the construction of `CASE` statements. For each case, specify the conditional and resulting expression in separate textboxes.

- See *Apply Conditional Transformations*.
- See *Case Transform*.

### Rename transform now supports multi-column rename

Use the `rename` transform to rename multiple columns in a single transform.

- See *Rename Columns*.
- See *Rename Transform*.

### Delete specified columns or delete the others

The `drop` transform now supports the option of deleting all columns except the ones specified in the transform. See *Drop Transform*.

### New string comparison functions

Compare two strings using Latin collation settings. See below.

### NOW function returns 24-hour time values

In Release 4.1.1 and earlier, the `NOW` function returned time values for the specified time zone in 12-hour time, which was confusing.

In Release 4.2 and later, this function returns values in 24-hour time.

### New Transforms

| Transform Name      | Documentation           |
|---------------------|-------------------------|
| <code>case</code>   | <i>Case Transform</i>   |
| <code>filter</code> | <i>Filter Transform</i> |

New Functions

| Function Name          | Documentation                          |
|------------------------|--|
| STRINGGREATERTHAN      | <i>STRINGGREATERTHAN Function</i>      |
| STRINGGREATERTHANEQUAL | <i>STRINGGREATERTHANEQUAL Function</i> |
| STRINGLESSTHAN         | <i>STRINGLESSTHAN Function</i>         |
| STRINGLESSTHANEQUAL    | <i>STRINGLESSTHANEQUAL Function</i>    |
| SUBSTITUTE             | <i>SUBSTITUTE Function</i>             |

# Changes to the APIs

## Contents:

- *Release 8.11*
  - *writesettings objects are no longer validated for API-based jobs*
- *Release 8.8*
  - *API rate limiting*
- *Release 8.7*
  - *Removal of old asset transfer endpoint*
  - *Deprecation of not useful connection endpoints*
- *Changes for Release 8.2*
- *Release 8.1*
  - *Customize relational connectors via API*
- *Release 8.0*
- *Release 7.7*
  - *credentialProvider no longer required for /v4/awsConfig*
  - */v4/awsConfigs PUT method has been deprecated*
- *Release 7.5*
  - *Flow sharing API now accepts user email addresses*
- *Release 7.1*
  - *Introducing in-app API reference documentation*
  - *API changes for authorization*
  - *Create Salesforce connection changes*
  - *Apply overrides to file-based data sources when you run a job*
- *Release 7.0*
  - *v3 endpoints are no longer available*
  - *Path to file-based imported datasets can be specified by URI*
  - *awsRoles GET method removed*
  - *Connections endpoints simplified*
  - *Publishing for a jobGroup has been simplified*
  - *WASB URI format has changed*
  - *Saving Databricks Personal Access Token has different response*
  - *Documentation for awsRoles object corrected*
- *Changes for Release 6.8.1*
  - *Improvements to AssetTransfer API documentation*
- *Changes for Release 6.8*
  - *Overrides to relational sources and targets*
  - *New flow object definition*
  - *Export and import macros*
- *Release 6.4*
  - *Request format for assigning connection permissions endpoint has changed*
  - *v4 version of password reset request endpoint*
  - *Changes to awsConfig object*
- *Release 6.3*
  - *Assign AWSConfigs to a user at create time*
- *Release 6.0*
  - *Error in Release 6.0.x API docs*
  - *Planned End of Life of v3 API endpoints*
- *Release 5.9*
  - *Introducing Access Tokens*
- *Release 5.1*
- *Release 5.0*
  - *Introducing v4 APIs*
- *Release 4.2*
  - *Create Hive and Redshift connections via API*
  - *WrangledDataset endpoints are still valid*

Review the changes to the publicly available REST APIs for Trifacta® for the current release and past releases.

## Release 8.11

### writesettings objects are no longer validated for API-based jobs

Beginning in this release, the settings in a writesettings object are no longer validated as part of job execution for scheduled and API-based jobs. A writesettings object defines the file-based output for a job.

**NOTE:** When this feature is enabled, you should first test scheduled and API-based jobs by running them through the Trifacta application. These validations are still performed during manual job execution. If the validation fails, the job can fail during transformation or publishing stages.

This feature is enabled by default. A workspace administrator can disable the feature if needed.

For more information, see *Workspace Settings Page*.

## Release 8.8

### API rate limiting

For some API endpoints that are heavily used or critical to the success use of the platform, rate limiting has been applied. API rate limiting sets the maximum number of requests that are accepted by an endpoint within a specified period of time. For example, you may be prevented from sending more than 50 requests to start a job or plan run within a minute.

**NOTE:** All rate-limited endpoints receive the default quota limit, which is determined based on the product tier. If a rate limit quota is reached, a small set of subsequent requests is enqueued. If the queue is filled before the end of the minute, subsequent requests return a 429 status code error.

**Tip:** Where applicable, the quotas for individual APIs are listed with the relevant API endpoint documentation.

See: <https://api.trifacta.com/ee/es.t/index.html#section/Overview/Rate-limiting>

## Release 8.7

### Removal of old asset transfer endpoint

The following endpoint for transferring assets between users has been removed from the product:

```
/v4/people/{to-userid:}/assetTransfer/{from-userid:}
```

The above endpoint no longer works.

In a previous release, this endpoint was replaced by an improved method of asset transfer. For more information, see "Asset transfer API using email addresses" below.

### Deprecation of not useful connection endpoints

The following endpoints have been deprecated:

```
/v4/connections/vendors  
/v4/connections/credentialTypes  
/v4/connections/:id/publish/info  
/v4/connections/:id/import/info
```

These endpoints were scheduled for deprecation in Release 8.2 and are not widely used.

**Tip:** To acquire the definition of a specific connection, you can try the following:

```
/v4/connectormetadata/:connectorname
```

See <https://api.trifacta.com/ee/es.t/index.html#operation/getConnectorConfig>

## Changes for Release 8.2

None.

## Release 8.1

### Customize relational connectors via API

Beginning in Release 8.1, you can customize aspects of each relational connection type available in your product edition through a set of APIs. Some terms:

- **connection:** In the Trifacta application and via API, you can create and manage connections between the platform and a specific datastore. A connection is the user-defined object that enables the connection to the datastore.
- **connector:** A connection interfaces with a connector, which is an underlying driver and its related configuration, that perform the actual connection. This configuration information includes runtime, publishing, and connection definitions.

**Tip:** All connections of the same type use the same underlying connector, including its configuration. Overrides that you apply to a connector apply to all current and new connections of that type in the workspace.

**NOTE:** In prior releases, Trifacta Self-Managed Enterprise Edition maintained relational connection type definitions on the Trifacta node. Beginning in this release, the default definitions are maintained in the new Connector Configuration Service database, and the existing node files contain the customized values for those definitions. You can continue to modify those files directly on the Trifacta node or make changes using the APIs listed below.

### Get connector identifier

To use these API endpoints, you must acquire the connector identifier. This value is the `vendor` value for a connection of the type. You can acquire this value in one of two ways:

- Create a connection of the type in the Trifacta application. Use the `/v4/connections/:id` endpoint with the GET method to acquire the connection information for your connection. Acquire the `vendor` value.  
<https://api.trifacta.com/ee/es.t/index.html#operation/listConnections>
- You may find the `vendor` values listed in the documentation. See *Connection Types*.

#### Get connector metadata information - defaults

The following endpoint returns the default metadata information for a specified connector type. This information is stored in the Connector Configuration Service database.

For the `:connectorId` value below, use the `vendor` value that you acquired above. For example, to acquire connector type definitions for MySQL connection type, use the value `mysql`.

|                      |   |
|----------------------|---|
| <b>Endpoint</b>      | <code>/v4/connectorMetadata/:connectorId/defaults</code>  |
| <b>Method</b>        | GET   |
| <b>Description</b>   | Get the default metadata for a connector without applying custom overrides. This metadata is used to defined connectivity, ingestion, and publishing for the connector. |
| <b>Documentation</b> | <a href="https://api.trifacta.com/ee/es.t/index.html#operation/getConnectorDefaults">https://api.trifacta.com/ee/es.t/index.html#operation/getConnectorDefaults</a>     |

#### Get connector metadata information - current values

The following endpoint acquires the current metadata for a specified connector type, which include the default values with any applicable overrides applied to them.

|                      |   |
|----------------------|---|
| <b>Endpoint</b>      | <code>/v4/connectorMetadata/:connectorId/</code>  |
| <b>Method</b>        | GET   |
| <b>Description</b>   | Get the consolidated metadata for a connector in a given workspace. This metadata is used to defined connectivity, ingestion, and publishing for the connector. |
| <b>Documentation</b> | <a href="https://api.trifacta.com/ee/es.t/index.html#operation/getConnectorConfig">https://api.trifacta.com/ee/es.t/index.html#operation/getConnectorConfig</a> |

#### Get connector metadata information - get overrides values

The following endpoint retrieves the overrides that have been applied to a specific connector.

|                      |   |
|----------------------|---|
| <b>Endpoint</b>      | <code>/v4/connectorMetadata/:connectorId/overrides</code>   |
| <b>Method</b>        | GET   |
| <b>Description</b>   | Get the metadata overrides for a connector in a given workspace. These overrides are applied to the base configuration for connectivity operations.                   |
| <b>Documentation</b> | <a href="https://api.trifacta.com/ee/es.t/index.html#operation/getConnectorOverrides">https://api.trifacta.com/ee/es.t/index.html#operation/getConnectorOverrides</a> |

#### Create overrides for a connector

The following endpoint applies the specified value or values as overrides to the connector.

|                    |   |
|--------------------|---|
| <b>Endpoint</b>    | <code>/v4/connectorMetadata/:connectorId/overrides</code>   |
| <b>Method</b>      | POST  |
| <b>Description</b> | The specified overrides are merged into the current set of overrides for the connector. A new entry is created if no overrides currently exist. |

|               |  |
|---------------|--|
| Documentation | <div> <b>Tip:</b> Overrides are specified in the request body. See the link below for more. </div> <a href="https://api.trifacta.com/ee/es.t/index.html#operation/updateConnectorOverrides">https://api.trifacta.com/ee/es.t/index.html#operation/updateConnectorOverrides</a> |
|---------------|--|

### Delete overrides for a connector

The following endpoint deletes all override values for a specified connector.

|               |   |
|---------------|---|
| Endpoint      | <code>/v4/connectorMetadata/:connectorId/overrides</code>   |
| Method        | DELETE  |
| Description   | All overrides are deleted. The connector reverts to the base configuration.   |
| Documentation | <a href="https://api.trifacta.com/ee/es.t/index.html#operation/deleteConnectorOverrides">https://api.trifacta.com/ee/es.t/index.html#operation/deleteConnectorOverrides</a> |

## Release 8.0

### Deprecation of prior asset transfer endpoint

The following endpoint has been available for transferring assets between users:

```
/v4/people/{to-userid:}/assetTransfer/{from-userid:}
```

This endpoint is not compatible with the forward direction of the platform.

**NOTE:** The above endpoint will be deprecated in a future release. Please switch to using the new endpoint for transferring assets between users.

For more information, see <https://api.trifacta.com/ee/es.t/index.html#operation/transferAssets>

## Release 7.7

### credentialProvider no longer required for /v4/awsConfig

In prior releases, API requests to the `/v4/awsConfig` endpoint using the PUT method could receive the following error:

```
{
  "exception":
  { "name": "ApiValidationFailed", "message": "Message does not adhere to API specification", "details":
    "'credentialProvider' field in request body must not be null" }}
}
```

The `credentialProvider` field is not required, and this requirement has been removed from the endpoint. <https://api.trifacta.com/ee/es.t/index.html#operation/updateAwsConfig>



## /v4/awsConfigs PUT method has been deprecated

In prior releases, the `/v4/awsConfigs` endpoint supported the use of a PUT method to modify AWS configuration objects.

This endpoint-method combination has been deprecated. The new PATCH method is the following:

**NOTE:** The PUT method for this endpoint is still accessible. In a future release, it will be removed from the platform. Please switch to using the PATCH method.

|              |   |
|--------------|---|
| Endpoint     | /v4/awsConfigs  |
| Method       | PATCH   |
| Request Body | <pre>{   "role": "&lt;my_iam_role_object_3&gt;" }</pre> |

<https://api.trifacta.com/ee/es.t/index.html#operation/updateAwsConfig>

For more information on configuring AWS access objects, see *API Workflow - Manage AWS Configurations*.

## Release 7.5

### Flow sharing API now accepts user email addresses

Beginning in Release 7.5, the API has been enhanced to allow insertion of a user's email address as part of the request body. Below is an example request:

|              |   |
|--------------|---|
| Endpoint     | /v4/flows/402/permissions/  |
| Method       | POST  |
| Request Body | <pre>{   "data": [     {       "email": "user@example.com",       "role": "collaborator",       "policy": "flow_editor"     }   ] }</pre> |

If the above returns a 201 - Success status message, then `user@example.com` has been given the role of `collaborator`, which uses the `flow_editor` policy, for flowId 402.

**Tip:** You can still use the internal identifier for the user, too.

<https://api.trifacta.com/ee/es.t/index.html#operation/shareFlow>

## Release 7.1

### Introducing in-app API reference documentation

Beginning in Release 7.1, API reference documentation is available directly from the Trifacta application. From the menu bar, select **Help menu > API Documentation**.

#### Key features:

- More endpoints and attributes are now available!
- Reference content comes directly from the codebase.
- Searchable and browsable

More content and features will be added to this new capability of the next few releases.

**NOTE:** The API reference doc that was published with the product documentation is no longer available. Please use the in-app API reference documentation.

**NOTE:** The API reference documentation may require enablement in your environment. For more information, see *API Reference*.

Workflow documentation is still available in the product documentation. For more information on these workflows, see *API Reference*.

### API changes for authorization

Release 7.1 introduces authorization, which provides finer-grained access controls to user-defined objects in the platform. For more information, see *Changes to User Management*.

As a result of authorization, the following changes have been applied to the listed API endpoints.

#### Connection permissions - POST

|          |                                  |
|----------|----------------------------------|
| Endpoint | /v4/connections/:id/permissions/ |
| Method   | POST                             |

#### Change:

When setting permissions, the request body now must include the `policyTag` parameter:

**NOTE:** This is required.

```
"policyTag": "connection_viewer",
```

| Parameter | Description   |
|-----------|---|
| policyTag | Defines the policy level to assign to the change in connection permissions. Accepted values: <ul style="list-style-type: none"><li>• connection_viewer = can use the connection</li></ul> |

For more information, see the API reference documentation:  
<https://api.trifacta.com/ee/es.t/index.html#operation/createConnectionPermission>

### Connection permissions - GET

|          |                                  |
|----------|----------------------------------|
| Endpoint | /v4/connections/:id/permissions/ |
| Method   | GET                              |

#### Change:

Response no longer includes `connectionPermissions` parameter.

**NOTE:** Above is removed.

Response does include the following new parameters:

```
"policyTag": "connection_viewer",  
"workspaceAdmin": false,  
"isCreatedBy": true,
```

| Parameter      | Description  |
|----------------|--|
| policyTag      | See previous.  |
| workspaceAdmin | If <code>true</code> , then the permitted user is a workspace admin.   |
| isCreatedBy    | If <code>true</code> , then the permitted user created the connection. |

For details on these parameters, please see the API Reference documentation:  
<https://api.trifacta.com/ee/es.t/index.html#operation/getConnectionPermissions>

### Connections - GET LIST

|          |                 |
|----------|-----------------|
| Endpoint | /v4/connections |
| Method   | GET             |

#### Change:

Response now includes an `authorizationPermission` field in the `associatedPeople` block for each connection:

```

{
  "data": [
    {
      "id": 14,
      ....
      "associatedPeople": {
        "data": [
          {
            "id": 1,
            ...
            "authorizationPermission": {
              "policyName": null,
              "policyTag": "workspace_admin",
              "resourceOperationList": [
                {
                  "operations": [
                    "execute",
                    "delete",
                    "update",
                    "read",
                    "create",
                    "share"
                  ],
                  "policyTag": "plan_author",
                  "resourceType": "plan"
                },
                {
                  "operations": [
                    "execute",
                    "delete",
                    "update",
                    "read",
                    "create",
                    "share"
                  ],
                  "policyTag": "flow_author",
                  "resourceType": "flow"
                }
              ]
            }
          }
        ]
      }
    }
  ]
}

```

For details on these parameters, please see the API Reference documentation:

<https://api.trifacta.com/ee/es.t/index.html#operation/listConnections>

#### Connections - GET

|                 |                     |
|-----------------|---------------------|
| <b>Endpoint</b> | /v4/connections/:id |
| <b>Method</b>   | GET                 |

#### Change:

Same change as previous.

For details on these parameters, please see the API Reference documentation:

<https://api.trifacta.com/ee/es.t/index.html#operation/getConnection>

#### People - GET LIST

|                 |            |
|-----------------|------------|
| <b>Endpoint</b> | /v4/people |
| <b>Method</b>   | GET        |

## Change:

Response now includes the following attributes:

- maximalPrivileges
- authorizationRoles

**NOTE:** Above attributes appear only if the authenticating user is an administrator or the specific user.

```
{
  "data": [
    {
      "id": 29,
      ....
      "maximalPrivileges": [
        {
          "operations": [
            "delete",
            "update",
            "read",
            "create",
            "share"
          ],
          "policyTag": "connection_author",
          "resourceType": "connection"
        },
        {
          "operations": [
            "execute",
            "delete",
            "update",
            "read",
            "create",
            "share"
          ],
          "policyTag": "plan_author",
          "resourceType": "plan"
        },
        {
          "operations": [
            "execute",
            "delete",
            "update",
            "read",
            "create",
            "share"
          ],
          "policyTag": "flow_author",
          "resourceType": "flow"
        }
      ],
      "authorizationRoles": [
        {
          "policyId": 13,
          "name": "default",
          "tag": null,
          "workspaceId": 1,
          "createdAt": "2020-04-22T07:50:04.502+0000",
          "updatedAt": "2020-04-22T07:50:04.502+0000",
          "resourceOperations": [
            {
              "operations": [
                "execute",
                "delete",
```

```

        "update",
        "read",
        "create",
        "share"
    ],
    "policyTag": "plan_author",
    "resourceType": "plan"
  },
  {
    "operations": [
      "execute",
      "delete",
      "update",
      "read",
      "create",
      "share"
    ],
    "policyTag": "flow_author",
    "resourceType": "flow"
  },
  {
    "operations": [
      "delete",
      "update",
      "read",
      "create",
      "share"
    ],
    "policyTag": "connection_author",
    "resourceType": "connection"
  }
]
}
],
}

```

For details on these parameters, please see the API Reference documentation:

<https://api.trifacta.com/ee/es.t/index.html#operation/listPeople>

#### People - GET

|          |                |
|----------|----------------|
| Endpoint | /v4/people/:id |
| Method   | GET            |

#### Change:

Same change as previous.

For details on these parameters, please see the API Reference documentation:

<https://api.trifacta.com/ee/es.t/index.html#operation/getPerson>

#### Create Salesforce connection changes

Beginning in Release 7.1, the structure of a Salesforce connection object has been changed. In prior releases, a Salesforce connection could not be shared without also sharing credentials. These changes allow a Salesforce connection to be shared without requiring the sharing of credentials.

Prior to Release 7.1, when creating a Salesforce connection using the APIs, the `/v4/connections` POST method required the insertion of a Salesforce security token as part of the `params` section of the request:

```
{
  "vendor": "salesforce",
  "vendorName": "salesforce",
  "type": "jdbc",
  "credentialType": "basic",
  "ssl": true,
  "name": "Salesforce",
  "description": "example_salesforce_connection",
  "disableTypeInference": true,
  "isGlobal": true,
  "credentialsShared": true,
  "host": "salesforce.example.com",
  "params": {
    "servername": "myServer",
    "securityToken": "string"
  },
  "credentials": [
    {
      "username": "string",
      "password": "string"
    }
  ]
}
```

Beginning in Release 7.1, the following changes have been applied to this request structure:

- Change `credentialType` to `securityToken`
- Move `securityToken` key from `params` to `credentials` area of the API request
- A host and port number value are not required.

```
{
  "vendor": "salesforce",
  "vendorName": "salesforce",
  "type": "jdbc",
  "credentialType": "securityToken",
  "ssl": true,
  "name": "Salesforce",
  "description": "example_salesforce_connection",
  "disableTypeInference": true,
  "isGlobal": true,
  "credentialsShared": true,
  "host": "salesforce.example.com",
  "params": {
    "servername": "myServer"
  },
  "credentials": [
    {
      "username": "string",
      "password": "string",
      "securityToken": "string"
    }
  ]
}
```

For more information:

- API: <https://api.trifacta.com/ee/es.t/index.html#operation/createConnection>
- UI: *Salesforce Connections*

## Apply overrides to file-based data sources when you run a job

You can now apply overrides to your file-based data sources when you run a job using the APIs. See *API Workflow - Run Job*.

## Release 7.0

### v3 endpoints are no longer available

The v3 version of the API endpoints are no longer available in the platform.

You must use the v4 versions. See *API Reference*.

### Path to file-based imported datasets can be specified by URI

In Release 6.8 and earlier, you could create an imported dataset from a file-based source using a request like the following:

```
{
  "path": "/tri-hdfs/uploads/1/4aee9852-cf92-47a8-8c6a-9ff2adeb3b4a/POS-r02.txt",
  "type": "hdfs",
  "bucket": null,
  "name": "POS-r02b.txt",
  "description": "POS-r02 - copy"
}
```

Beginning in this release, you can specify the source using a single `uri` value:

```
{
  "uri": "hdfs:///tri-hdfs/uploads/1/4aee9852-cf92-47a8-8c6a-9ff2adeb3b4a/POS-r02.txt",
  "name": "POS-r02b.txt",
  "description": "POS-r02 - copy"
}
```

**NOTE:** The prior format using `path` and `type` is still supported.

For more information, see

<https://api.trifacta.com/ee/es.t/index.html#operation/createImportedDataset>

### awsRoles GET method removed

Prior to Release 7.0, the platform supported the ability to retrieve `awsRole` objects based on `awsRoleId` value.

In Release 7.0 and later, this endpoint has been disabled in the product.

**Workaround:** You can still acquire `awsRole` identifiers and objects via API. For more information, see <https://api.trifacta.com/ee/es.t/index.html#operation/listAwsRoles>

### Connections endpoints simplified

Prior to Release 7.0, the request to create a connection looked like the following:



```
{
  "connectParams": {
    "vendor": "sqlserver",
    "vendorName": "sqlserver",
    "host": "sqlserver.example.com",
    "port": "1433"
  },
  "host": "sqlserver.example.com",
  "port": 1433,
  "vendor": "sqlserver",
  "params": {
    "connectStrOpts": ""
  },
  "ssl": false,
  "vendorName": "sqlserver",
  "name": "sqlserver_test2",
  "description": "",
  "type": "jdbc",
  "isGlobal": false,
  "credentialType": "basic",
  "credentialsShared": true,
  "disableTypeInference": false,
  "credentials": [
    {
      "username": "<username>",
      "password": "<password>"
    }
  ]
}
```

In the above, the `connectParams` information is duplicated elsewhere in the request.

Beginning in Release 7.0, the `connectParams` elements of the `Connections` endpoint have been removed.

The above request now looks like the following:

```
{
  "host": "sqlserver.example.com",
  "port": 1433,
  "vendor": "sqlserver",
  "params": {
    "connectStrOpts": ""
  },
  "ssl": false,
  "vendorName": "sqlserver",
  "name": "sqlserver_test2",
  "description": "",
  "type": "jdbc",
  "isGlobal": false,
  "credentialType": "basic",
  "credentialsShared": true,
  "disableTypeInference": false,
  "credentials": [
    {
      "username": "<username>",
      "password": "<password>"
    }
  ]
}
```

### Required changes:

- The `params` and `name` attributes are now required when creating a new connection.

- If you have downstream scripts that utilize the `connectParams` objects, these must be updated to reference the corresponding attributes after upgrade to Release 7.0 or later.

For more information, see

<https://api.trifacta.com/ee/es.t/index.html#operation/getConnection>

## Publishing for a jobGroup has been simplified

After you have run a job, you can publish the results from the job using the following endpoint:

```
/v4/jobGroups/<id>/publish
```

Prior to Release 7.0, the request to submit to the endpoint looked like the following:

```
{
  "connection": {
    "id": 1
  },
  "path": ["default"],
  "table": "test_table",
  "action": "create",
  "inputFormat": "pqt",
  "flowNodeId": 27
}
```

However, the `flowNodeId` value is not necessary, as it can be inferred from the `id` value that is part of the endpoint URI.

Beginning in Release 7.0, the request no longer requires the `flowNodeId` value:

```
{
  "connection": {
    "id": 1
  },
  "path": ["default"],
  "table": "test_table",
  "action": "create",
  "inputFormat": "pqt"
}
```

**NOTE:** No changes are required to existing scripts. Values submitted for the `flowNodeId` or the `output ObjectId` are ignored.

For more information, see

<https://api.trifacta.com/ee/es.t/index.html#operation/publishJobGroup>

## WASB URI format has changed

In Release 6.8 and earlier, the supported format for WASB URIs was the following:

```
wasbs://blobhost.blob.core.windows.net/container/path/to/file.csv
```

This format was not an Azure standard. Beginning in Release 7.0, the following standards-based format is supported:

```
wasbs://container@blobhost.blob.core.windows.net/path/to/file.csv
```

**NOTE:** Any references in your API tooling to the previous WASB UI format must be updated to the new format.

### Saving Databricks Personal Access Token has different response

In Release 6.8 and earlier, when saving a Databricks Personal Access Token, the response looked like the following:

|                      |                  |
|----------------------|------------------|
| <b>Response Code</b> | 204 - No Content |
| <b>Response Body</b> | Empty.           |

Beginning in Release 7.0, the response looks like the following:

|                      |  |
|----------------------|--|
| <b>Response Code</b> | 200 - Ok   |
| <b>Response Body</b> | <pre>{   "message": "Databricks Personal Access Token has been stored for user &lt;user_email&gt;" }</pre> |

For more information, see

<https://api.trifacta.com/ee/es.t/index.html#operation/saveDatabricksAccessToken>

### Documentation for awsRoles object corrected

In Release 6.8 and earlier, the following issues appeared in the documentation for the awsRoles object. These issues have been corrected in this release:

#### Inadequate permissions:

The POST and PUT methods for the awsRoles endpoint require either of the following user permissions:

- Workspace admin
- Trifacta admin

#### Optional workspaceid/personid attributes:

For the POST and PUT methods for the awsRoles endpoint, you can optionally assign the role to a workspace or person.

See

<https://api.trifacta.com/ee/es.t/index.html#operation/createAwsRole>

See

<https://api.trifacta.com/ee/es.t/index.html#operation/updateAwsRole>

## Changes for Release 6.8.1

### Improvements to AssetTransfer API documentation

Before disabling or deleting a user's account, an admin may wish to transfer the user's assets to the ownership of a different user. For more information on how to transfer all assets via API, see <https://api.trifacta.com/ee/es.t/index.html#operation/transferAssets>

## Changes for Release 6.8

### Overrides to relational sources and targets

Through the APIs, you can now apply overrides to relational sources and targets during job execution or deployment import.

#### jobGroups

When you are running a job, you can override the default publication settings for the job using `overrides` in the request. For more information, see *API Workflow - Run Job*.

#### Deployments

When you import a flow package into a deployment, you may need to remap the source and output of the flow to use production versions of your data. This capability has been present in the product for file-based sources and targets. Now, it's available for relational sources and targets. For more information, see *Define Import Mapping Rules*.

### New flow object definition

Release 6.8 introduces a new version of the flow object definition. This new version will support cross-product and cross-version import and export of flows in the future. For more information see *Changes to the Object Model*.

**NOTE:** The endpoints to use to manage flow packages remain unchanged. Similarly, the methods used to define import mapping rules remains unchanged. The API responses that contain flow definitions has changed. See below.

### Export and import macros

Beginning in Release 6.8, you can export and import macro definitions via API.

- **Export:** <https://api.trifacta.com/ee/es.t/index.html#operation/getMacroPackage>
- **Import:** <https://api.trifacta.com/ee/es.t/index.html#operation/importMacroPackage>

## Release 6.4

### Request format for assigning connection permissions endpoint has changed

For this endpoint:

```
/v4/connections/<cid>/permissions/
```

where:

- <cid> is in the internal connection identifier.

The request payload format has changed.

Before Release 6.4, the request format was as follows:

```
[
  {
    "personId": 3,
    "role": "readOnly"
  }
]
```

Beginning in Release 6.4, the request format is as follows:

```
{
  "data": [
    {
      "person": {
        "id": 3
      },
      "role": "readOnly"
    }
  ]
}
```

**NOTE:** The old request format is no longer supported.

For more information, see <https://api.trifacta.com/ee/es.t/index.html#operation/createConnectionPermission>

## v4 version of password reset request endpoint

To assist in migration from the command-line interface to using the APIs, a v4 version of an API endpoint has been made available to allow for administrators to generate password reset codes.

## Changes to awsConfig object

**NOTE:** No action is required.

In Release 6.0, the awsConfig object was introduced to enable the assignment of AWS configurations to individual users (per-user auth) via API. This version of the awsConfig object supported a mapping of a single IAM role to an awsConfig object.

Beginning in Release 6.4, per-user authentication now supports mapping of multiple possible IAM roles to an individual user's configuration. To enable this one-to-many mapping, the awsRoles object was introduced.

- An awsRoles object creates a one-to-one mapping between an IAM role and an awsConfig object.
- An awsConfig object can have multiple awsRoles assigned to it.

Changes to awsConfig object:

- The `role` field in the object has been replaced by `activeRoleId`, which maps to the active role for the configuration object.
- For each `role` reference in the awsConfig objects, a corresponding awsRole object has been created and mapped to it.

Beginning in Release 6.4, you can create, edit, and delete `awsRoles` objects, which can be used to map an AWS IAM role ARN to a specified AWSConfig object. You can map multiple `awsRoles` to a single `awsConfig`.

For more information, see *API Workflow - Manage AWS Configurations*.

## Release 6.3

### Assign AWSConfigs to a user at create time

Beginning in Release 6.3, you can assign an AWSConfig object to a user when you create the object. This shortcut reduces the number of REST calls that you need to make.

**NOTE:** For security reasons, AWSConfig objects must be assigned to users at the time of creation. Admin users can assign to other users. Non-admin users are automatically assigned the AWSConfig objects that they create.

Prior to Release 6.3, AWSConfig objects were assigned through the following endpoint. Example:

```
/v4/people/2/awsConfigs/6
```

**NOTE:** This endpoint has been removed from the platform. Please update any scripts that reference the above endpoint to manage AWS configuration assignments through the new method described in the following link.

See *API Workflow - Manage AWS Configurations*.

## Release 6.0

### Error in Release 6.0.x API docs

In Release 6.0 - Release 6.0.2, the online and PDF versions of the documentation referenced the following endpoint: API JobGroups Get Status v4. According to the docs, this endpoint was triggered in this manner:

|          |  |
|----------|--|
| Method   | GET  |
| Endpoint | <code>/v4/jobGroups/&lt;id&gt;/status</code> |

This endpoint exists in v3 of the API endpoints. It does not exist in v4.

Instead, you should monitor the `status` field for the base GET endpoint for jobGroups. For more information, see <https://api.trifacta.com/ee/es.t/index.html#operation/getJobGroup>

### Planned End of Life of v3 API endpoints

In Release 6.0, the v3 API endpoints are supported.

**In the next release of Trifacta Self-Managed Enterprise Edition after Release 6.0, the v3 API endpoints will be removed from the product (End of Life).**

You must migrate to using the v4 API endpoints before upgrading to the next release after Release 6.0.

## Release 5.9

### Introducing Access Tokens

Each request to the API endpoints of the Trifacta platform requires submission of authentication information. In Release 5.1 and earlier:

- A request could include clear-text username/password combinations. This method is not secure.
- A request could include a browser cookie. This method does not work for well for use cases outside of the browser (e.g. scripts).

Beginning in Release 5.9, API users can manage authentication using access tokens. These tokens obscure any personally identifiable information and represent a standards-based method of secure authentication.

**NOTE:** All previous methods of API authentication are supported in this release. Access tokens is the preferred method of authentication.

The basic process works in the following manner:

1. API user requests generation of a new token.
  - a. This initial request must contain a valid username and password.
  - b. Request includes expiration.
  - c. Token value is returned in the response.
2. The token value inserted into the Authorization header of each request to the platform.
3. User monitors current time and expiration time of the token. At any time, the user can request a new token to be generated using the same endpoint used in the initial request.

Access tokens can be generated via API or the Trifacta application.

**NOTE:** This feature must be enabled in your instance of the platform. See *Enable API Access Tokens*.

- **API:** For more information, see <https://api.trifacta.com/ee/es.t/index.html#operation/createApiAccessToken>
- **Trifacta application:** For more information, see *Access Tokens Page*.

For more information on API authentication methods, see <https://api.trifacta.com/ee/es.t/index.html#section/Authentication>

## Release 5.1

None.

## Release 5.0

### Introducing v4 APIs

**NOTE:** This feature is in Beta release.

Release 5.0 signals the introduction of version 4 of the REST APIs.

**NOTE:** At this time, a very limited number of v4 REST APIs are publicly available. Where possible, you should continue to use the v3 endpoints. For more information, see *API Reference*.

v4 conventions

The following conventions apply to v4 and later versions of the APIs:

- Parameter lists are consistently enveloped in the following manner:

```
{  "data": [    {      ...    }  ]}
```

- Field names are in camelCase and are consistent with the resource name in the URL or with the embed URL parameter.
- From early API versions, foreign keys have been replaced with identifiers like the following:

| v3 and earlier             | v4 and later                         |
|----------------------------|--------------------------------------|
| <pre>"createdBy": 1,</pre> | <pre>"creator": {  "id": 1  },</pre> |
| <pre>"updatedBy": 2,</pre> | <pre>"updater": {  "id": 2  },</pre> |

- Publication endpoint references database differently. This change is to make the publishing endpoint for relational targets more flexible in the future.

| v3 and earlier                   | v4 and later                     |
|----------------------------------|----------------------------------|
| <pre>"database": "dbName",</pre> | <pre>"path": [ "dbName" ],</pre> |

Release 4.2

Create Hive and Redshift connections via API

You can create connections of these types via API:

- Only one global Hive connection is still supported.
- You can create multiple Redshift connections.

WrangledDataset endpoints are still valid

In Release 4.1.1 and earlier, the WrangledDataset endpoints enabled creation, modification, and deletion of a wrangled dataset object, which also created the associated recipe object.

In Release 4.2, wrangled datasets have been removed from the application interface. However, the WrangledDataset API endpoints remain, although they now apply to the recipe directly.

The following endpoints and methods are still available:

**NOTE:** In a future release, these endpoints may be migrated to recipe-based endpoints. API users should review this page for each release.



For more information, see *Changes to the Object Model*.

# Changes to Configuration

## Contents:

- *Release Updates*
    - *Release 8.11*
    - *Release 8.2*
    - *Release 8.0*
    - *Release 7.10*
    - *Release 7.8*
    - *Release 7.6*
    - *Release 7.1*
    - *Release 6.5*
    - *Release 6.4.1*
    - *Release 6.4*
    - *Release 6.0*
  - *Configuration Mapping*
- 

To assist administrators in managing configuration, these section provides a per-release set of updates to configuration.

- Platform settings appear in the Admin Settings page. From the left nav bar, select **User menu > Admin console > Admin settings**. See *Admin Settings Page*.
  - For more information on configuration in general, see *Platform Configuration Methods*.
- Settings pertaining to the workspace environment and its users are available through a different page. From the left nav bar, select **User menu > Admin console > Workspace settings**. See *Workspace Settings Page*.

## Release Updates

### Release 8.11

#### Sample sizes can be configured to be up to 40MB in size

Beginning in this release, sample sizes can be configured to be up to 40MB in size.

**NOTE:** The Trifacta platform now permits creation and storage of samples up to 40 MB in size. This size is four times the size of samples in previous releases. Details on how to change this size are below.

#### Changes in configuration:

- New settings allow for:
  - setting the default and maximum size of sample displayed in browser
    - users can override the actual size of the sample downloaded to their browser based on their own experience
  - setting the maximum size of samples stored on the backend
- Previous settings on configuring sample size have been deprecated.

#### New settings:

---

| Setting                        | Default value | Description and Notes  |
|--------------------------------|---------------|--|
| webapp.client.defaultLoadLimit | 10485760      | <p>Sets the default maximum number of bytes that can be loaded into the browser for samples. This default figure remains at 10 MB.</p> <p><b>NOTE:</b> On a per-recipe basis, users can override this setting through the Transformer page. See <i>Change Recipe Sample Size</i>.</p>  |
| webapp.client.maxLoadLimit     | 41943040      | <p>Sets the maximum number of bytes that can be loaded into the browser for samples. This limit is now set to 40 MB.</p> <p><b>NOTE:</b> This value cannot be overridden by users. Users can set the sample size in their browser up to this limit and no higher.</p>  |
| webapp.sampleOutputLimit       | 41943040      | <p>Sets the requested size of samples in bytes that are stored in the base storage layer for each sample. This limit is now set to 40 MB.</p> <p><b>NOTE:</b> This parameter defines the storage size of samples on the backend. Default storage size is four times larger than in previous releases.</p> <p><b>NOTE:</b> For datasources that must be decompressed or converted during ingest, the actual storage volume may be larger than this limit.</p> |

#### Old settings:

**NOTE:** In most environments, these settings were left at their default values. They are now deprecated.

| Setting                                  | Default value | Description and Notes  |
|--|---------------|--|
| webapp.client.loadLimit                  | 10485760      | <p>Defined the number of bytes that could be downloaded to the web browser, which limited the size of samples.</p> <p><b>NOTE:</b> Setting is now deprecated and is no longer available in the Admin Settings page.</p>      |
| webapp.client.maxInternalFormatJsonBytes | 39321600      | <p>Defined the number of JSON bytes that could be loaded into the browser from internal objects, including samples.</p> <p><b>NOTE:</b> Setting is now deprecated and is no longer available in the Admin Settings page.</p> |

For more information, see *Configure Application Limits*.

## Release 8.2

### Photon running environment now enabled through Admin console

You can now enable or disable Photon execution through the Admin console. For more information, see *Workspace Settings Page*.

### Removed settings

The features for the following settings are now a part of the core product and can no longer be disabled. These settings have been removed.

| Settings  | Setting Description  | Notes   |
|---|--|---|
| Detect maximum column count in XLSX sheet                   | When you have enabled the Apache POI method for converting Excel files, you can enable this feature to force the conversion service to detect the maximum number of columns in an Excel sheet before beginning the conversion. | Set to true by default, so this setting was no longer needed. |
| Enable Apache POI based converter for Excel data conversion | When this setting is enabled, the conversion service uses the Java-based Apache POI converter to convert Excel data for ingestion into the product.  | Set to true by default, so this setting was no longer needed. |

## Release 8.0

### Collaboration settings apply to both flows and plans

In previous releases, the collaboration features (export, import, and sharing) for flows and plans were managed under different settings. These configuration flags to enable the following capabilities have been consolidated:

**NOTE:** When plans have been enabled, the following flags apply to both flows and plans.

| Old Setting  | New Setting | Doc for Flows       | Doc for Plans       |
|--------------|-------------|---------------------|---------------------|
| Flow export  | Export      | <i>Export Flow</i>  | <i>Export Plan</i>  |
| Flow import  | Import      | <i>Import Flow</i>  | <i>Import Plan</i>  |
| Flow sharing | Sharing     | <i>Share a Flow</i> | <i>Share a Plan</i> |

For more information on modifying these settings, see *Workspace Settings Page*.

## Release 7.10

### Ability to list users is now configurable

Beginning in this release, workspace administrators can choose whether or not individual users are able to see lists of all users in the workspace. When enabled, for example, users who are trying to share their flows can review the list of all workspace users from which they can select the users to share.

**NOTE:** Accessible information can include user email addresses. Some workspace administrators may choose to disable this feature for security reasons. Users are still able to select users if their full email addresses are known.

For more information, see *Workspace Settings Page*.

## Release 7.8

### JDBC Connection pooling disabled

Beginning in Release 7.8, JDBC connection pooling is disabled for all environments. This implementation of JDBC-based connection pools has been a source of functional and performance issues, which led to jobs failing. When JDBC connection pooling is disabled:

- The connections made at the platform level are now created on an on-demand basis. For example, if you have one SQL Server Connection created through the user interface, there may be multiple JDBC connections to your SQL Server instance if multiple jobs are executing at the same time.
- The ingest and publishing process may take a bit longer, since the connections are now made on-demand. Transformation jobs should not be affected, since they are executed on the running environment.
- The number of physical JDBC connections should more closely match the number of jobs running concurrently.

Connection pooling may be replaced in a future release.

**NOTE:** This feature is disabled for all environments, even if it had been enabled in previous releases.

For more information, see *Configure Data Service*.

## Release 7.6

### Migrated parameters

The following parameters have been migrated from the Admin Settings page (sourced from `trifacta-conf.json`) into the Workspace Settings page (sourced from Configuration Service):

| Old <code>trifacta-conf.json</code> property | New Configuration Service property | Description  |
|--|------------------------------------|--|
| <code>aws.s3.enabled</code>                  | Enable S3 Connectivity             | <p>When set to <code>enabled</code>, base S3 connectivity is possible for the workspace.</p> <div><p><b>During upgrade, this setting may be disabled in some deployments. If your pre-upgrade environment was connected to S3, please verify that this setting is enabled in the Workspace Settings Page. See <i>Workspace Settings Page</i>.</b></p></div> <p>For more information, see <i>S3 Access</i>.</p> |

For more information, see *Workspace Settings Page*.

### Use of Spark native libraries now automatic for EMR and Databricks

Prior to Release 7.6, you could configure Spark to use Spark libraries that were either 1) on the local Trifacta node or 2) available on the cluster and provided by the Spark vendor.

Beginning in Release 7.6, Spark automatically uses the vendor libraries for EMR and Databricks clusters. If you are using a Hadoop version of Spark, you can still choose to use the local libraries.

| Setting                                    | Old value                       | New Value                      |
|--|---------------------------------|--------------------------------|
| <code>spark.useVendorSparkLibraries</code> | Default is <code>false</code> . | Default is <code>true</code> . |

**NOTE:** This value is ignored for EMR and Databricks clusters.

For more information:

- *Configure for Spark*
- *Configure for EMR*
- *Configure for Azure Databricks*

## Release 7.1

### Unused parameters

| Setting   | Changed | Setting Description  | Notes  |
|---|---------|--|--|
| feature.parameterization.maxNumberOfFilesForExecution.databricksSpark | removed | Maximum number of parameterized source files that are permitted to be executed as part of an Azure Databricks job. | Method by which parameterization is now done no longer requires a union operation between source datasets, so this setting was no longer needed. |

### Migrate Data Service properties to Configuration Service

Prior to Release 7.1, a number of properties related to the Data Service have been moved from the `application.properties` configuration file on the Trifacta node into the Configuration Service, where these properties can be better managed across product versions and editions.

**NOTE:** These properties are not intended for customer modification. If you believe that you need to make changes to these property values, please contact *Alteryx Customer Success Services*.

| Old application.properties property                 | New Configuration Service property                    |
|---|---|
| com.trifacta.dataservice.oracle.whitelist           | dataservice.vendors.oracle.import.whitelist           |
| com.trifacta.dataservice.oracle.greylis             | dataservice.vendors.oracle.import.greylis             |
| com.trifacta.dataservice.oracle.ciphers             | dataservice.vendors.oracle.ciphers                    |
| com.trifacta.dataservice.postgres.greylis           | dataservice.vendors.postgres.import.greylis           |
| com.trifacta.dataservice.postgres.whitelist         | dataservice.vendors.postgres.import.whitelist         |
| com.trifacta.dataservice.databricks.whitelist       | dataservice.vendors.databricks.import.whitelist       |
| com.trifacta.dataservice.hive.whitelist             | dataservice.vendors.hive.import.whitelist             |
| com.trifacta.dataservice.glue.whitelist             | dataservice.vendors.glue.import.whitelist             |
| com.trifacta.dataservice.sqlserver.whitelist        | dataservice.vendors.sqlserver.import.whitelist        |
| com.trifacta.dataservice.sqlserver.greylis          | dataservice.vendors.sqlserver.import.greylis          |
| com.trifacta.dataservice.sqldatawarehouse.whitelist | dataservice.vendors.sqldatawarehouse.import.whitelist |
| com.trifacta.dataservice.sqldatawarehouse.greylis   | dataservice.vendors.sqldatawarehouse.import.greylis   |
| com.trifacta.dataservice.snowflake.whitelist        | dataservice.vendors.snowflake.import.whitelist        |
| com.trifacta.dataservice.snowflake.greylis          | dataservice.vendors.snowflake.import.greylis          |
| com.trifacta.dataservice.redshift.whitelist         | dataservice.vendors.redshift.import.whitelist         |
| com.trifacta.dataservice.teradata.whitelist         | dataservice.vendors.teradata.import.whitelist         |
|   |   |

|  |  |
|--|--|
| com.trifacta.dataservice.teradata.greylist                         | dataservice.vendors.teradata.import.greylist                 |
| com.trifacta.dataservice.trifacta.postgresql.fetchSize             | dataservice.vendors.postgres.import.fetchSize                |
| com.trifacta.dataservice.trifacta.redshift.fetchSize               | dataservice.vendors.redshift.import.fetchSize                |
| com.trifacta.dataservice.trifacta.oracle.fetchSize                 | dataservice.vendors.oracle.import.fetchSize                  |
| com.trifacta.dataservice.trifacta.sqlserver.fetchSize              | dataservice.vendors.sqlserver.import.fetchSize               |
| com.trifacta.dataservice.trifacta.sources.databricks.enabled       | dataservice.vendors.databricks.enabled                       |
| com.trifacta.dataservice.trifacta.sources.hive.enabled             | dataservice.vendors.hive.enabled                             |
| com.trifacta.dataservice.trifacta.sources.glue.enabled             | dataservice.vendors.glue.enabled                             |
| com.trifacta.dataservice.trifacta.sources.redshift.enabled         | dataservice.vendors.redshift.enabled                         |
| com.trifacta.dataservice.trifacta.sources.postgres.enabled         | dataservice.vendors.postgres.enabled                         |
| com.trifacta.dataservice.trifacta.sources.sqlserver.enabled        | dataservice.vendors.sqlserver.enabled                        |
| com.trifacta.dataservice.trifacta.sources.teradata.enabled         | dataservice.vendors.teradata.enabled                         |
| com.trifacta.dataservice.trifacta.sources.oracle.enabled           | dataservice.vendors.oracle.enabled                           |
| com.trifacta.dataservice.trifacta.sources.bigquery.enabled         | dataservice.vendors.bigquery.enabled                         |
| com.trifacta.dataservice.trifacta.sources.sqldatawarehouse.enabled | dataservice.vendors.sqldatawarehouse.enabled                 |
| com.trifacta.dataservice.trifacta.sources.snowflake.enabled        | dataservice.vendors.snowflake.enabled                        |
| com.trifacta.dataservice.databricks.publish.type.integer           | dataservice.vendors.databricks.publish.typeMap.integer       |
| com.trifacta.dataservice.databricks.publish.type.string            | dataservice.vendors.databricks.publish.typeMap.string        |
| com.trifacta.dataservice.databricks.publish.type.bool              | dataservice.vendors.databricks.publish.typeMap.bool          |
| com.trifacta.dataservice.databricks.publish.type.float             | dataservice.vendors.databricks.publish.typeMap.float         |
| com.trifacta.dataservice.databricks.publish.type.datetime          | dataservice.vendors.databricks.publish.typeMap.datetime      |
| com.trifacta.dataservice.hive.publish.type.integer                 | dataservice.vendors.hive.publish.typeMap.integer             |
| com.trifacta.dataservice.hive.publish.type.string                  | dataservice.vendors.hive.publish.typeMap.string              |
| com.trifacta.dataservice.hive.publish.type.bool                    | dataservice.vendors.hive.publish.typeMap.bool                |
| com.trifacta.dataservice.hive.publish.type.float                   | dataservice.vendors.hive.publish.typeMap.float               |
| com.trifacta.dataservice.hive.publish.type.datetime                | dataservice.vendors.hive.publish.typeMap.datetime            |
| com.trifacta.dataservice.redshift.publish.type.integer             | dataservice.vendors.redshift.publish.typeMap.integer         |
| com.trifacta.dataservice.redshift.publish.type.string              | dataservice.vendors.redshift.publish.typeMap.string          |
| com.trifacta.dataservice.redshift.publish.type.bool                | dataservice.vendors.redshift.publish.typeMap.bool            |
| com.trifacta.dataservice.redshift.publish.type.float               | dataservice.vendors.redshift.publish.typeMap.float           |
| com.trifacta.dataservice.redshift.publish.type.datetime            | dataservice.vendors.redshift.publish.typeMap.datetime        |
| com.trifacta.dataservice.snowflake.publish.type.integer            | dataservice.vendors.snowflake.publish.typeMap.integer        |
| com.trifacta.dataservice.snowflake.publish.type.string             | dataservice.vendors.snowflake.publish.typeMap.string         |
| com.trifacta.dataservice.snowflake.publish.type.bool               | dataservice.vendors.snowflake.publish.typeMap.bool           |
| com.trifacta.dataservice.snowflake.publish.type.float              | dataservice.vendors.snowflake.publish.typeMap.float          |
| com.trifacta.dataservice.snowflake.publish.type.datetime           | dataservice.vendors.snowflake.publish.typeMap.datetime       |
| com.trifacta.dataservice.snowflake.publish.type.time               | dataservice.vendors.snowflake.publish.typeMap.time           |
| com.trifacta.dataservice.snowflake.publish.type.date               | dataservice.vendors.snowflake.publish.typeMap.date           |
| com.trifacta.dataservice.sqldatawarehouse.publish.type.integer     | dataservice.vendors.sqldatawarehouse.publish.typeMap.integer |

|  |   |
|--|---|
| com.trifacta.dataservice.sqldatawarehouse.publish.type.string          | dataservice.vendors.sqldatawarehouse.publish.typeMap.string   |
| com.trifacta.dataservice.sqldatawarehouse.publish.type.bool            | dataservice.vendors.sqldatawarehouse.publish.typeMap.bool     |
| com.trifacta.dataservice.sqldatawarehouse.publish.type.float           | dataservice.vendors.sqldatawarehouse.publish.typeMap.float    |
| com.trifacta.dataservice.sqldatawarehouse.publish.type.datetime        | dataservice.vendors.sqldatawarehouse.publish.typeMap.datetime |
| com.trifacta.dataservice.sqldatawarehouse.publish.type.time            | dataservice.vendors.sqldatawarehouse.publish.typeMap.time     |
| com.trifacta.dataservice.azure.database.resource                       | dataservice.azure.databaseResource                            |
| com.trifacta.dataservice.trifacta.hive.forcekint                       | dataservice.vendors.hive.forcekint                            |
| com.trifacta.dataservice.trifacta.sources.snowflake.intermediateformat | dataservice.vendors.snowflake.intermediateformat              |
| com.trifacta.dataservice.teradata.disable_ordering                     | dataservice.vendors.teradata.disableOrdering                  |
| com.trifacta.dataservice.hive.publish.disableDistributedBy             | dataservice.vendors.hive.publish.disableDistributedBy         |

## Release 6.5

| Admin Settings or <code>trifacta-conf.json</code> setting | Workspace setting                     | Update   |
|---|---------------------------------------|--|
| n/a   | Show output directory in profile view | This parameter has been hidden, as it is now enabled by default. |
| n/a   | Show upload directory in profile view | This parameter has been hidden, as it is now enabled by default. |

## Release 6.4.1

Following parameter was moved in this release:

| Admin Settings or <code>trifacta-conf.json</code> setting | Workspace setting       | Update  |
|---|-------------------------|---|
| webapp.connectivity.customSQLQuery.enabled                | Enable custom SQL query | Feature is now enabled at the workspace or tier level. For more information on this feature, see <i>Enable Custom SQL Query</i> . |
| webapp.enableTypecastOutput                               | Schematized output      | Feature is now enabled at the workspace or tier level. For more information, see <i>Miscellaneous Configuration</i> .             |

## Release 6.4

Parameters from Release 6.0 that are no longer available in the Workspace Settings page are now enabled by default for all users.

Some new features for this release may be enabled or disabled through the Workspace Settings page.

See *Workspace Settings Page*.

## Release 6.0

Initial release of the Workspace Settings page. See below for configuration mapping.

## Configuration Mapping

The following mapping between old `trifacta-conf.json` settings and new Workspace settings is accurate for the current release.



| Admin Settings or <code>trifacta-conf.json</code> setting | Workspace setting   | Update  |
|---|---|---|
| <code>webapp.walkthrough.enabled</code>                   | Product walkthroughs                                      |   |
| <code>webapp.session.durationInMins</code>                | Session duration  |   |
| <code>webapp.enableDataDownload</code>                    | Sample downloads  |   |
| <code>webapp.enableUserPathModification</code>            | Allow the user to modify their paths                      |   |
| <code>webapp.enableSelfServicePasswordReset</code>        | Enable self service password reset                        |   |
| <code>webapp.connectivity.customSQLQuery.enabled</code>   | Enable custom SQL query                                   |   |
| <code>outputFormats.Parquet</code>                        | Parquet output format                                     |   |
| <code>outputFormats.JSON</code>                           | JSON output format  |   |
| <code>outputFormats.CSV</code>                            | CSV output format   |   |
| <code>outputFormats.Avro</code>                           | Avro output format  |   |
| <code>outputFormats.TDE</code>                            | TDE output format   |   |
| <code>feature.scheduling.enabled</code>                   | Enable Scheduling feature                                 |   |
| <code>feature.scheduling.schedulingManagementUI</code>    | Scheduling management UI (experimental)                   | Now hidden.   |
| <code>feature.scheduling.upgradeUI</code>                 | Show users a modal to upgrade to a plan with Scheduling   | Now hidden.   |
| <code>feature.sendFlowToOtherPeople.enabled</code>        | Allow users to send copies of flows to other users        | As of Release 6.10, this feature has been removed from the product. |
| <code>feature.flowSharing.enabled</code>                  | Enable Flow Sharing feature                               |   |
| <code>feature.flowSharing.upgradeUI</code>                | Show users a modal to upgrade to a plan with Flow Sharing | Now hidden.   |
| <code>feature.enableFlowExport</code>                     | Allow users to export their flows                         |   |
| <code>feature.enableFlowImport</code>                     | Allow users to import flows into the platform             |   |
| <code>feature.hideAddPublishingActionOptions</code>       | Forbid users to add non-default publishing actions        |   |
| <code>feature.hideUnderlyingFileSystem</code>             | Hide underlying file system to users                      |   |
| <code>feature.publishEnabled</code>                       | Enable publishing   |   |
| <code>feature.rangeJoin.enabled</code>                    | Enable UI for range join                                  |   |
| <code>feature.showDatasourceTabs</code>                   | Show datasource tabs in the application                   |   |
| <code>feature.showFileLocation</code>                     | Show file location  |   |
| <code>feature.showOutputHomeDirectory</code>              | Show output directory in profile view                     |   |
| <code>feature.showUploadDirectory</code>                  | Show upload directory in profile view                     |   |
| <code>metadata.branding</code>                            | Branding to use in-product for this deployment            | Now hidden.   |
| <code>webapp.connectivity.enabled</code>                  | Enable Connectivity feature                               |   |
| <code>webapp.connectivity.upgradeUI</code>                | Show users a modal to upgrade to a plan with Connectivity |   |
|   |   |   |

|                                 |                  |  |
|---------------------------------|------------------|--|
| feature.apiAccessTokens.enabled | API Access Token |  |
|---------------------------------|------------------|--|

# Changes to the Object Model

## Contents:

- *Release 7.1*
    - *Introducing Plans*
  - *Release 6.8*
    - *Version 2 of flow definition*
  - *Release 6.4*
    - *Macros*
  - *Release 6.0*
  - *Release 5.1*
  - *Release 5.0*
    - *Datasets with parameters*
  - *Release 4.2*
    - *Wrangled datasets are removed*
    - *Recipes can be reused and chained*
    - *Introducing References*
    - *Introducing Outputs*
    - *Flow View Differences*
    - *Connections as a first-class object*
- 

For more information on the objects available in the platform, see *Object Overview*.

## Release 7.1

### Introducing Plans

Beginning in Release 7.1, you can build plans, which are sequences of flow tasks (job executions). When a flow task completes, the next one is executed, and so on. Plans are a very useful means of automating job execution. For more information, see *Overview of Operationalization*.

See *Plans Page*.

## Release 6.8

### Version 2 of flow definition

This release introduces a new specification for the flow object.

**NOTE:** This version of the flow object now supports export and import across products and versions of those products in the future. There is no change to the capabilities and related objects of a flow.

Beginning in Release 6.8:

- You can export a flow from one product and imported it into another. For example, you can develop a flow in Trifacta and then import it into Trifacta Wrangler, assuming that the product receiving the import is on the same build or a later one.

**NOTE:** Cloud-based products, such as free Trifacta Wrangler are updated on a periodic basis, as often as once a month. These products are likely to be on a version that is later than your installed

version of Trifacta Self-Managed Enterprise Edition. For compatibility reasons, you should develop your flows in your earliest instance of Trifacta Self-Managed Enterprise Edition on Release 6.8 or later.

- You can export a flow from Release 6.8 or later of Trifacta and later import into Release 7.0 after upgrading the platform.

**NOTE:** You cannot import a pre-Release 6.8 flow into a Release 6.8 or later instance of Trifacta. You should re-import those flows before you upgrade to Release 6.8 or later.

## Release 6.4

### Macros

This release introduces **macros**, which are reusable sequences of parameterized steps. These sequences can be saved independently and references in other recipes in other flows. See *Overview of Macros*.

## Release 6.0

None.

## Release 5.1

None.

## Release 5.0

### Datasets with parameters

Beginning in Release 5.0, imported datasets can be augmented with parameters, which enables operationalizing sampling and jobs based on date ranges, wildcards, or variables applied to the input path. For more information, see *Overview of Parameterization*.

## Release 4.2

In Release 4.2, the object model has undergone the following revisions to improve flexibility and control over the objects you create in the platform.

### Wrangled datasets are removed

In Release 3.2, the object model introduced the concepts of imported datasets, recipes, and wrangled datasets. These objects represented data that you imported, steps that were applied to that data, and data that was modified by those steps.

In Release 4.2, the wrangled dataset object has been removed in place of two objects listed below. All of the functionality associated with a wrangled dataset remains, including the following actions. Next to these actions are the new object with which the action is associated.

| Wrangled Dataset action  | Release 4.2 object |
|--------------------------|--------------------|
| Run or schedule a job    | Output object      |
| Preview data             | Recipe object      |
| Reference to the dataset | Reference object   |

**NOTE:** At the API level, the `wrangledDataset` endpoint continues to be in use. In a future release, separate endpoints will be available for recipes, outputs, and references. For more information, see *API Reference*.

These objects are described below.

## Recipes can be reused and chained

Since recipes are no longer tied to a specific wrangled dataset, you can now reuse recipes in your flow. Create a copy with or without inputs and move it to a new flow if needed. Some cleanup may be required.

This flexibility allows you to create, for example, recipes that are applicable to all of your datasets for initial cleanup or other common wrangling tasks.

Additionally, recipes can be created from recipes, which allows you to create chains of recipes. This sequencing allows for more effective management of common steps within a flow.

## Introducing References

Before Release 4.2, reference datasets existed and were represented in the user interface. However, these objects existed in the downstream flow that consumes the source. If you had adequate permissions to reference a dataset from outside of your flow, you could pull it in as a reference dataset for use.

In Release 4.2, a **reference** is a link between a recipe in your flow to other flows. This object allows you to expose your flow's recipe for use outside of the flow. So, from the source flow, you can control whether your recipe is available for use.

This object allows you to have finer-grained control over the availability of data in other flows. It is a dependent object of a recipe.

**NOTE:** For multi-dataset operations such as union or join, you must now explicitly create a reference from the source flow and then union or join to that object. In previous releases, you could directly join or union to any object to which you had access.

## Introducing Outputs

In Release 4.1, outputs became a configurable object that was part of the wrangled dataset. For each wrangled dataset, you could define one or more publishing actions, each with its own output types, locations, and other parameters. For scheduled executions, you defined a separate set of publishing actions. These publishing actions were attached to the wrangled dataset.

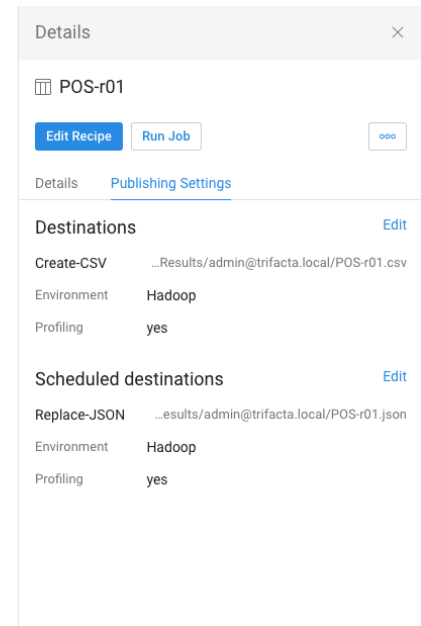
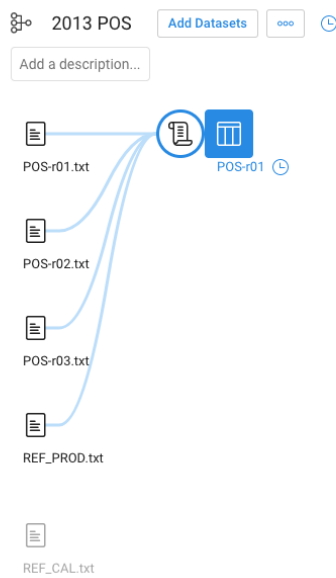
In Release 4.2, an **output** is a defined set of scheduled or ad-hoc publishing actions. With the removal of the wrangled dataset object, outputs are now top-level objects attached to recipes. Each output is a dependent object of a recipe.

## Flow View Differences

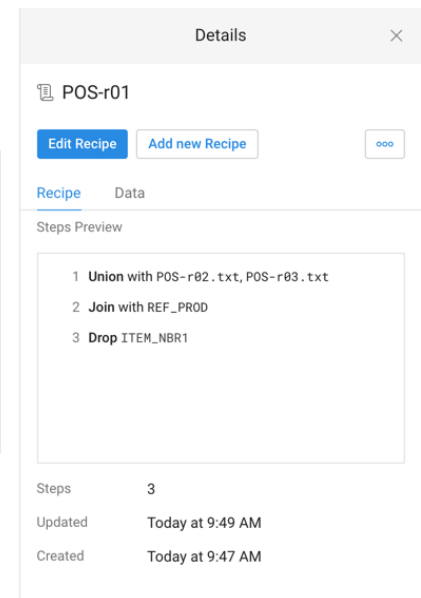
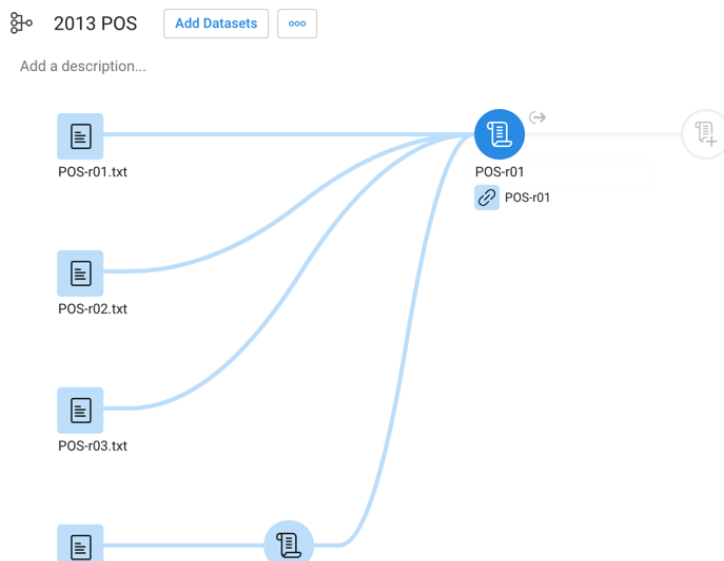
Below, you can see the same flow as it appears in Release 4.1 and Release 4.2. In each Flow View:

- The same datasets have been imported.
- POS-r01 has been unioned to POS-r02 and POS-r03.
- POS-r01 has been joined to REF-PROD, and the column containing the duplicate join key in the result has been dropped.
- In addition to the default CSV publishing action (output), a scheduled one has been created in JSON format and scheduled for weekly execution.

## Release 4.1 Flow View



## Release 4.2 Flow View



## Flow View differences

- Wrangled dataset no longer exists.
- In Release 4.1, scheduling is managed off of the wrangled dataset. In Release 4.2, it is managed through the new output object.
  - Outputs are configured in a very similar manner, although in Release 4.2, the tab is labeled, "Destinations."
  - No changes to scheduling UI.
- Like the output object, the reference object is an externally visible link to a recipe in Flow View. This object just enables referencing the recipe object in other flows.
- See *Flow View Page*.

## Other differences

- In application pages where you can select tabs to view object types, the available selections are typically: All, Imported Dataset, Recipe, and Reference.
- Wrangled datasets have been removed from the Dataset Details page, which means that the job cards for your dataset runs have been removed.
  - These cards are still available in the Jobs page when you click the drop-down next to the jjob entry.
  - The list of jobs for a recipe is now available through the output object in Flow View. Select the object and review the job details through the right panel.
- In Flow View and the Transformer page, context menu items have changed.

## Connections as a first-class object

In Release 4.1.1 and earlier, connections appeared as objects to be created or explored in the Import Data page. Through the left navigation bar, you could create or edit connections to which you had permission to do so. Connections were also selections in the Run Job page.

- Only administrators could create public connections.
- End-users could create private connections.

In Release 4.2, the Connections Manager enables you to manage your personal connections and (if you're an administrator) global connections. Key features:

- Connections can be managed like other objects.
- Connections can be shared, much like flows.
  - When a flow with a connection is shared, its connection is automatically shared.
  - For more information, see *Overview of Sharing*.
- Release 4.2 introduces a much wider range of connectivity options.
  - Multiple Redshift connections can be created through this interface. In prior releases, you could only create a single Redshift connection.

**NOTE:** Beginning in Release 4.2, all connections are initially created as **private connections**, accessible only to the user who created. Connections that are available to all users of the platform are called, public connections. You can make connections public through the Connections page.

For more information, see *Connections Page*.

# Improvements to the Type System

## Contents:

- *Release 8.11*
    - *Mismatched values are no longer published as null values in CSV outputs*
  - *Release 8.7*
    - *Snowflake date publishing improvements*
  - *Release 8.2*
  - *Release 8.1*
    - *Running environments now support output of Datetime values as Datetime/Timestamp values in Parquet files*
  - *Release 8.0*
    - *Data type inference and row split inference run on more data*
  - *Release 7.5*
    - *PII - Improved matching for social security numbers*
    - *PII - Improved and expanded matching for credit card numbers*
    - *Improved type inference for relational sources*
  - *Release 6.0 and earlier*
    - *General Improvements in Typecasting*
    - *Datetime changes*
- 

This section provides information on improvements to the Trifacta® type system.

## Release 8.11

### Mismatched values are no longer published as null values in CSV outputs

In prior releases, when a file was published in CSV format, any values that were mismatches for a column's data type were written as null values, which could lead to loss of data that was meaningful to downstream systems.

Beginning in this release, mismatched values are written out in CSV format as String values by default.

**NOTE:** The ability to write out mismatched values in CSV outputs is enabled by default in new flows and CSV publishing actions. For existing CSV outputs, the prior behavior is maintained.

**NOTE:** This capability applies only to CSV outputs at this time. In the future, it will be applied to other non-schematized outputs, such as JSON.

**Tip:** When visual profiling is enabled, you can still identify the values in the generated results that are mismatched for their column data types.

As needed, you can configure the ability to write out mismatches for CSV outputs for individual publishing actions. For more information, see *File Settings*.

## Release 8.7

### Snowflake date publishing improvements

In prior releases, the Trifacta application did not fully support Datetime values on publication. It published Date type values to Snowflake as follows:



| Publishing Action | Create/Drop Table         | Append/Truncate |                 |
|-------------------|---------------------------|-----------------|-----------------|
| TargetTable       |                           | Date            | Datetime        |
| Date              | DateTime Appends 00:00:00 | Append 00:00:00 | Append 00:00:00 |

Beginning in this release, the Trifacta application publishes Date values as follows:

| Publishing Action | Create/Drop | Append/Truncate |                           |
|-------------------|-------------|-----------------|---------------------------|
| TargetTables      |             | Date            | Datetime                  |
| Date              | Date        | Date            | Datetime Appends 00:00:00 |

For more information, see *Snowflake Data Type Conversions*.

## Release 8.2

None.

## Release 8.1

### Running environments now support output of Datetime values as Datetime/Timestamp values in Parquet files

In prior releases, Datetime values in the Trifacta application were written as String values for Parquet outputs.

Beginning in this release, you can optionally enable the generation of Datetime/Timestamp values as the outputs of Datetime columns.

**NOTE:** To ensure consistency with prior releases, this feature is disabled by default.

For more information on enabling this feature, see *Miscellaneous Configuration*.

## Release 8.0

### Data type inference and row split inference run on more data

When an dataset is imported into the Trifacta application, a larger volume of data is read from it for the following processes:

**NOTE:** The following is applied to datasets that do not contain schema information.

- **Split row inference:** Patterns in the data are used to determine the end of a row of data. When a larger volume of data is read, there should be more potential rows to review, resulting in better precision on where to split the data into separate rows in the application.
- **Type inference:** Patterns of data in the same column are used to determine the best Trifacta data type to assign to the imported dataset. A larger volume of data means that the application has more values for the same column from which to infer the appropriate data type.

**NOTE:** An increased data volume should result in a more accurate split row and data type inferencing. For pre-existing datasets, this increased volume may result in changes to the row and column type definitions when a dataset is imported.

**Tip:** For datasets that are demarcated by quoted values, you may experience a change in how columns are typed.

If you notice unexpected changes in column data types or in row splits in your datasets:

1. **Type inference:** You should move your recipe panel cursor to the top of the dataset to see if you must reassign data types.
2. **Split row inference:** Create a new imported dataset, disabling type inference in the import settings. Check the `splitrows` transform to see if it is splitting the rows appropriately. For more information, see *Import Data Page*.

### Change Trifacta application load limits

As needed, you can modify the limits that the Trifacta application uses during the data type and split row inference processes. For more information, see *Configure Type Inference*.

## Release 7.5

### PII - Improved matching for social security numbers

In prior releases, Personally Identifiable Information (PII) for social security numbers was identified based only on the length of values, which matched too broadly.

In this release, the constraints on matching of SSN values has been tightened when applied to PII.

**Tip:** PII detection is applied in generated log entries and in collaborative suggestions. When matching PII patterns are detected in data that is surface in these two areas, a mask is applied over the values for security reasons.

For more information, see *Social Security Number Data Type*.

For more information, see *Data Type Validation Patterns*.

### PII - Improved and expanded matching for credit card numbers

In prior releases, PII for credit card numbers was identified base on 16-digit values.

In this release, the matching constraints have been expanded to include 14-digit credit card values.

Also, the constraints around valid 16-digit numbers have been improved with better recognition around values for different types of credit cards. In the following table, you can see lists of valid test numbers for different credit card types and can see how detection of these values has changed between releases.

| Test Number         | Credit Card Type | Is Detected 7.4? | Is detected 7.5? |
|---------------------|------------------|------------------|------------------|
| 2222 4053 4324 8877 | Mastercard       | No               | Yes              |
| 2222 9909 0525 7051 | Mastercard       | No               | Yes              |
| 2223 0076 4872 6984 | Mastercard       | No               | Yes              |
| 2223 5771 2001 7656 | Mastercard       | No               | Yes              |
| 5105 1051 0510 5100 | Mastercard       | Yes              | Yes              |
| 5111 0100 3017 5156 | Mastercard       | Yes              | Yes              |

|                     |                  |     |     |
|---------------------|------------------|-----|-----|
| 5204 7400 0990 0014 | Mastercard       | Yes | Yes |
| 5420 9238 7872 4339 | Mastercard       | Yes | Yes |
| 5455 3307 6000 0018 | Mastercard       | Yes | Yes |
| 5506 9004 9000 0444 | Mastercard       | Yes | Yes |
| 5553 0422 4198 4105 | Mastercard       | Yes | Yes |
| 5555 5555 5555 4444 | Mastercard       | Yes | Yes |
| 4012 8888 8888 1881 | Visa             | Yes | Yes |
| 4111 1111 1111 1111 | Visa             | Yes | Yes |
| 6011 0009 9013 9424 | Discover         | Yes | Yes |
| 6011 1111 1111 1117 | Discover         | Yes | Yes |
| 3714 496353 98431   | American Express | Yes | Yes |
| 3782 822463 10005   | American Express | Yes | Yes |
| 3056 9309 0259 04   | Diners           | No  | Yes |
| 3852 0000 0232 37   | Diners           | No  | Yes |
| 3530 1113 3330 0000 | JCB              | Yes | Yes |
| 3566 0020 2036 0505 | JCB              | Yes | Yes |

For more information, see *Credit Card Data Type*.

For more information, see *Data Type Validation Patterns*.

## Improved type inference for relational sources

In prior releases, when you generated outputs, the typecasting for the outputs was determined by the data types that were inferred by the application. So, if a column contained only "Yes" or "No" values, then the application is likely to have inferred the column data type as Boolean.

The above presented problems for relational sources for the following reasons:

1. Relational sources could include schema information, which should override what the application inferred. Type inferencing can optionally be disabled for schematized sources from within the application. For more information on disabling type inference for a schematized source, see *Import Data Page*.
2. For some relational sources, datasets are ingested into the backend datastore and stored there as CSV files. These CSV files are then used as the source for the imported datasets. In these cases, the original source schema information was lost, which meant that the application's type inference was applied. This applied to the following data sources:
  - a. Snowflake
  - b. Redshift
  - c. SQL Datawarehouse
  - d. Alation
  - e. Waterline

Beginning in this release, the schemas from relational datasources that are ingested to the backend datastore are now used for generated outputs, unless the type was being forcibly set to something else during the recipe step. At the time of original import, the schema of the relational datasource is stored as part of the ingest process; this schema is stored separately.

**NOTE:** If you created recipe steps that forcibly change a column's data type from within the application to be different from the source data type of your relational source, you may need to revise these recipe steps or remove them altogether.

During publication, Trifacta maps its internal data types to the data types of the publishing target using an internal mapping per vendor. For more information, see *Type Conversions*.

## Release 6.0 and earlier

### General Improvements in Typecasting

#### Mismatched data types

Where there are mismatches between inputs and the expected input data type, the following values are generated for the mismatches:

| Source data type   | Output if mismatched                      |
|--|---|
| Primitive data types: <ul style="list-style-type: none"><li>• Integer</li><li>• Decimal</li><li>• Boolean</li><li>• Arrays</li><li>• Maps</li></ul>  | null value, if mismatched                 |
| Datetime   | null value, if mismatched                 |
| Other non-primitive data types, including: <ul style="list-style-type: none"><li>• SSN</li><li>• Phone Number</li><li>• Email Address</li><li>• Credit Card</li><li>• Gender</li><li>• IP Address</li><li>• URL</li><li>• HTTP Code</li><li>• Zip Code</li></ul> | Converted to string values, if mismatched |
| String   | Anything can be a String value.           |

State values and custom data types are converted to string values, if they are mismatched.

#### Three-value logic for null values

The Trifacta Photon running environment has been augmented to use three-value logic for null values.

When values are compared, the result can be `true` or `false` in most cases.

If a null value was compared to a null value in the Trifacta Photon running environment:

- In Release 3.0 and earlier, this evaluated to `true`.
- In Release 3.1 and later, this evaluates to an unknown (null) value.

This change aligns the behavior of the running environment with that of SQL and Hadoop Pig.

### Improved handling of null values

Assume that the column `nuller` contains null values and that you have the following transform:

```
derive value:(nuller >= 0)
```

Prior to Release 3.1, the above transform generated a column of `true` values.

In Release 3.1 and later, the transform generates a column of null values.

### More consistent evaluation of null values in ternaries

In the following example, `a_null_expression` always evaluates to a null value.

```
derive value: (a_null_expression ? 'a' : 'b')
```

In Release 3.0, this expression generated `b` for all inputs on the Trifacta Photon running environment and a null value on Hadoop Pig.

In Release 3.1 and later, this expression generates a null value for all inputs on both running environments.

**Tip:** Beginning in Release 3.1, you can use the `if` function instead of ternary expressions. Ternaries may be deprecated at some point in the future. For more information, see *IF Function*.

For example, you have the following dataset:

| MyStringCol           |
|-----------------------|
| This works.           |
| You can't break this. |
| Not broken yet.       |

You test each row for the presence of the string `can't`:

```
derive value: if(find(MyStringCol, 'can\t',true,0) > -1, true, false) as:'MyFindResults'
```

The above transform results in the following:

| MyStringCol           | MyFindResults |
|-----------------------|---------------|
| This works.           |               |
| You can't break this. | true          |
| Not broken yet.       |               |

In this case, the value of `false` is not written to the other columns, since the `find` function returns a null value. This null value, in turn, nullifies the entire expression, resulting in a null value written in the new column.

You can use the following to locate the null values:

```
derive value:isnull(MyFindResults) as:'nullInMyFindResults'
```

## Datetime changes

### Raw date and time values must be properly formatted

**NOTE:** Upgraded recipes continue to function properly. However, if you edit the recipe step in an upgraded system, you are forced to fix the formatting issue before saving the change.

Before this release, you could create a transform like the following:

```
derive value:date(2016,2,15)
```

This transform generated a column of map values, like the following:

```
{"year":"2016","month":"2","date":"15"}
```

Beginning this release, the above command is invalid, as the date values must be properly formatted prior to display. The following works:

```
derive value:dateformat(date(2016,2,15),'yyyy-MM-dd')
```

This transform generates a column of Datetime values in the following format:

```
2016-02-15
```

### Time:

Before this release:

```
derive value:time(11,34,58)
```

Prior release output:

```
{"hours":"11","minutes":"34","seconds":"58"}
```

This release:

```
derive value:dateformat(time(11,34,58),'HH-mm-ss')
```

This release's output:

```
11-34-58
```

- See *DATEFORMAT Function*.
- See *UNIXTIMEFORMAT Function*.

### Date formatting functions supports 12-hour time only if AM/PM indicator is included

Beginning in this release, the `unixtimeformat` and `dateformat` functions requires an AM/PM indicator (a) if the date formatting string uses a 12-hour time indicator (h or hh).

Valid for earlier releases:

```
derive value: unixtimeformat(myDate, 'yyyy-MM-dd hh:mm:ss') as:'myUnixDate'
```

Valid for this release and later:

```
derive value: unixtimeformat(myDate, 'yyyy-MM-dd hh:mm:ss a') as:'myUnixDate'
```

These references in recipes fail to validate in this release or later and must be fixed.

- See *DATEFORMAT Function*.
- See *Unixtimeformat Function*.

#### Un-inferable formats from dateformat and unixtimeformat functions are written as strings

If a formatting string is not a datetime format recognized by Trifacta, the output is generated as a string value.

This change was made to provide clarity to some ambiguous conditions.

#### Colon as a delimiter for date values is no longer supported

Beginning in this release, the colon (:) is no longer supported as a delimiter for date values. It is still supported for time values.

| myDateValue | Recognized?                |
|-------------|----------------------------|
| 02:03:2016  | No                         |
| 02:03:16    | Recognized as a time value |

When data such as the above is imported, it may not be initially recognized by the Trifacta application as Datetime type.

To fix, you might apply the following transform:

```
replace col:myDateValue with:'-' on:`-` global:true
```

The new column values are more likely to be inferred as Datetime values. If not, you can choose the appropriate Datetime format from the data type drop-down for the column. See *Data Grid Panel*.

# Changes to User Management

## Contents:

- *Release 8.7*
    - *Workspace admin role can be assigned*
  - *Release 8.1*
    - *Fine-grained sharing permissions on individual objects*
  - *Release 7.10*
  - *Release 7.9*
    - *Manage Users section has been deprecated*
  - *Release 7.6*
    - *New user management page*
    - *Users section of Admin Settings page is disabled*
  - *Release 7.5*
    - *Roles Management*
    - *Workspace owner role is removed*
    - *Authorization Overview*
  - *Release 7.1*
    - *Workspace admin is a super user*
    - *All upgraded Trifacta admins are now workspace admins*
    - *Admin can edit any global connection*
    - *Menu items unavailable due to account roles*
    - *Logging*
    - *Authorization changes to APIs*
- 

This section covers changes between release on the following topics:

- Authorization to the platform
  - User roles
  - Permissions of roles
- Required permissions
- Authentication methods
- User management

## Release 8.7

### Workspace admin role can be assigned

Beginning in Release 8.7, users can be assigned the workspace admin role through the Users page.

**NOTE:** The Workspace admin role enables users to manage all workspace objects and configuration in the *Workspace Settings Page*.

The Workspace admin role is separate from the Platform admin role, which allows an admin to access the *Admin Settings Page*, among other things.

For more information, see *Users Page*.



## Release 8.1

### Fine-grained sharing permissions on individual objects

Beginning in this release, you can change the permissions to a shared object for individual users. These fine-grained permissions can be assigned at the time of sharing by the object's Owner or a workspace admin. They can also be changed at a later time.

**NOTE:** Workspace-level permissions that are defined through a user's assigned roles still apply. These permissions define the maximum and default level of permissions that can be assigned when an object is shared.

**NOTE:** In this release, fine-grained sharing permissions apply to flows and connections only.

For more information, see *Overview of Sharing*.

## Release 7.10

## Release 7.9

### Manage Users section has been deprecated

All user management functions have been moved to the Workspace Users page. The following configuration items were migrated in this release:

- Enable Platform admin permission
- Specify Hadoop principal (if applicable)
- Specify Kerberos principal (if applicable)

These configuration items were the last ones that were handled through the Manage Users section. Please manage users through the Workspace Users page or the appropriate API endpoints.

- For more information, see *Users Page*.
- For more information, see *API Reference*.

## Release 7.6

### New user management page

The Workspace Users page now centralizes user management tasks in a single dedicated interface.

**Except for some uses, the Users area of the Admin Settings page is no longer needed. It has been disabled by default. See below for details.**

For more information on the new page, see *Users Page*.

You can also explore details of individual users in another new page. See *User Details Page*.

### Users section of Admin Settings page is disabled

Except for the following situations, all user management functions of the Admin Settings page have been migrated to the Workspace Users page. The Users section should be re-enabled for the following situations:

- Configure user principal values for integration with:
  - Enterprise SSO
  - A Hadoop cluster
  - Kerberos security
- Configure user accounts for the Trifacta admin role

**NOTE:** The Trifacta admin role is a super-user role across the entire platform. It should be assigned to a small number of accounts.

## Release 7.5

### Roles Management

Beginning in Release 7.5, workspace administrators can create and assign roles to workspace users. Each role contains zero or more privileges.

- A **role** is a set of privileges that you can assign to workspace users. Workspace users may have one or more roles.

**NOTE:** Each current user or newly created user is automatically assigned the `default` role, which grants a set of privileges for all governed object types in the workspace.

- A **privilege** is a level of access to a type of user-defined workspace object, such as flows.

Roles are created and assigned through the Roles page in the Admin console. For more information, see *Roles Page*.

### Workspace owner role is removed

As of Release 7.5, the workspace owner role has been removed from user access.

The privileges of this role have been collapsed into the `admin` role for workspaces, which has full capability to administer the workspace.

### Authorization Overview

Workspace roles govern access to workspace objects.

Platform roles govern access to platform capabilities.

For more information on these distinctions, see *Overview of Authorization*.

## Release 7.1

Release 7.1 introduces **role-based access controls (RBAC)**, in which access to Trifacta resources are managed at finer-grained levels. This release introduces the basic RBAC framework and the following key changes.

**NOTE:** Over the next few releases, additional capabilities will be added to the basic RBAC framework, enabling administrators to provide better and more closely defined access to objects. Check back to this section with each upgrade.

### Workspace admin is a super user

Beginning in Release 7.1, the workspace admin is a super-user of the product.

**NOTE:** In this release, the workspace admin user has owner access to user-created objects, such as flows and connections, within the workspace.

A **workspace** is a set of users and their objects, such as flows and connections. For more information, see *Workspace Admin Permissions*.

### All upgraded Trifacta admins are now workspace admins

**NOTE:** If you are upgrading Trifacta, any Trifacta admin users are now workspace admin users. A single workspace is supported in your instance of Trifacta. Additional workspaces are not supported.

**NOTE:** Any user who is granted the admin role is also granted the workspace admin role, which enables owner-level access to user-created objects in the workspace.

### Admin can edit any global connection

After an administrator has made a connection global (available to all users):

- Any administrator can edit the connection.
- All users can use the connection (existing functionality)
- The connection cannot be made private again (existing functionality). Connection must be deleted and recreated.

### Menu items unavailable due to account roles

Beginning in this release, menu items may not be displayed to specific users because of their current role assignments.

**NOTE:** This behavior had existed in previous releases. In this release and future releases, workspace admins may receive inquiries about menu option availability. A user's assigned roles could be a likely source for why a menu option is not available to the user.

## Logging

Logs from the authorization service may provide insight into access problems. These logs are available for download through the support bundle. For more information, see *Support Bundle Contents*.

## **Authorization changes to APIs**

Some API endpoints now include information that is specific to the changes in this release for authorization. See *Changes to the APIs*.

# Release Notes 9.2

## Contents:

- *Release 9.2*
    - *What's New*
    - *Changes in System Behavior*
    - *Deprecated*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 9.1*
    - *What's New*
    - *Changes in System Behavior*
    - *Deprecated*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 9.0*
    - *What's New*
    - *Changes in System Behavior*
    - *Deprecated*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 8.11*
    - *What's New*
    - *Changes in System Behavior*
    - *Deprecated*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 8.10*
    - *What's New*
    - *Changes in System Behavior*
    - *Deprecated*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 8.9*
    - *What's New*
    - *Changes in System Behavior*
    - *Deprecated*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 8.8*
    - *What's New*
    - *Changes in System Behavior*
    - *Deprecated*
    - *Key Bug Fixes*
    - *New Known Issues*
- 

## Release 9.2

August 3, 2022

### What's New

### Connectivity:

- Support for Google Sheets connections using OAuth2 authentication.

**NOTE:** This connection type is not supported for integration with a Hadoop-based running environment.

- OAuth 2.0 authentication must be enabled. For more information, see *Enable OAuth 2.0 Authentication*.
- You must also create an OAuth 2.0 client for the Trifacta application. See *OAuth 2.0 for Google Sheets*.
- For more information on creating the connection object, see *Google Sheets Connections*.

### Connectivity:

- Support for Google Analytics connections using OAuth2 authentication.
  - OAuth 2.0 authentication must be enabled. For more information, see *Enable OAuth 2.0 Authentication*.
  - You must also create an OAuth 2.0 client for the Trifacta application. See *OAuth 2.0 for Google Analytics*.
  - For more information on creating the connection object, see *Google Analytics Connections*.

### Lock/unlock column data type:

You can now lock or unlock a column's data type. When the data type is locked, the Trifacta application no longer attempts to infer the column's data type when subsequent recipe steps are applied.

**Tip:** You can unlock the individual 's column data type through column menu. To the left of the column name, you can click the icon and select **Automatically update** to change the column's data type. For more information, see *Column Menus*.

**Tip:** As an early step in your recipe, you can use the Advanced column selector in the Change column data type transformation to specify locking of the data types for all columns.

For more information, see *Change Column Data Type*.

### Install:

- Support for Dockerized installs for AWS and Azure deployments.

**NOTE:** You cannot upgrade from an Dockerized on-premises installation to Dockerized AWS or Azure. You must perform a fresh install.

For more information, see *Install for Docker*.

## Changes in System Behavior

### Set column data type transformation locks the column's type by default:

Starting in this release, the column data type is locked by default when you change the column data type.

**NOTE:** This change in behavior does not affect recipe steps that were defined before this release. Column data types continue to be re-inferred after those recipe steps. For those steps, you can edit them and mark them as locking the data type, if preferred.

If required, you can unlock the column's data type. For more information, see *Change Column Data Type*.

### Generate an initial sample:

When generating an initial sample from a set of files in a directory, the maximum number of files that can be read is now limited to 50.

- Previously, the Trifacta application read files until either 10MB of data or all matching files had been scanned.
- This change is to limit the number of files that must be read for various operations in the Transformer page. It only applies to generating the initial sample type. Other sampling types, such as random sample, can scan the full set of files.

As needed, an administrator can change this maximum limit.

- For more information, see *Workspace Settings Page*.

For more information on sampling, see *Overview of Sampling*.

### Email notifications:

In a future release, the setting for email notifications based on job success will default to **Default (Any Jobs)** at the project or workspace level and at the flow level. This change means that the user who executes a job and others who have access to the flow receive, by default, an email notification whenever a job executes for flows where email notification settings have never been modified. As part of this change, each email will contain a richer set of information about the job that was executed.

If needed, this new default setting can be modified: Individual users can override these settings for individual flows. For more information, see *Manage Flow Notifications Dialog*.

### New Azure Key Vault secret permission:

If you are deploying the Trifacta platform into Azure and are using an Azure Key Vault, the Recover secret permission is now strongly recommended. Soon, this secret permission will be required for all Azure key vaults. For more information, see *Configure Azure Key Vault*.

### Deprecated

### Hortonworks support:

The Hortonworks Data Platform (HDP) is no longer supported. Please switch to using a supported version of Cloudera (CDH or CDP). For more information, see *Product Support Matrix*.

**NOTE:** As part of the above deprecation, the following Hortonworks-only integrations are no longer available for the Trifacta platform:

- Hive Warehouse Connector
- High Availability for Hive
- Zookeeper Quorum

### High Availability on AWS:

- In prior releases, documentation was published for installing the Trifacta platform in a highly available environment on AWS using an EKS (Kubernetes) cluster.
- This method of high availability deployment is no longer supported. For more information, see *End of Life and Deprecated Features*.
- You can still deploy an HA installation in AWS using the standard on-premises method. For more information, see *Install for High Availability*.

### Custom data types from dictionary files:

- Custom types via dictionary are no longer supported.

**Tip:** As a replacement, you can create a custom type using regular expressions. This method is suitable for data types that follow a describable pattern or have relatively few unique values. For more information, see *Create Custom Data Types Using RegEx*.

**Tip:** For custom data types that contain a large number of discrete values, you can import your dictionary file as an imported dataset and then perform a left join between the column to validate and the imported dataset to determine if the column values are valid. For more information, see *Validate Column Values against a Dataset*.

For more information, see *End of Life and Deprecated Features*.

### Planned deprecation

In a subsequent release of Trifacta Self-Managed Enterprise Edition:

- Support for Java 8 will be deprecated. Customers must migrate to using Java 11. Additional instructions will be provided.

**Tip:** For Release 9.2, Java 11 is supported at runtime only.

- Java 11 requires Spark 3.x. When Java 8 is deprecated, support for Spark 2.x will be deprecated. Customers must migrate to using Spark 3.x. Additional instructions will be provided.
- These changes have the following implications:
  - Cloudera clusters do not support Spark 3.x. Customers using these running environments must migrate to Cloudera Data Platform.
  - Some deployments of EMR can migrate to using Spark 3.x in this release. For more information, see *Configure for EMR*.
  - Some deployments of Databricks can migrate to using Spark 3.x in this release. For more information:
    - *Configure for AWS Databricks*
    - *Configure for Azure Databricks*
- For more information on these changes, please contact *Alteryx Support*.



## Key Bug Fixes

| Ticket   | Description  |
|----------|--|
| TD-70522 | Cannot import converted files such as Excel, PDF, or JSON through SFTP connections.                                    |
| TD-69279 | Test Connection button fails a ValidationFailed error when editing a working connection configured with SSH tunneling. |
| TD-69004 | Patch httpd to version 2.4.52  |
| TD-66185 | Flatten transformation cannot handle multi-character delimiters.   |

## New Known Issues

None.

## Release 9.1

March 10, 2022

## What's New

### Databricks:

Support for Databricks 10.x using Spark 3.2.0.

**NOTE:** Some versions of Databricks have been deprecated. See below.

- For more information, see *Configure for AWS Databricks*.
- For more information, see *Configure for Azure Databricks*.

### Job execution:

The Trifacta application can check for changes to your dataset's schemas before jobs are executed and optionally halt job execution to prevent data corruption.

**Tip:** Schema validation can be overridden for individual jobs. For more information, see *Run Job Page*.

- For more information, see *Overview of Schema Management*.

### Sample Job IDs:

When a sample is collected, a job ID is generated and displayed in the Trifacta application. These job IDs enable you to identify the sample jobs.

- For more information, see *Generate a Sample*.
- For more information, see *Samples Panel*.
- For more information, see *Sample Jobs Page*.

### Import:

For long-loading Parquet datasets, you can monitor the ingest process as you continue your work.

**NOTE:** This feature is in Beta release.

**NOTE:** This feature may require enablement in your deployment. For more information, see *Configure JDBC Ingestion*.

For more information, see *Flow View Page* .

**Changes in System Behavior**

**Snowflake execution:**

Support for splitting columns on multiple delimiters during Snowflake pushdown execution. See *Snowflake Running Environment*.

**Nginx:**

Upgraded to Nginx 1.20.1.

**Deprecated**

**Databricks:**

**NOTE:** Supported for Databricks 8.3 has been deprecated. You can downgrade to Databricks 7.3 LTS or upgrade to Databricks 10.

If you have enabled user mode for ADLS Gen2, you must upgrade to Databricks 10. See *ADLS Gen2 Access*.

For more information, see *End of Life and Deprecated Features*.

**NOTE:** Support for Databricks 5.5 LTS has been deprecated. For more information, see *End of Life and Deprecated Features*.

**Key Bug Fixes**

| Ticket   | Description  |
|----------|--|
| TD-60881 | For ADLS datasets, parameter indicators in Flow View are shifted by one character. |

## New Known Issues

None.

## Release 9.0

January 31, 2022

## What's New

### Snowflake Running Environment:

For flows whose data sources and outputs are hosted in Snowflake, you can now push down execution of transformation jobs to Snowflake for much faster execution. For more information, see *Snowflake Running Environment*.

### Connectivity:

Build connections to accessible REST API endpoints.

**This feature is disabled by default. For more information about enabling REST API connectivity in your environment, please contact *Alteryx Support*.**

For more information, see *REST API Connections*.

### Dataset Schema Refresh:

You can now refresh your imported datasets with the current schema information from the source file or table. Schema refresh enables you to capture any changes to the columns in your dataset.

- For more information, see *Overview of Schema Management*.
- Dataset schemas can be refreshed through the following pages:
  - See *Library Page*.
  - See *Dataset Details Page*.
  - In Flow View. For more information, see *View for Imported Datasets*.

## Changes in System Behavior

None.

## Deprecated

None.

## Key Bug Fixes

| Ticket   | Description   |
|----------|---|
| TD-68162 | Flow parameters cannot be displayed or edited in the Transformer page and cannot be embedded in recipe steps. |

## New Known Issues

None.

## Release 8.11

January 10, 2022

### What's New

#### Performance:

- Improved performance during design time through intelligent caching of recipe steps.

**NOTE:** This feature is in Beta release.

- A workspace administrator may need to enable this feature in your workspace. See *Workspace Settings Page*.
- Improvements in job execution performance, due to skipping some output validation steps for file-based outputs.

**NOTE:** When scheduled or API jobs are executed, no validations are performed of any writesettings objects. Issues with these objects may cause failures during transformation or publishing stages of job execution.

- A workspace administrator can disable this feature if needed. See *Workspace Settings Page*.

#### EMR:

- Updated supported versions of EMR to address log4j2 issues:
  - EMR 6: Please upgrade to EMR 6.2.1.
  - EMR 5: Please upgrade to EMR 5.30.2.
  - For more information, see <https://docs.trifacta.com/display/PUB/Trifacta+Alert+TD-67372+-+0-Day+Exploit+in+log4j2+for+Self+Managed>

### Changes in System Behavior

#### Sample sizes can be increased up to 40MB

Prior to this release, the size of a sample was capped at 10MB. This size represented:

- the actual size of the sample object stored in the base storage layer
- the default maximum size of the sample displayed in the Trifacta application. This sample size can be reduced from 10MB, if needed.

Beginning in this release:

- The actual size of the stored sample has increased to 40MB.

**NOTE:** On backend storage, sample sizes are now four times larger than in previous releases. For datasources that require decompression or conversion, actual storage sizes may exceed this 40 MB limit.

**NOTE:** These limits are configurable. See *Changes to Configuration*.

- The size of the sample displayed for a recipe can be configured to be up to 40MB in size by individual users.

For more information, see *Change Recipe Sample Size*.

### Data type mismatches can now be written out in CSV format

Beginning in this release, for CSV outputs mismatched values are written as regular values by default. In prior releases, mismatched values were written as null values in CSV outputs.

See *Improvements to the Type System*.

## Deprecated

### Embedded onboarding tour is deprecated

In prior releases, Trifacta Self-Managed Enterprise Edition included a tour that walked users through the application. When the tour was enabled, it was accessible through the Home page.

This tour has been superseded by other resources, including tours available through in-app messaging, example flows and datasets, and training.

For more information, see *End of Life and Deprecated Features*.

## Key Bug Fixes

None.

## New Known Issues

| Ticket   | Description  |
|----------|--|
| TD-68162 | <p>Flow parameters cannot be displayed or edited in the Transformer page and cannot be embedded in recipe steps.</p> <div><b>Workaround:</b> To edit your flow parameters, select Parameters from the Flow View context menu.</div> <div><b>NOTE:</b> There is no current workaround for embedding in recipe steps. While your existing parameters should continue to work at execution time, avoid changing names of your flow parameters or editing recipe steps in which they are referenced. New flow parameters cannot be used in recipes at this time.</div> |

## Release 8.10

November 22, 2021

### What's New

#### Session Management:

You can view the current and recent sessions for your account in the Trifacta application. As needed, you can revoke any unfamiliar devices or sessions. For more information, see *Sessions Page*.

#### Changes in System Behavior

##### Ingestion:

Maximum permitted record length has been increased from 1 MB to 20 MB. For more information, see *Working with JSON v2*.

**NOTE:** This new default limit can be modified if you are experiencing performance issues. See *Configure Application Limits*.

##### Publishing:

Improvements to publishing of Trifacta Date values to Snowflake. For more information, see *Improvements to the Type System*.

##### Split transformation:

When splitting a column based on positions, the positions no longer need to be listed in numeric order. See *Changes to the Language*.

### Deprecated

None.

### Key Bug Fixes

None.

### New Known Issues

None.

## Release 8.9

October 26, 2021

### What's New

#### Plans:

- Create plan tasks to delete files and folders from file-based backend storage.

For more information, see *Create Delete Task*.

- You can now reference output metadata from within your plans. See *Plan Metadata References*.

#### **Collaboration:**

- You can view the list of collaborators and their corresponding avatars on shareable objects, such as Flows, Plans, and Connections pages.
  - For more information, see *Flows Page*.
  - For more information, see *Connections Page*.
  - For more information, see *Plans Page*.

#### **Sampling:**

- Adjust the size of samples loaded in the browser for your current recipe to improve performance and address low-memory conditions. See *Change Recipe Sample Size*.

#### **Changes in System Behavior**

None.

#### **Deprecated**

None.

#### **Key Bug Fixes**

None.

#### **New Known Issues**

None.

#### **Release 8.8**

October 12, 2021

#### **What's New**

##### **Flow View:**

- Updated context panel for output objects. For more information, see *View for Outputs*.

#### **Changes in System Behavior**

##### **Import:**

Improvements have been made in how double quotes are handled in CSV files during import to align Trifacta with other systems that support CSV import.

- Example values in source CSV file:

```
""My product""",In stock,"16,000",0.05
```

Note that the value 16,000 must be double-quoted, since the value contains a comma, which is the field delimiter.

- Previously, this value appeared in the Transformer page in columns as the following:

| c1                   | c2       | c3       | c4   |
|----------------------|----------|----------|------|
| " " "My product" " " | In stock | "16,000" | 0.05 |

- As of this version, the Trifacta application handles the values in a better manner when displaying them in the Transformer page:

| c1           | c2       | c3     | c4   |
|--------------|----------|--------|------|
| "My product" | In stock | 16,000 | 0.05 |

- c1:** Escaped values (triple double-quotes) in the source no longer render in the application as triple double-quotes and are represented as quoted values.
- c3:** Note that the double quotes in c3 have been stripped. Leading and trailing quotes are trimmed if the quotes are balanced within a cell.

**NOTE:** This change in behavior applies only to newly created imported datasets sourced from a CSV file. Existing imported datasets should not be affected. However, if a newly imported dataset is transformed by a previously existing recipe that compensated for the extra quotes in the Transformer page, the effects on output data could be unpredictable. These recipes and their steps should be reviewed.

This change does apply to any newly imported dataset sourced from CSV and may cause the data to change. For example, if you export an older flow and import into a new workspace or project, this change in parsing behavior applies to the datasets that are newly created in the new environment. Recipes may require review upon import.

- When results are generated in CSV, output files should continue to reflect the formatting of the source data before import. See above.

**Tip:** You can also choose the Include quotes option when creating a CSV output.

- When profiling is enabled, values that appear in CSV as " " are now marked as missing.

## Deprecated

None.

## Key Bug Fixes

None.

## New Known Issues

| Ticket   | Description  |
|----------|--|
| TD-63974 | In imported datasets sourced from CSV files, double quotes that are escaped with a backslash ( \ "backslash-escaped value\ " ) can cause remainder of row to be compressed into a single cell. |



# Release Notes 8.7

## Contents:

- *Release 8.7*
    - *What's New*
    - *Changes in System Behavior*
    - *Deprecated*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 8.6*
    - *What's New*
    - *Changes in System Behavior*
    - *Deprecated*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 8.5*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 8.4*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 8.3*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
- 

## Release 8.7

November 12, 2021

### What's New

#### EMR:

- Updated supported versions of EMR to address log4j2 issues. Please upgrade to EMR 5.30.2.
  - For more information, see <https://docs.trifacta.com/display/PUB/Trifacta+Alert+TD-67372+-+0-Day+Exploit+in+log4j2+for+Self+Managed>

#### Install:

- Support for CentOS/RHEL 8.4 installation. See *Product Support Matrix* .

#### Databases:

- Support for PostgreSQL 11 on Azure.

**NOTE:** PostgreSQL 12 is not supported in Azure at this time. Please install PostgreSQL 11. You can modify the installation commands for PostgreSQL 12 referenced in the database installation documentation to use PostgreSQL 11. For more information, see *Install Databases*.

- Support for PostgreSQL 12 for all other deployments.

**NOTE:** Support for PostgreSQL 9.6 has been deprecated. Unless you are installing the Trifacta platform in Azure, please install PostgreSQL 12. Azure deployments require PostgreSQL 11 for this release.

- See *Product Support Matrix*.

### Browsers:

- Update to supported browsers:
  - Mozilla Firefox is generally supported.
  - Microsoft Edge is now supported.

**NOTE:** This feature is in Beta release.

- New versions of supported browsers are now supported.
- For more information, see *Browser Requirements*.

### Plans:

- Create plan tasks to deliver messages to a specified Slack channel.

For more information, see *Create Slack Task*.

### Single Sign-On:

- If using SAML integration, you can now configure the security algorithm for the Trifacta application to use. See *Configure SSO for SAML*.

### Connectivity:

- Access to S3 is now managed using the native AWS SDK.

### Connectivity:

- Append Trifacta user identifiers to SQL queries to enable auditing through your database logs. For more information, see *Configure Connectivity*.

### Parameterization:

- For ADLS Gen2 storage, you can now parameterize the user info, host name, and path value fields as separate parameters.
  - For more information, see *Overview of Parameterization*.
  - For more information, see *ADLS Gen2 Access*.

### Recipe panel:

- Some enhancements to flagging steps for review. See *Flag for Review*.

## Azure:

- Support for installation on Azure Gov Cloud.

**NOTE:** The Azure environment must be set to `US_GOV`.

For more information, see *Configure for Azure*.

## Databricks:

- Support for Databricks 7.x and 8.x.

**NOTE:** Databricks 7.3 and Databricks 8.3 are recommended.

- Support for Databricks cluster creation via cluster policies.
- Store a user-defined set of secret information such as credentials in Databricks Secrets.
  - For more information, see *Configure for Azure Databricks*.
  - For more information, see *Configure for AWS Databricks*.

## Performance:

- Improved performance when previewing or performing quick scan samples on Parquet files.

**NOTE:** This feature may require enabling in your environment.

See *Configure Photon Running Environment*.

## Logging:

- Format of webapp access logs can now be configured like other Java-based platform services. For more information, see *Configure Logging for Services*.

## Trifacta node:

- NodeJS upgraded to 14.17.5.

## Changes in System Behavior

None.

## Deprecated

### HDInsight no longer supported:

- HDInsight 3.5 and 3.6 are no longer supported.

### HDP 3.0 deprecated:

- Please upgrade to HDP 3.1.

**NOTE:** In a future release, support for Hortonworks Data Platform (HDP) will be deprecated. Please migrate to using a different supported running environment. For more information, see *Product Support Matrix*.

## API:

- Deprecated API endpoint to transfer assets between users has been removed from the platform. This endpoint was previously replaced by an improved method of transfer.
- Some connection-related endpoints have been deprecated. These endpoints have little value for public use.
- For more information, see *Changes to the APIs*.

## Key Bug Fixes

| Ticket   | Description   |
|----------|---|
| TD-65753 | Some platform services do not bind to localhost only.   |
| TD-65502 | Datasets from parameters are improperly being permitted to be referenced in recipes and returns an error during job execution.  |
| TD-63953 | HYPER format publishing jobs remain in queued state on Azure.   |
| TD-63564 | Schedules created by a flow collaborator with editor access stop working if the collaborator is removed from the flow.<br><br>Collaborators with viewer access cannot create schedules. |

## New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-63974 | In imported datasets sourced from CSV files, double quotes that are escaped with a backslash ( <code>\ "backslash-escaped value\"</code> ) can cause remainder of row to be compressed into a single cell.  |
| TD-63517 | Unpivoting a String column preserves null values in Spark but converts them to empty strings in Photon. Running jobs on the different running environments generates different results.<br><br><div> <p><b>Workaround:</b> After the unpivot step, you can add an Edit with formula step. Set the columns to all of the columns in the unpivot and add the following formula, which converts all missing values to null values:</p> <pre>if(ismissing(\$col),NULL(),\$col)</pre> </div> |

## Release 8.6

August 2, 2021

### What's New

#### Collaboration:

- Flow editors and plan collaborators can be permitted to schedule jobs.  
See *Workspace Settings Page*.

#### Databricks:

- Support for Databricks 8.3.
  - For more information, see *Configure for Azure Databricks*.
  - For more information, see *Configure for AWS Databricks*.
- Support for per-user access to ADLS Gen2 for running Databricks jobs. For more information, see *ADLS Gen2 Access*.

#### Performance:

- Conversion jobs are now processed asynchronously.
- Better management of file locking and concurrency during job execution.

This feature must be enabled in your environment and requires installation of Redis server. See *Configure for Redis*.

#### Job execution:

- Support for execution of pre- and post-SQL scripts on Hive datastores.
  - This feature may need to be enabled in your environment. For more information, see *Workspace Settings Page*.
  - For more information, see *Create Output SQL Scripts*.

#### Better Handling of JSON files

The Trifacta application now supports the regularly formatted JSON files during import. You can now import flat JSON records contained in a single array object. With this, each array is treated as a single line and imported as a new row. For more information, see *Working with JSON v2*.

#### Changes in System Behavior

##### Compression scheme for imported files can now inferred

Prior to this release, when a file was imported, the Trifacta application inferred any compression applied to the file based on the filename extension at the end of the file. For example, if the filename ended with `.gz`, then the file was passed through the internal code for decompressing GZIP files.

Beginning in this release, the Trifacta application can be configured to detect the applied compression based on reading in the first few bytes of the file. Based on the data signature, the application passes the file for decompression to the appropriate code.

For more information on enabling, see *Miscellaneous Configuration*.

#### Deprecated

None.

#### Key Bug Fixes

None.

#### New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-63564 | <p>Schedules created by a flow collaborator with editor access stop working if the collaborator is removed from the flow.</p> <div> <p><b>Tip:</b> Flow owners can delete the schedule and create a new one. When this issue is fixed, the original schedule will continue to be executed under the flow owner's account.</p> </div> <p>Collaborators with viewer access cannot create schedules.</p> |

## Release 8.5

June 28, 2021

## What's New

### Parameterization:

- Create environment parameters to ensure that all users of the project or workspace use consistent references.

**NOTE:** You must be a workspace administrator or project owner to create environment parameters.

**Tip:** Environment parameters can be exported from one project or workspace and imported into another, so that these references are consistent across the enterprise.

- For more information, see *Environment Parameters Page*.
- For more information on parameters in general, see *Overview of Parameterization*.
- Parameterize names of your storage buckets using environment parameters.
  - See *Create Dataset with Parameters*.
  - See *Create Outputs*.
  - For more information on parameters, see *Overview of Parameterization*.

### Flow View:

- Click a node to see its lineage within the flow. See *Flow View Page*.

### Job execution:

- Define SQL scripts to execute before data ingestion or after publication for file-based or table-based jobs.
  - This feature may need to be enabled in your environment.

For more information, see *Workspace Settings Page*.

- For more information, see *Create Output SQL Scripts*.

### Connectivity:

- Connect to your relational database systems hosted on Amazon RDS. In the Connections page, click the **A mazon RDS** card for your connection type.

For more information, see *Create Connection Window*.

### Developer:

Download and install the Python SDK, which enables you to leverage the visual tools of the Trifacta application to transform data in your existing Python pipelines.

**NOTE:** This is an Alpha release. Do not deploy the Python SDK in production environments.

For more information, see *Python SDK*.

### **Job execution:**

You can choose to ignore the recipe errors before job execution and then review any errors in the recipe through the Job Details page.

- For more information, see *Run Job Page*.
- For more information, see *Job Details Page*.

### **Language:**

- NUMVALUE function can be used to convert a String value formatted as a number into an Integer or Decimal value.
- NUMFORMAT function now supports configurable grouping and decimal separators for localizing numeric values.
- For more information, see *Changes to the Language*.

### **Changes in System Behavior**

None.

### **Key Bug Fixes**

None.

### **New Known Issues**

None.

### **Release 8.4**

May 24, 2021

### **What's New**

### **Collaboration:**

- You can receive email notifications whenever a plan or a flow is shared with you by the owner.
  - This feature may need to be enabled in your environment.  
For more information, see *Workspace Settings Page*.
  - For more information, see *Email Notifications Page*.

### **Language:**

- New content on parsing functions. See *Changes to the Language*.

### **Changes in System Behavior**

None.

## Key Bug Fixes

| Ticket   | Description   |
|----------|---|
| TD-60881 | Incorrect file path and missing file extension in the application for parameterized outputs     |
| TD-60382 | Date format M/d/yyyy is handled differently by PARSEDATE function on Trifacta Photon and Spark. |

## New Known Issues

None.

## Release 8.3

April 26, 2021

## What's New

### Spark:

- Support for Spark 3.0.1 for AWS and Azure Databricks.

**NOTE:** Spark 3.0.1 is supported on Databricks for version 7.3 only.

**NOTE:** Execution of jobs sourced from Databricks Tables is now supported for Spark 3.0.1.

- See *Configure for AWS Databricks*.
- See *Configure for Azure Databricks*.
- Support for Spark 3.0.1 on EMR. See *Configure for EMR*.
- For more information, see *Configure for Spark*.

### EMR:

- Support for EMR 6.2 using Spark 3.0.1.

**NOTE:** EMR 6.1 and 6.0 are not supported in this release. For more information, see *Supported Deployment Scenarios for AWS*.

### Connectivity:

You can publish results to external S3 buckets through an Access key and Secret key. For more information, see *External S3 Connections*.

### Job execution:

Introducing new filter pushdowns to optimize the performance of your flows during job execution. For more information, see *Flow Optimization Settings Dialog*.



**Job results:**

You can now preview job results and download them from the Overview tab of the Job details page. For more information, see *Job Details Page* .

**Tip:** You can also preview job results in Flow View. See *View for Outputs*.

**Changes in System Behavior**

**Improved method of JSON import**

Beginning in this release, the Trifacta application now uses the conversion service to ingest JSON files during import. This improved method of ingestion can save significant time wrangling JSON into records.

**NOTE:** The new method of JSON import is enabled by default but can be disabled as needed.

For more information, see *Working with JSON v2*.

Flows that use imported datasets created using the old method continue to work without modification.

**NOTE:** It is likely that support for the v1 version of JSON import is deprecated in a future release. You should switch to using the new version as soon as possible. For more information on migrating your flows and datasets to use the new version, see *Working with JSON v1*.

Future work on support for JSON is targeted for the v2 version only. For more information on using the old version and migrating to the new version, see *Working with JSON v1* .

**Key Bug Fixes**

| Ticket   | Description  |
|----------|--|
| TD-60701 | Most non-ASCII characters incorrectly represented in visual profile downloaded in PDF format |
| TD-60660 | Azure SSO redirects to the Home page instead of the target page after login.                 |
| TD-59854 | Datetime column from Parquet file incorrectly inferred to the wrong data type on import.     |

**New Known Issues**

None.

# Release Notes 8.2

## Contents:

- *Release 8.2.2*
    - *What's New*
    - *Changes in System Behavior*
    - *Deprecated*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 8.2.1*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 8.2*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Fixes*
    - *New Known Issues*
  - *Release 8.1*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 8.0*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 7.10*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 7.9*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 7.8*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 7.7*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
- 

## Release 8.2.2

March 25, 2022

## What's New

### Databricks:

- Support for Databricks 7.x and 8.x.

**NOTE:** Databricks 7.3 and Databricks 8.3 are recommended.

- Support for Databricks cluster creation via cluster policies.
- Store a user-defined set of secret information such as credentials in Databricks Secrets.
  - For more information, see *Configure for Azure Databricks*.
  - For more information, see *Configure for AWS Databricks*.

### Changes in System Behavior

#### Publishing:

Improvements to publishing of Trifacta Date values to Snowflake. For more information, see *Improvements to the Type System*.

#### Nginx:

- Upgraded to Nginx 1.20.1.

### Deprecated

None.

### Key Bug Fixes

| Ticket   | Description  |
|----------|--|
| TD-69201 | Vulnerability scan detected compromised versions of log4j on the Hadoop dependency jars  |
| TD-69052 | Job fails using Spark when using parameterized files as input  |
| TD-69004 | Patch httpd to version 2.4.52  |
| TD-68085 | Trifacta unavailable due to update lock on plantasksnapshotruns  |
| TD-67953 | Remove log4j dependencies from Java projects   |
| TD-67747 | CVE-2021-44832: Apache Log4j2 vulnerable to RCE via JDBC Appender when attacker controls configuration                                 |
| TD-67677 | EMR spark job fails with error "org.apache.spark.sql.AnalysisException: Cannot resolve column name" if flow optimizations are enabled. |
| TD-67640 | Intermittent failure to publish to Tableau in Fileconverter.   |
| TD-67572 | EMR spark job fails with error "org.apache.spark.sql.AnalysisException: Cannot resolve column name"                                    |
| TD-67558 | CVE-2021-45105: Log4j vulnerability (denial of service)  |
| TD-67531 | Glue jobs not working after upgrade to Release 8.2   |
|          | CVE-2021-45046: Log4j vulnerability  |

|          |  |
|----------|--|
| TD-67455 |  |
| TD-67410 | CVE-2021-23017: Nginx v.1.20.0 security vulnerability  |
| TD-67388 | Nest function failing  |
| TD-67372 | Patch/update Log4J (RCE 0-day exploit found in log4j)  |
| TD-67329 | Publish failing with " <i>java.io.IOException: No FileSystem for scheme: sftp</i> "  |
| TD-66779 | Output home directory is not picked correctly for job runs in wasb/adls-gen2   |
| TD-66160 | SSLHandshakeException when accessing Databricks table  |
| TD-66025 | Glue connection not working on after upgrade to Release 8.2  |
| TD-65696 | In Azure environment, changing the user output/upload directory only persists the path and not the container name/account storage. |
| TD-65331 | Writing to ADLS failing in SSL Handshake to TLSv1.1  |
| TD-65286 | Trifacta jobs fail at Transform stage with Optimizer Service exception   |
| TD-65058 | Unable to upgrade due to migration failure   |
| TD-64627 | Trifacta failing due to concurrent DB transaction  |
| TD-64528 | Upgrade to Release 8.2 failed to load dictionaries   |
| TD-64281 | /change-password page fails to load.   |
| TD-64171 | Cannot import parameterized datasets that include files with zero and non-zero byte sizes together.                                |
| TD-63981 | Start/stop scripts should not modify any config/database settings during startup.  |
| TD-63867 | Jobs are not triggering for a parameterized datasets with zero-byte file sizes.  |
| TD-63493 | Unable to cancel a plan run  |
| TD-60881 | Incorrect path shown when using parameterized output path  |
| TD-59706 | No vertical scroll when there are too many connections on Import page  |
| TD-58576 | Cannot read property 'expandScriptLines' of undefined when flow node's activeSampleId is pointing to failed (null) sample.         |

## New Known Issues

None.

## Release 8.2.1

August 13, 2021

## What's New

### EMR:

- Updated supported versions of EMR to address log4j2 issues. Please upgrade to EMR 5.30.2.
  - For more information, see <https://docs.trifacta.com/display/PUB/Trifacta+Alert+TD-67372+-+0-Day+Exploit+in+log4j2+for+Self+Managed>

### EMR:

- Updated supported versions of EMR to address log4j2 issues. Please upgrade to EMR 5.30.2.
  - For more information, see <https://docs.trifacta.com/display/PUB/Trifacta+Alert+TD-67372+-+0-Day+Exploit+in+log4j2+for+Self+Managed>

### Databricks:

- Support for Databricks 8.3.
  - For more information, see *Configure for Azure Databricks*.
  - For more information, see *Configure for AWS Databricks*.
- Support for per-user access to ADLS Gen2 for running Databricks jobs. For more information, see *ADLS Gen2 Access*.

### Trifacta node:

- NodeJS upgraded to 14.16.0.

## Changes in System Behavior

None.

## Key Bug Fixes

| Ticket   | Description  |
|----------|--|
| TD-62689 | <p>Nginx returns Bad Request Status: 400 error, due to duplicate entries in <code>/etc/nginx/conf.d/trifacta.conf</code> for:</p> <pre>proxy_set_header Host \$host;</pre> <p><b>Tip:</b> Workaround is to delete the second entry in the file manually.</p> |

## New Known Issues

None.

## Release 8.2

June 11, 2021

## What's New

### Preferences:

- Re-organized user account, preferences, and storage settings to streamline the setup process. See *Preferences Page*.

### API:

### Connectivity:

- Support for custom SQL queries on Databricks Tables. See *Databricks Tables Connections*.

### Databricks:

Support for Databricks 7.3, using Spark 3.0.1.

**NOTE:** Databricks 5.5 LTS is scheduled for end of life in July 2021. An upgrade to Databricks 7.3 is recommended.

**NOTE:** In this release, Spark 3.0.1 is supported for use with Databricks 7.3 only.

- For more information, see *Configure for Azure Databricks*.
- For more information, see *Configure for AWS Databricks*.

### Plan metadata references:

Use metadata values from other tasks and from the plan itself in your HTTP task definitions.

- For more information, see *Create HTTP Task*.
- For more information, see *Plan Metadata References*.

### Improved accessibility of job results:

The Jobs tabs have been enhanced to display the list of latest and the previous jobs that have been executed for the selected output.

For more information, see *View for Outputs*.

### Sample Jobs Page:

You can monitor the status of all sample jobs that you have generated. Project administrators can access all sample jobs in the workspace. For more information, see *Sample Jobs Page*.

### Install:

Support for Nginx 1.20.0 on the Trifacta node. See *System Requirements*.

### Changes in System Behavior

### Java service classpath changes:

**NOTE:** This required update applies only to customers who have modified their Java service classpaths to include `/etc/hadoop/conf`.

In deployments on a Hadoop edge node, the classpath values for Java-based services may have been modified to include the following:

```
/etc/hadoop/conf
```

As of this release, symlinks must be created to locations within the Trifacta install directory to replace the above path modifications.

**NOTE:** Before you before the following update, please create a backup of `/etc/hadoop/conf` first.

In the following example, all files in the `etc/hadoop/conf` directory are updated with symlinks to the proper directory in the `conf` directory of files.

```
for file in `ls /etc/hadoop/conf`; do ln -sf /etc/hadoop/conf/$file /opt/trifacta/conf/hadoop-site/$file; done
```

### Running Environment:

Cloudera 5.x, including Cloudera 5.16, is no longer supported. Please upgrade to a supported version of Cloudera 6.x.

- For more information on supported version, see *Product Support Matrix*.
- For more information, see *End of Life and Deprecated Features*.

### Catalog integrations end of life:

The following catalog integrations are no longer available in the platform:

- Alation
- Waterline
- Cloudera Navigator

For more information, see *End of Life and Deprecated Features*.

### API:

The following API endpoints are scheduled for deprecation in a future release:

**NOTE:** Please avoid using the following endpoints.

```
/v4/connections/vendors  
/v4/connections/credentialTypes  
/v4/connections/:id/publish/info  
/v4/connections/:id/import/info
```

These endpoints have little value for public use.

## Key Fixes

| Ticket   | Description  |
|----------|--|
| TD-59854 | Datetime column from Parquet file incorrectly inferred to the wrong data type on import.                   |
| TD-59658 | IAM roles passed through SAML does not update after Hotfix upgrade   |
| TD-59633 | Enabled session tag feature but running into "The security token included in the request is invalid" error |
| TD-59331 | When include quotes option is disabled on an output, Databricks still places quotes around empty values.   |
| TD-59128 | BOM characters at the beginning of a file causing multiple headers to appear in Transformer Page.          |
| TD-58932 | Cannot read file paths with colons from EMR Spark jobs   |
| TD-58694 | Very large number of files generated during Spark job execution  |
| TD-58523 | Cannot import dataset with filename in Korean alphabet from HDFS.  |

## New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-60701 | Most non-ASCII characters incorrectly represented in visual profile downloaded in PDF format. |

## Release 8.1

February 26, 2021

### What's New

**In-app messaging:** Be sure to check out the new in-app messaging feature, which allows us to share new features and relevant content to Trifacta users in your workspace. The user messaging feature can be disabled by workspace administrators if necessary. See *Workspace Settings Page*.

### Install:

- Support for PostgreSQL 12.X for Trifacta databases on all supported operating systems.

**NOTE:** Beginning in this release, the latest stable release of PostgreSQL 12 can be installed with the Trifacta platform. Earlier versions of PostgreSQL 12.X can be installed manually.

**NOTE:** Support for PostgreSQL 9.6 is deprecated for customer-managed Hadoop-based deployments and AWS deployments. PostgreSQL 9.6 is supported only for Azure deployments. When Azure supports PostgreSQL 12 or later, support for PostgreSQL 9.6 will be deprecated in the subsequent release of Trifacta.

- For more information, see *Install Databases for PostgreSQL*.
- See *Product Support Matrix*.

### Security:



- Support for SSL secure access to the Trifacta databases. For more information, see

*Enable SSL for Databases.*

#### Databases:

- New databases:
  - The Secure Token Service database is used for managing the tokens used by the secure token service.
  - The Connector Configuration Service database stores the connection configuration information for a workspace's available connectors (connection types).
  - These databases are installed and managed in conjunction with the other Trifacta databases. See *Install Databases.*

#### Connectivity:

- For AWS-based installations, you can create multiple read-only S3 connections through the Trifacta application. These connections use key and secret pair combinations to access specific S3 buckets. For more information, see *External S3 Connections*.
- You can enable logging of events from the CData driver underlying your supported relational connections. For more information, see *Configure Connectivity*.

#### Authorization:

- Support for fine-grained access controls to S3 and EMR through AWS session tags in your IAM roles. For more information, see *Configure AWS Per-User Auth for Temporary Credentials*.

#### Sharing:

- Define permissions on individual objects when they are shared.

**NOTE:** Fine-grained sharing permissions apply to flows and connections only.

For more information, see *Changes to User Management*.

#### API:

- Apply job-level overrides to AWS Databricks or Azure Databricks job executions via API. See *API Workflow - Run Job*.
- Customize connection types (connectors) to ensure consistency across all connections of the same type and to meet your enterprise requirements. For more information, see *Changes to the APIs*.

#### Running environment:

- Support for Databricks as a running environment for Trifacta hosted in AWS.
  - For more information, see *Configure for AWS Databricks*.

- Integration with AWS Secrets Manager is required. For more information, see *Configure for AWS Secrets Manager*.
- Support for connections to Databricks Tables from AWS Databricks. See *Databricks Tables Connections*.
- Support for job throttling for user clusters.
  - See *Configure for AWS Databricks*.
  - See *Configure for Azure Databricks*.
- Support for custom driver-specific instance pools on Databricks.
  - See *Configure for Azure Databricks*.
  - See *Configure for AWS Databricks*.

### Publishing:

- Support for publishing to Datetime values in Parquet outputs. For more information, see *Improvements to the Type System*.

### Macro updates:

You can replace an existing macro definition with a macro that you have exported to your local desktop.

**NOTE:** Before you replace the existing macro, you must export a macro to your local desktop. For more information, see *Export Macro*.

For more information, see *Macros Page*.

### Sample Jobs Page:

You can monitor the status of all sample jobs that you have generated. Project administrators can access all sample jobs in the workspace. For more information, see *Sample Jobs Page*.

### Specify column headers during import

You can specify the column headers for your dataset during import. For more information, see *Import Data Page*.

### Services:

- The Secure Token Service is used for managing tokens for third-party systems, such as Azure Key Vault and OAuth 2.0 authentication. See *Configure Secure Token Service*.
- The Connector Configuration Service manages the storage and retrieval of connection type information for the workspace. See *Configure Connector Configuration Service*.

### Changes in System Behavior

**NOTE:** CDH 6.1 is no longer supported. Please upgrade to the latest supported version. For more information, see *Product Support Matrix*.

**NOTE:** HDP 2.6 is no longer supported. Please upgrade to the latest supported version. For more information, see *Product Support Matrix*.

### Support for custom data types based on dictionary files to be deprecated:

**NOTE:** The ability to upload dictionary files and use their contents to define custom data types is scheduled for deprecation in a future release. This feature is limited and inflexible. Until an improved feature can be released, please consider using workarounds. For more information, see

*Validate Your Data.*

You can create custom data types using regular expressions.

### **Strong consistency management now provided by AWS S3:**

Prior to this release, S3 sometimes did not accurately report the files that had been written to it, which resulted in consistency issues between the files that were written to disk and the files that were reported back to the Trifacta application.

As of this release, AWS has improved S3 with strong consistency checking, which removes the need for the product to maintain a manifest file containing the list of files that have been written to S3 during job execution.

**NOTE:** As of this release, the S3 job manifest file is no longer maintained. All configuration related to this feature has been removed from the product. No additional configuration is needed.

For more information, see <https://aws.amazon.com/s3/consistency/>.

For more information on integration with S3, see *S3 Access*.

### **Installation of database client is now required:**

Before you install or upgrade the database or perform any required database cross-migrations, you must install the appropriate database client first.

**NOTE:** Use of the database client provided with each supported database distribution is now a required part of any installation or upgrade of the Trifacta platform.

For more information:

- See *Install Databases for PostgreSQL*.
- See *Install Databases for MySQL*.

### **Job logs collected asynchronously for Databricks jobs:**

In prior releases, the Trifacta application reported that a job failed only after the job logs had been collected from the Databricks cluster. This log collection process could take a while to complete, and the job was reported as in progress when it had already failed.

Beginning in this release, collection of Databricks job logs for failed jobs happens asynchronously. Jobs are now reported in the Trifacta application as soon as they are known to have failed. Log collection happens in the background afterward.

- See *Configure for AWS Databricks*.
- See *Configure for Azure Databricks*.

### **Catalog integrations now deprecated:**

Integrations between Trifacta and Alation and Waterline services are now deprecated. For more information, see *End of Life and Deprecated Features*.

## Key Bug Fixes

| Ticket   | Description   |
|----------|---|
| TD-56170 | The Test Connection button for some relational connection types does not perform a test authentication of user credentials.   |
| TD-54440 | <p>Header sizes at intermediate nodes for JDBC queries cannot be larger than 16K.</p> <p>Previously, the column names for JDBC data sources were passed as part of a header in a GET request. For very wide datasets, these GET requests often exceeded 16K in size, which represented a security risk.</p> <p>The solution is to turn these GET requests into ingestion jobs.</p> <div><b>NOTE:</b> To mitigate this issue, JDBC ingestion and JDBC long loading must be enabled in your environment. For more information, see <i>Configure JDBC Ingestion</i>.</div> |

## New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-58818 | <p>Cannot run jobs on some builds HDP 2.6.5 and later. There is a known incompatibility between HDP 2.6.5.307-2 and later and the Hadoop bundle JARs that are shipped with the Trifacta installer .</p> <div><b>Solution:</b> The solution is to use an earlier compatible version.</div> |
| TD-58523 | <p>Cannot import dataset with filename in Korean alphabet from HDFS.</p> <div><b>Workaround:</b> You can upload files with Korean characters from your desktop. You can also add a 1 to the end of the file on HDFS, and it can then be imported.</div>                                   |
| TD-55299 | Imported datasets with encodings other than UTF-8 and line delimiters other than \n may generate empty outputs on Spark or D atafLOW running environments.  |
| TD-51516 | Input data containing BOM (byte order mark) characters may cause Spark or Dataflow running environments to read data improperly and/or generate invalid results.  |

## Release 8.0

January 26, 2021

### What's New

#### APIs:

- Individual workspace users can be permitted to create and use their own access tokens for use with the REST APIs. For more information, see *Workspace Settings Page*.

#### Connectivity:

- Support for connections to SharePoint Lists. See *SharePoint Connections* .
- Support for using OAuth2 authentication for Salesforce connections.

**NOTE:** Use of OAuth2 authentication requires additional configuration. For more information, see *OAuth 2.0 for Salesforce*.

See *Salesforce Connections*.

- Support for re-authenticating through connections that were first authenticated using OAuth2.

#### Import:

- Improved method for conversion and ingestion of XLS/XSLX files. For more information, see *Import Excel Data*.

#### Recipe development:

- The Flag for Review feature enables you to set review checkpoints in your recipes. You can flag recipe steps for review by other collaborators for review and approval. For more information, see *Flag for Review*.

#### Update Macros:

- Replace / overwrite an existing macro's steps and inputs with a newly created macro.
- Map new macro parameters to the existing parameters before replacing.
- Edit macro input names and default values as needed.
  - For more information, see *Create or Replace Macro*.
  - For more information, see *Overview of Macros*.

#### Job execution:

- You can enable the Trifacta application to apply SQL filter pushdowns to your relational datasources to remove unused rows before their data is imported for a job execution. This optimization can significantly improve performance as less data is transferred during the job run. For more information, see *Flow Optimization Settings Dialog*.
- Optimizations that were applied during the job run now appear in the Job Details Page. See *Job Details Page*.

#### Changes in System Behavior

None.

#### Key Bug Fixes

| Ticket   | Description   |
|----------|---|
| TD-57354 | Cannot import data from Azure Databricks. This issue is caused by an incompatibility between TLS v1.3 and Java 8, to which it was backported. |
| TD-57180 | AWS jobs run on Photon to publish to HYPER format fail during file conversion or writing.   |

#### New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-56170 | <p>The Test Connection button for some relational connection types does not perform a test authentication of user credentials.</p> <div><p><b>Workaround:</b> Append the following to your Connect String Options:</p><pre>;ConnectOnOpen=true</pre><p>This option forces the connection to validate user credentials as part of the connection. There may be a performance penalty when this option is used.</p></div> |

## Release 7.10

December 21, 2020

### What's New

**Tip:** Check out the new in-app tours, which walk you through the steps of wrangling your datasets into clean, actionable data.

#### Import:

- The maximum permitted size of a file uploaded through Trifacta application has been increased from 100 MB to 1 GB.

#### Plan View:

- **Import and Export Plans:** You can import and export plans from one environment, workspace, or projects to others.

For more information, see *Export Plan*.

For more information, see *Import Plan*.

- **Share Plans:** Share plans with one or more users to work together on the same plan. For more information, see *Share a Plan*.
- **Email notifications:** Send email notifications to plan owners and collaborators based on the status of execution of plans. For more information, see *Manage Plan Notifications Dialog*.

#### Authentication:

- Support for authentication to third-party datastores using OAuth2. For more information, see *Enable OAuth 2.0 Authentication*.
- Support for using OAuth2 authentication for Snowflake connections. For more information, see *OAuth 2.0 for Snowflake*.

#### Connectivity:

- Improved Salesforce connection type.

For more information, see *Salesforce Connections*.

#### Language:

- New function for calculating end-of-month values. See *Changes to the Language*.

#### API:

- **Experimental feature:** Export Python Pandas code to generate the transformation steps required to produce a defined output object.
- **NOTE:** This feature can be changed or removed from the platform at any time without notice. Do not deploy it in a production environment.

- For more information, see *API Workflow - Wrangle Output to Python*.

## Changes in System Behavior

### Rebuild custom UDF JARs for Databricks clusters

Previously, UDF files were checked for consistency based upon the creation time of the JAR file. However, if the JAR file was passed between Databricks nodes in a high availability environment or between services in the platform, this timestamp could change, which could cause job failures due to checks on the created-at timestamps.

Beginning in this release, the platform now inserts a build-at timestamp into the custom UDF manifest file when the JAR is built. This value is fixed, regardless of the location of the copy of the JAR file.

**NOTE:** Custom UDF JARs that were created using earlier releases of the platform and deployed to a Databricks cluster must be rebuilt and redeployed as of this release. For more information on troubleshooting the error conditions, see *Java UDFs*.

### Custom credential provider JAR no longer required for EMR access

In prior releases of Trifacta, integration with EMR required the deployment of a custom credential provider JAR file provided by the customer as part of the initial bootstrap of the EMR cluster. As of this release, this JAR file is no longer required. Instead, it is provided by the Trifacta platform directly.

**NOTE:** If your deployment of the Trifacta platform integrates with AWS Glue, you must still provide and deploy a custom credentials JAR file. For more information, see *AWS Glue Access*.

For more information on integrating with EMR, see *Configure for EMR*.

### Upgrade nodeJS

On the Trifacta node, the version of nodeJS has been upgraded to nodeJS 14.15.4 LTS. For more information, see *System Requirements*.

### Data type and row split inference utilize more data

When a dataset is loaded, the Trifacta application now reads in more data before the type inference system and row splitting transformations analyze the data to break it into rows and columns. This larger data size should result in better data inference in the system.

**NOTE:** Types and row splits on pre-existing datasets may be affected by this change.

For more information, see *Improvements to the Type System*.

## Key Bug Fixes

| Ticket   | Description   |
|----------|---|
| TD-54742 | Access to S3 is disabled after upgrade.   |
| TD-53527 | When importing a dataset via API that is sourced from a BZIP file stored on S3, the columns may not be properly split when the platform is permitted to detect the structure. |

## New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-57180 | <p>AWS jobs run on Photon to publish to HYPER format fail during file conversion or writing.</p> <div><b>Workaround:</b> Run the job on the Spark running environment instead.</div>  |
| TD-56830 | <p>Receive <code>malformed_query</code>: enter a filter criterion when importing table from Salesforce.</p> <div><b>NOTE:</b> Some Salesforce tables require mandatory filters when they are queried. Mandatory filters are not currently supported for Salesforce connections.</div> |

## Release 7.9

November 16, 2020

### What's New

#### Plan View:

- **Execute Plan using status rules:** Starting in Release 7.9, you can execute tasks based on the previous task execution result. For more information, see *Create a Plan*.
- **Execute Parallel Plan tasks:** In previous releases, plans were limited to a sequential order of task execution. Beginning in Release 7.9, you can create branches in the graph into separate parallel nodes, enabling the corresponding tasks to run in parallel. This feature enables you to have a greater level of control of your plans' workflows. For more information, see *Create a Plan*.
- **Zoom options:** Zoom control options and keyboard shortcuts have been introduced in the plan canvas. For more information, see *Plan View Page*.
- **Filter Plan Runs:** Filter your plan runs based on dates or plan types. For more information, see *Plan Runs Page*.

#### Transform Builder:

- An All option has been added for selecting columns in the Transform Builder. For more information, see *Changes to the Language* page.

### Changes in System Behavior

#### Manage Users section has been deprecated:

In previous releases, user management functions were available through the Manage Users section of the Admin Settings page. These functions have been migrated to the Workspace Settings page, where all of the previous functions are now available. The Manage Users section has been deprecated.

- For more information, see *Changes to User Management*.
- For more information, see *Users Page*.

#### Better license management:

In prior releases, the Trifacta application locked out all users if the number of active users exceeded the number permitted by the license. This situation could occur if users were being added via API, for example.

Beginning in this release, the Trifacta application does not block access when the number of licensed users is exceeded.



**NOTE:** If you see the notification banner about license key violations, please adjust your users until the banner is removed. If you need to adjust the number of users associated with your license key, please contact *Alteryx Support*.

For more information, see *License Key*.

### Trifacta Photon jobs now use ingestion for relational sources:

When a job is run on Trifacta Photon, any relational data sources are ingested into the backend datastore as a preliminary step during sampling or transformation execution. This change aligns Trifacta Photon job execution with future improvements to the overall job execution framework. No additional configuration is required.

**Tip:** Jobs that are executed on the Trifacta Server are executed in an embedded running environment, called **Trifacta Photon**. Quick Scan samples are automatically executed in Trifacta Photon.

For more information on ingestion, see *Configure JDBC Ingestion*. **Job results page changes:**

- The dependencies tab is renamed as dependency graph tab.
- The old flow view in the dependency graph tab is replaced with the new flow view. For more information, see *Job Details Page*.

### Key Bug Fixes

| Ticket   | Description   |
|----------|---|
| TD-55125 | Cannot copy flow. However, export and import of the flow enables copying. |
| TD-53475 | Missing associated artifact error when importing a flow.                  |

### New Known Issues

None.

## Release 7.8

October 19, 2020

### What's New

#### Plans:

- Create HTTP tasks for your plans, which can be configured to issue a request to an API endpoint over HTTP.
  - For more information, see *Plan View for HTTP Tasks*.
  - For more information on plans, see *Overview of Operationalization*.
- The viewport position and zoom level are now preserved when returning to a given flow.

#### Publishing:

- Improved performance when publishing to Tableau Server.
- Configure publishing chunk sizes as needed. For more information, see *Configure Data Service*.

#### Language:

- Rename columns now supports uppercase or lowercase characters or shorten column names to a specified character length from the left or right. For more information, see *Changes to the Language*.

#### Connectivity:

- IAM support for Redshift connections.

**NOTE:** To enable use of an existing IAM role for Redshift, additional permissions must be added. For more information, see *Required AWS Account Permissions*.

For more information, see *Amazon Redshift Connections*.

## Changes in System Behavior

#### JDBC connection pooling disabled:

**NOTE:** The ability to create connection pools for JDBC-based connections has been disabled. Although it can be re-enabled if necessary, it is likely to be removed in a future release. For more information, see *Changes to Configuration*.

#### TDE format has been deprecated:

Tableau Server has deprecated support for the TDE file format. As of this release, all outputs and publications to Tableau Server must be generated using HYPER, the replacement format for TDE.

- Any flow that uses TDE format is automatically switched to use HYPER format during the upgrade process.
- Any flow that is imported into the upgraded environment is automatically switched to using the HYPER format.

For more information, see *Tableau Hyper Data Type Conversions*.

#### Enhanced Flow and Flow View menu options :

The context menu options for Flow View and Flow have been renamed and reorganized for a better user experience.

- For more information, see *Flows Page*.
- For more information, see *Flow View Page*.

## Key Bug Fixes

None.

## New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-54030 | <p>When creating custom datasets from Snowflake, columns containing time zone data are rendered as null values in visual profiles, and publishing back to Snowflake fails.</p> <p><b>Workaround:</b> In your SELECT statement applied to a Snowflake database, references to time zone-based data must be wrapped in a function to convert it to UTC time zone. For more information, see <i>Create Dataset with SQL</i>.</p> |

## Release 7.7

September 21, 2020

### What's New

#### Flow View:

- Automatically organize the nodes of your flow with a single click. See *Flow View Page*.

### Changes in System Behavior

#### Deprecated Parameter History Panel Feature

As a part of collaborative suggestions enhancement, the support for Parameter History panel is deprecated from the software. For more information on collaborative suggestions feature, see *Overview of Predictive Transformation*.

#### Classic Flow View no longer available

In Release 7.6, an improved version of Flow View was released. At the time of release, users could switch back to using the classic version.

Beginning in this release, the classic version of Flow View is no longer available.

**Tip:** The objects in your flows that were created in classic Flow View may be misaligned in the new version of Flow View. You can use auto-arrange to re-align your flow objects.

For more information, see *Flow View Page*.

### Key Bug Fixes

| Ticket   | Description  |
|----------|--|
| TD-53318 | Cannot publish results to relational targets when flow name or output filename or table name contains a hyphen (e.g. my - filename.csv). |

### New Known Issues

None.

# Release Notes 7.6

## Contents:

- *Releases 7.6.2*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Fixes*
    - *New Known Issues*
  - *Release 7.6.1*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Fixes*
    - *New Known Issues*
  - *Release 7.6*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
- 

## Releases 7.6.2

May 12, 2021

### What's New

#### EMR:

- Updated supported versions of EMR to address log4j2 issues. Please upgrade to EMR 5.30.2.
  - For more information, see <https://docs.trifacta.com/display/PUB/Trifacta+Alert+TD-67372+-+0-Day+Exploit+in+log4j2+for+Self+Managed>

#### Security:

- Support for SSL secure access to the Trifacta databases. For more information, see *Enable SSL for Databases*.
  - Support for SSL access to Amazon RDS instances of the Trifacta databases. For more information, see *Install Databases on Amazon RDS*.

#### Authorization:

- Support for fine-grained access controls to S3 and EMR through AWS session tags in your IAM roles. For more information, see *Configure AWS Per-User Auth for Temporary Credentials*.

#### Trifacta node:

- NodeJS upgraded to 12.22.1.

### Changes in System Behavior

None.

## Key Fixes

| Ticket   | Description   |
|----------|---|
| TD-60881 | Incorrect file path and missing file extension in the application for parameterized outputs                         |
| TD-60378 | Join and Union user interfaces are taking too long to respond.  |
| TD-60187 | Snowflake publishing fails during validation when both Stage database and External Stage are in use.                |
| TD-59658 | IAM roles passed through SAML does not update after Hotfix upgrade  |
| TD-59633 | Enabled session tag feature but running into "The security token included in the request is invalid" error          |
| TD-59249 | After ad-hoc publishing, job cleanup process deletes user's output directory on HDFS.                               |
| TD-59229 | Uploaded CSV file fails with Parquet schema error.  |
| TD-58932 | Cannot read file paths with colons from EMR Spark jobs  |
| TD-58591 | Copied flow does not include output objects.  |
| TD-58433 | Some recipe steps missing in copied flow  |
| TD-58036 | Custom SQL query from Hive fails to run.  |
| TD-57792 | Trifacta Photon jobs take a long time to write output without chunked encoding.                                     |
| TD-57653 | Monitoring details not visible in dataset details and flow view   |
| TD-57528 | Slow ingest time for small XSLX files   |
| TD-57512 | Java-vfs-service is becoming unresponsive when fetching files for a parameterized dataset comprised of 10,000 files |
| TD-57264 | Transformation engine crashes when specifying Group By parameter for List function.                                 |
| TD-56739 | Memory leak in java-vfs-service   |
| TD-53375 | S3 browsing 400 error   |

## New Known Issues

| Ticket   | Description  |
|----------|--|
| TD-59854 | Datetime column from Parquet file incorrectly inferred to the wrong data type on import.<br><br><b>Workaround:</b> Use the column drop-down to change the data type to Datetime. |

## Release 7.6.1

December 7, 2020

### What's New

#### Install:

- Support for PostgreSQL 12.3 for Trifacta databases on all supported operating systems.

**NOTE:** This feature is in Beta release.

**NOTE:** Support for PostgreSQL 9.6 will be deprecated in a future release.

- For more information, see *Install Databases for PostgreSQL*.
- See *Product Support Matrix*.

#### **AWS:**

- Support for configurable endpoints in AWS GovCloud. See *Configure for AWS*.

#### **Changes in System Behavior**

##### **Installation of database client is now required:**

Beginning in this release, before you install or upgrade the database or perform any required database cross-migrations, you must install the appropriate database client first.

**NOTE:** Use of the database client provided with each supported database distribution is now a required part of any installation or upgrade of the Trifacta platform.

**NOTE:** The MySQL database client cannot be provided by Alteryx. It must be downloaded and installed separately. As a result, installation or upgrade of a Docker environment using MySQL requires additional support. For more information, please contact *Alteryx Customer Success Services*.

For more information:

- See *Install Databases for PostgreSQL*.
- See *Install Databases for MySQL*.

##### **Catalog support to be deprecated:**

**NOTE:** Integrations with Alation and Waterline catalogs are likely to get deprecated in a future release.

##### **Support for custom data types based on dictionary files to be deprecated:**

**NOTE:** The ability to upload dictionary files and use their contents to define custom data types is scheduled for deprecation in a future release. This feature is limited and inflexible. Until an improved feature can be released, please consider using workarounds. For more information, see *Validate Your Data*.

##### **Maintenance release updater script is deprecated:**

The maintenance release updater script has been deprecated. This script could be used for performing maintenance upgrades:

- Release X.Y.1 to Release X.Y.2
- Hot Fixes

#### **Key Fixes**

None.

## New Known Issues

| Ticket   | Description  |
|----------|--|
| TD-57354 | <p>Cannot import data from Azure Databricks. This issue is caused by an incompatibility between TLS v1.3 and Java 8, to which it was backported.</p> <p>This issue is known to impact Marketplace installs of Trifacta Self-Managed Enterprise Edition and can impact on-premises installs.</p> <div><b>Workaround:</b> The solution is to downgrade Java on the Trifacta node to openJDKv1.8.0_242 or earlier. Java 8 is required. After you have downgraded, restart the platform. For more information, see <i>System Requirements</i>.</div>   |
| TD-56632 | <p>Non-default admin users are not automatically granted full workspace admin privileges on upgrade. These users may be able to see Workspace Settings and Admin Settings but are not granted access to edit roles and users.</p> <div><b>Workaround:</b> Login as the default admin user. Select <b>User menu &gt; Admin Console &gt; Roles</b>. For the Workspace Admin role, select <b>Assign Role</b>. Assign the role to the non-default admin users.</div> <p>For more information:</p> <ul style="list-style-type: none"><li>• See <i>Roles Page</i>.</li><li>• See <i>Users Page</i>.</li><li>• See <i>Changes to User Management</i>.</li></ul> |
| TD-54742 | <p>Access to S3 is disabled after upgrade.</p> <div><b>Workaround:</b> This issue is caused by the migration of the S3 enablement setting into the Workspace Settings page. To address, set <code>Enable S3 Connectivity</code> to <code>true</code>. For more information, see <i>Workspace Settings Page</i>.</div>  |
| TD-53527 | <p>When importing a dataset via API that is sourced from a BZIP file stored on a backend datastore such as S3, WASB, or ADLS Gen1/Gen2, the columns may not be properly split when the platform is permitted to detect the structure.</p> <div><b>Workaround:</b> Import the dataset via UI. If you must still import via API, please change <code>webapp.loadLimitForSplitInference</code> to 900000. See <i>Admin Settings Page</i>.</div>   |

## Release 7.6

September 7, 2020

### What's New

#### New Flow View is now generally available:

- Drag and drop to reposition objects on the Flow View canvas, and zoom in and out to focus on areas of development.
- Perform joins and unions between objects on the Flow View canvas.
- Search for flow objects by name or by type.
- Annotate the canvas with notes.

**Tip:** The relative position of objects on the flow view canvas is preserved between screen updates. On refresh, the window on the canvas is repositioned based on the leftmost object on the canvas to focus on the flow to other objects from that one.

**NOTE:** Classic Flow View is no longer available.

See *Flow View Page*.

#### Install:

- Support for PostgreSQL 12.3 for Trifacta databases on CentOS/RHEL 7.

**NOTE:** This feature is in Beta release.

- For more information, see *Install Databases for PostgreSQL*.
- See *Product Support Matrix*.
- Support for Cloudera Data Platform.

**NOTE:** Installation requirements for Cloudera Data Platform are consistent with installation for CDH. The Trifacta platform must be installed on a pre-existing Cloudera Data Platform.

There are minor differences in configuration. For more information, see *Configure for Cloudera*.

- Support for high availability on AWS.
  - On-premises installations can also be deployed in a highly available environment.
  - For more information, see *Install for High Availability*.
- Support for high availability integration with EMR clusters. For more information, see *Configure for EMR*.

#### Job execution:

- Support for Spark 2.4.6. For more information, see *Configure for Spark*.
- Support for EMR 5.30.1.

**NOTE:** Avoid EMR 5.30.0. Instead, please use EMR 5.30.1.

See *Configure for EMR*.

#### Import:

- For long-loading relational datasets, you can monitor the ingest process through Flow View as you continue your work.

**NOTE:** This feature may require enablement in your deployment. For more information, see *Configure JDBC Ingestion*.

For more information, see *Flow View Page*.

- Improved performance when browsing databases for tables to import.

**Tip:** Performance improvements are due to limiting the volume of table metadata that is imported when paging through available tables. This metadata can be retrieved when you hover over a table in the database browser.

For more information, see *Database Browser*.

- Logical and physical optimizations when reading from relational sources during job execution, which includes column pruning push-down among other enhancements.

**NOTE:** This feature may need to be enabled in your workspace. See *Workspace Settings Page*.



This feature applies to the following relational connections in this release:

- *Oracle Database Connections*
- *Microsoft SQL Server Connections*
- *PostgreSQL Connections*
- *Amazon Redshift Connections*
- *Snowflake Connections*
- *Microsoft SQL Data Warehouse Connections*

#### Flow View:

- Configure advanced settings on your flow and its job executions. See *Flow Optimization Settings Dialog*.

#### Plan View:

- Apply overrides to recipe parameters for your plans. See *Plan View Page*.
- New Plan Runs page:
  - Monitor status of all of your plan runs and drill into details.
  - Download logs for plan runs and individual flow tasks in the run.
  - See *Plan Runs Page*.

#### Transformer Page:

- Collaborative suggestions allow users within a workspace to receive suggestions based on the transformations that have been recently created by themselves or by all members of the workspace. As more users generate transformations, the relevance of these suggestions to the data in the workspace continues to improve.
  - This feature must be enabled for individual users or across the workspace. See *Workspace Settings Page*.
  - When the feature is enabled, individual users can choose to not share their transformation data. See *User Profile Page*.
  - For more information, see *Overview of Predictive Transformation*.
- Create and edit flow parameters and their overrides while editing your recipe. See *Transformer Page*.
  - For more information on editing flow parameters, see *Manage Parameters Dialog*.

#### Job Execution:

- Support for job cancellation on EMR clusters. See *Jobs Page*.

**NOTE:** Additional configuration may be required. For more information, see *Configure for EMR*.

- Azure Databricks enhancements:
  - Support for creating clusters using instance pools across multiple Databricks workspaces using instance pooling and Databricks pool names.
  - Manage jobs on Azure Databricks to prevent reaching Databricks workspace limits.
  - See *Configure for Azure Databricks*.
- When profiling is enabled for a Spark job, the transformation and profiling steps are combined into a single task, which optimizes the execution of transform and profiling tasks for a Spark job. For more information, see *Configure for Spark*.

#### Authorization:

- Workspace administrators can now create and assign roles to govern access to types of objects in the workspace. For more information, see *Changes to User Management*.

#### Single Sign-On:

- Support for configurable Azure AD endpoint and authority, including Gov Cloud. For more information, see *Configure SSO for Azure AD*.

### Connectivity:

- Improved performance for Oracle and SQL Server connections. These performance improvements will be applied to other relational connections in future releases.
- Improved performance when reading from Hive with many partitions into the Transformer page.

### Language:

- New approximation functions for median, percentile, and quartile based on a very fast algorithm.
- New functions to encode and decode base64 strings.
- New weekday name function.
- New rolling window date functions.
- New Kth-largest functions.
- New conditional minimum, maximum, and mode date functions.
- See *Changes to the Language*.

### Documentation:

- Additional connect string options and troubleshooting information has been included for specific relational connections. For more information, see *Connection Types*.

## Changes in System Behavior

### End of Life for Wrangler Enterprise desktop application

The Wrangler Enterprise desktop application is no longer available for installation and is not supported for use with the product. Please use one of the supported browser versions instead. For more information, see *Browser Requirements*.

### Users section of Admin Settings is disabled

In previous releases, the Users section of the Admin Settings page was used to manage users.

- Beginning in this release, the above area has been replaced by the Workspace Users page, where almost all user management tasks can be performed.
- The Users section must still be used for assigning the Trifacta admin platform role and the SSO, Hadoop, or Kerberos principals through the Trifacta application. It can be enabled as needed.
- For more information, see *Changes to User Management*.

### CentOS/RHEL 7.1 and 7.2 deprecated

Please upgrade to a supported distribution of either operating system. For more information, see *System Requirements*.

### S3 access uses Java VFS service

Access to S3 is now managed through the Java-based virtual file system. For more information, see *Configure Java VFS Service*.

**NOTE:** No configuration changes are required for upgrading customers. For more information, see *S3 Access*.

## Schema information is retained

When schematized datasources are ingested, schema information is now retained for publication of job results.

**NOTE:** In prior releases, you may have set column data types manually because this schema information was lost during the ingest process. You may need to remove these manual steps from your recipe. For more information, see *Improvements to the Type System*.

## Enhanced PII masking

For social security numbers and credit card numbers, the methods by which these values are determined for purposes of masking sensitive Personally Identifiable Information (PII) has been expanded and improved. For more information, see *Improvements to the Type System*.

## Updated credential types for connections via API

- Redshift connections now require a different credential type.
- Snowflake connections now require a different credential type.
- See *Changes to the APIs*.

## New components

Optimizer service and database: During job execution on relational sources, the optimizer service assists in managing SQL queries efficiently so that smaller volumes of data are retrieved for the job. Queries are stored in the related database.

## Key Bug Fixes

| Ticket   | Description   |
|----------|---|
| TD-52221 | API: Unable to update awsConfig objects in per-user or per-workspace modes.   |
| TD-51229 | When an admin user shares a flow that the admin user owns, a <code>Failed to share flow with selected user</code> error message may be displayed, even though the flow was successfully shared. |
| TD-48915 | Inserting special characters in an output filename results in a validation error in the the application and job failures.   |
| TD-47696 | Platform appears to fail to restart properly through Admin Settings page due to longer restarts of individual services.   |
| TD-49559 | Cannot select and apply custom data types through column Type menu.   |
| TD-47473 | Uploaded files (CSV, XLS, PDF) that contain a space in the filename fail to be converted.   |
| TD-34840 | Platform fails to provide suggestions for transformations when selecting keys from an object with many of them.   |

## New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-51840 | Import of dataset from Alation catalog hangs. <div><b>NOTE:</b> The Alation catalog integration is not working in Release 7.6. For more information, please contact <i>Alteryx Support</i>.</div> |
| TD-      |   |

50942

If a flow is unshared with you, you cannot see or access the datasources for any jobs that you have already run on the flow. You can still access the job results.

# Release Notes 7.1

## Contents:

- *Release 7.1.2*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 7.1.1*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 7.1*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
- 

## Release 7.1.2

November 25, 2020

### What's New

This release provides fixes to key issues.

### Publishing:

- Improved performance when publishing to Tableau Server.
- Configure publishing chunk sizes as needed. For more information, see *Configure Data Service*.

### Changes in System Behavior

None.

### Key Bug Fixes

| Ticket   | Description  |
|----------|--|
| TD-55714 | Receiving a 502 Internal Server error when attempting to use to a tested SFTP server connection. <div><b>NOTE:</b> In this case, the issue relates to how the batch job runner service authenticates to the SFTP server. For more information on configuration options, see <i>SFTP Connections</i>.</div> |
| TD-55125 | Cannot copy flow. However, export and import of the flow enables copying.  |
| TD-53475 | Missing associated artifact error when importing a flow.   |
| TD-52737 | After it is created, SFTP connection displays a blank page when opening from Import Data page.   |

## New Known Issues

None.

## Release 7.1.1

August 21, 2020

## What's New

- Support for PostgreSQL 12.3 for Trifacta databases.

**NOTE:** For this release, PostgreSQL 12.3 is supported for supported versions of CentOS/RHEL 7 only. See *Product Support Matrix*.

**NOTE:** In a future release, support for PostgreSQL 9.6 will be deprecated. For more information, see *Upgrade Databases for PostgreSQL*.

## Azure Databricks:

- Support for configurable Azure AD endpoint and authority for SSO validation. For more information, see *Configure SSO for Azure AD*.

## Changes in System Behavior

### Cloudera support:

If you are upgrading your cluster to CDH 6.3.3, please set the following property to the value listed below:

```
"spark.version": "2.4.cdh6.3.3.plus",
```

Save your changes and restart the platform. For more information, see *Admin Settings Page*.

For more information, see *Configure for Spark*.

## Key Bug Fixes

| Ticket   | Description  |
|----------|--|
| TD-53062 | After upgrade, imported recipe has UDF steps converted to comments.                                |
| TD-52738 | On Azure Databricks, creating a stratified sample fails.   |
| TD-52686 | Cannot run Azure Databricks jobs on ADLS-Gen1 cluster in user mode.                                |
| TD-52614 | UnknownHostException error when generating Azure Databricks access token from Secure Token Service |
| TD-51903 | Cannot import some Parquet files into the platform.  |
| TD-51681 | Import data page is taking too long to load.   |
| TD-51537 | Closing the connections search bar removes search bar and loses sort order.                        |

|          |   |
|----------|---|
| TD-51306 | On upgrade, Spark is incorrectly parsing files of type "UTF-8 Unicode (with BOM)."  |
| TD-51218 | Import rules not working for remapping of WASB bucket name. For more information, see <i>Define Import Mapping Rules</i> .  |
| TD-51166 | Cannot import flow due to missing associated flownode error.  |
| TD-50945 | Server Save error when deleting a column.   |
| TD-50906 | Transformation engine unavailable due to prior crash  |
| TD-50791 | After upgrade, you cannot edit recipes or run jobs on recipes that contain the optional <code>replaceOn</code> parameter is not used in Replace transformation.   |
| TD-50703 | Optional file cleanup generates confusing error logging when it fails.  |
| TD-50642 | When modifying file privileges, the platform makes assumptions about database usernames.  |
| TD-50530 | On upgrade, the migration framework for the authorization service is too brittle for use with Amazon RDS database installations.  |
| TD-50525 | When flows are imported into the Deployment Manager, additional characters are inserted into parameterized output paths, causing job failures.  |
| TD-50522 | PostgreSQL connections may experience out of memory errors due to incorrectly specified fetch size and vendor configuration.  |
| TD-50516 | Can't import a flow that contains a reference in a flow webhook task to a deleted output.   |
| TD-50508 | Generic Hadoop folder is missing in <code>hadoop-deps</code> folder.  |
| TD-50496 | After upgrade, you cannot publish as a single-file to WASB to replace an existing output destination.   |
| TD-50495 | After upgrade, users cannot load recipes due to Requested Data Not Found error when loading samples.  |
| TD-50466 | <p>After upgrading Cloudera cluster to version 6.3.3, you cannot run jobs due to the following error:</p> <pre>class not found exception: java.lang.NoClassDefFoundError: org/apache/spark/sql/execution/datasources/csv/CSVOptions</pre> <p>Please see "Cloudera support" above.</p> |
| TD-50446 | <p>During upgrade, cross-migration fails for authorization service and its database with the following error:</p> <pre>Cross migration failed. Make sure the authorization DB is reset.</pre>   |
| TD-50164 | After upgrade, ad-hoc publish to Hive fails.  |
| TD-49991 | After upgrade, you cannot unzip downloaded log files.   |
| TD-49973 | After upgrade, cross-migration validation fails for "groupsPolicies."   |
| TD-49692 | Tripache Vulnerabilities - CVE-2020-1927  |

## New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-59854 | Datetime column from Parquet file incorrectly inferred to the wrong data type on import.<br><br><b>Workaround:</b> Use the column drop-down to change the data type to Datetime.                |
| TD-51229 | When an admin user shares a flow that the admin user owns, a <code>Failed to share flow with selected user</code> error message may be displayed, even though the flow was successfully shared. |

## Release 7.1

May 4, 2020

### What's New

#### In-app chat:

**Have a question about the product?** Use the new in-app chat feature to explore content or ask a question to our support staff. If you need assistance, please reach out!

#### Troubleshooting:

- Users can download log files related to their current session through the application. See *Download Logs Dialog*.
  - Administrators have a separate admin dialog that enables log download by time frame, job identifier, or session identifier. See *Admin Download Logs Dialog*.

#### Install:

**NOTE:** If you are installing or upgrading a deployment of Trifacta that uses or will use a remote database service, such as Amazon RDS, for hosting the Trifacta databases, please contact *Alteryx Customer Success Services*. For this release, additional configuration may be required.

- Support for installation on CentOS/RHEL 8. See *System Requirements*.

**NOTE:** SSO using SAML is not supported on CentOS/RHEL 8. See *Configure SSO for SAML*.

**NOTE:** Support for CentOS/RHEL 6 has been deprecated. Please upgrade to CentOS/RHEL 8.

- Support for installation on CentOS/RHEL 7.7. See *System Requirements*.
- Support for EMR 5.28.1 and EMR 5.29.0

**NOTE:** EMR 5.28.0 is not supported, due to *Spark compatibility issues*.

**NOTE:** Support for EMR 5.8 - EMR 5.12 is deprecated. For more information, see *End of Life and Deprecated Features*.



- Support for Azure Databricks 6.2. See *Configure for Azure Databricks*.
- Support for installation on Ubuntu 18.04 (Bionic Beaver). See *System Requirements*.

**NOTE:** Support for installation on Ubuntu 14.04 (Trusty) has been deprecated. See *End of Life and Deprecated Features*.

- Support for CDH 6.0 is deprecated. See *End of Life and Deprecated Features*.

#### Spark:

- Support for Spark 2.2.x versions is deprecated. See *End of Life and Deprecated Features*.
- Improved performance for Spark profiling on Datetime and numeric columns with low number of discrete values.

#### Kerberos:

- Support for access to Kerberized clusters. See *Configure for EMR*.

#### Connectivity:

- Improved performance for Oracle, SQL Server, and DB2 connections. These performance improvements will be applied to other relational connections in future releases.

**NOTE:** For more information on enabling this feature, please contact *Alteryx Customer Success Services*.

- Azure Databricks Tables:
  - Support for read/write on Delta tables.
  - Support for read/write on external tables.
  - Support for read from partitioned tables.

**NOTE:** To enable these additional read/write capabilities through Databricks Tables, the underlying connection was changed to use a Simba driver. In your connection definition, any Connect String Options that relied on the old Hive driver may not work. For more information, see *Configure for Azure Databricks*.

#### Import:

- Ingestion of large relational datasets is no longer a blocking operation. For more information, see *Configure JDBC Ingestion*.
  - Track progress of large-scale ingestion in Flow View and the Library page.
    - See *Flow View Page*.
    - See *Import Data Page*.

#### Workspace:

- Redesigned Settings and Help menus. See *Home Page*.
  - User settings are now modified through Preferences. See *Preferences Page*.
  - Administrators now have a dedicated admin area. See *Admin Console*.

#### Plans:

- Introducing plans. A plan is a sequence of tasks on one or more flows that can be scheduled.

**NOTE:** In this release, the only type of task that is supported is Run Flow.

- For more information on plans, see *Plans Page*.
- For more information on orchestration in general, see *Overview of Operationalization*.

### Flow View:

- Introducing new Flow View. The Flow View page has been redesigned to improve the user experience and overall productivity.

**NOTE:** This feature is in Beta release.

- Enhancements include:
  - Drag and drop to reposition objects on the Flow View canvas, and zoom in and out to focus on areas of development.
  - Perform joins and unions between objects on the Flow View canvas.
  - Annotate the canvas with notes.
- You can toggle between new and classic views through the context menu in the corner of Flow View. See *Flow View Page*.
- As needed, Trifacta administrators can disable access to the new Flow View completely. See *Miscellaneous Configuration*.
- Create flow parameters that you can reference in your flow. Flow parameters can be string literals, Trifacta patterns, or regular expression patterns.

**NOTE:** For this release, flow parameters can be applied into your recipes only.

- As needed, you can apply overrides to the parameters in your flow or to downstream flows.

**NOTE:** Flow parameters do not apply to datasets or output objects, which have their own parameters. However, if you specify an override at the flow level, any parameters within the flow that use the same name receive the override value, including output object parameters and datasets with parameters.

- See *Manage Parameters Dialog*.
- For more information on parameters, see *Overview of Parameterization*.
- Monitor job progress through each phase in the Jobs panel. See *Flow View Page*.

### Transformer Page:

- Improved performance when loading the Transformer page and when navigating between the Flow View and Transformer pages.
- Join steps are now created in a larger window for more workspace. See *Join Window*.
- New column selection UI simplifies choosing columns in your transformations. See *Transform Builder*.
- Faster and improved method of surfacing transform suggestions based on machine learning.

### Job Execution:

**NOTE:** Azure Databricks 5.3 and 5.4 are no longer supported. Please upgrade to Azure Databricks 5.5 LTS or 6.x. See *End of Life and Deprecated Features*.

- Apply overrides to Spark properties for individual job execution. See *Enable Spark Job Overrides*.
- Execute jobs from SFTP sources on EMR and Azure Databricks. See *SFTP Connections*.

#### Job Details:

- When visual profiling is enabled for a job, you can now download your visual profile in PDF format. See *Job Details Page*.

#### Publishing:

- Support for generating results and publishing to Tableau Hyper format.

**NOTE:** Tableau TDE format will be deprecated in a future release. Please switch to using Tableau Hyper format.

- If you have upgraded to Tableau Server 10.5 or later, you may have a mix of TDE and Hyper files stored on the server. You can automatically upgrade the TDE files to Hyper, if needed. For more information, see [https://help.tableau.com/current/online/en-us/extracting\\_upgrade.htm](https://help.tableau.com/current/online/en-us/extracting_upgrade.htm).
- If you are on Tableau Server 10.5 or later and you append to a TDE file, the file is automatically converted to Hyper format. This conversion cannot be reverted.
- See *Tableau Server Connections*.

#### Language:

- New functions to parse values against specific data types.
- New functions for calculating working days between two valid dates.
- New two-column statistical functions.
- See *Changes to the Language*.

#### Documentation:

- New content on the getting started with sampling. See *Sampling Basics*.
  - Feature overview: *Overview of Sampling*
  - Best practices: <https://community.trifacta.com/s/article/Best-Practices-Managing-Samples-in-Complex-Flows>

#### Changes in System Behavior

##### Wrangler Enterprise desktop application:

**The Wrangler Enterprise desktop application is no longer available in the software distribution and has been deprecated. Please switch to a supported browser version. For more information, see *Browser Requirements*.**

**A Release 6.8 version of the Wrangler Enterprise desktop application can be made available upon request. For more information, please contact *Alteryx Support*.**

#### Authorization:

- All Trifacta admin users are now workspace admins.
- All workspace admins now have access to all user-created objects within the workspace.

**NOTE:** Workspace administrators can access some types of user-created objects in the workspace with the same level of access as the object owner. Under some conditions, workspace

admins may have access to source datasets and generated results. See *Workspace Admin Permissions*.

- For more information, see *Changes to User Management*.

#### API Documentation:

- API reference documentation is now available directly through the application. This release includes more supported endpoints and documented options. To access, select **Help menu > API Documentation**.

**NOTE:** API reference content is no longer available with the product documentation. Please use the in-app reference documentation instead.

- Workflow documentation is still available with the product documentation. For more information, see *API Reference*.
- For details, see *Changes to the APIs*.

#### Trifacta node:

- Upgrade to NodeJS 12.16.1.

**NOTE:** This dependency is specific to the Trifacta platform. For this release, a separate installation of Trifacta dependencies is required for installing or upgrading the platform.

See *Install on CentOS and RHEL*.

See *Install on Ubuntu*.

- See *System Requirements*.
- See *System Dependencies*.

#### APIs:

- The v3 version of the API endpoints are no longer available in the platform. You must use v4 endpoints. See *API Reference*.
  - Simplified connections endpoints.
- The format of the supported WASB URIs has changed.

**NOTE:** If you were using the APIs to interact with WASB resources, you must update your resources to use the new format. See *Changes to the APIs*.

- See *Changes to the APIs*.

#### Custom dictionaries:

In a future release, custom dictionaries that rely on an uploaded file will be deprecated. The specific release vehicle has not been determined yet.

- Deprecation only affects the ability to create custom types using a file. Where possible, you can and should continue to create custom times using regular expressions. For more information, see *Create Custom Data Types Using RegEx*.
- The file-based feature will be replaced by a standardization-based option.
- Beginning in this release, this feature is disabled by default.

#### Parameter overrides:

- If you have upgraded to Release 7.1 or later, any parameter overrides that you have specified in your flows can be modified in the Overrides tab of the Manage Parameters dialog.
- For more information, see *Manage Parameters Dialog*.

### WASB and ADLS:

- Configuration to enable WASB and ADLS access has been streamlined and simplified.

**NOTE:** No action is required for upgrading customers.

See *WASB Access*.

See *ADLS Gen1 Access*.

### Secure Token Service:

- The default port number for the secure token service has been changed from 8090. The new default port number is 41921.

**NOTE:** Your upgraded installation is forced to use this new port number. You can modify the value after installation or upgrade.

### Sharing:

- The Send a Copy feature is no longer available in the product. Instead, you can make a copy of the flow and share it. See *Flow View Page*.

### Language:

- All MODE functions return the lowest value in a set of values if there is a tie in the evaluation. See *Changes to the Language*.

### Key Bug Fixes

| Ticket   | Description   |
|----------|---|
| TD-48245 | <p>By default, under SSO manual logout and session expiration logout redirect to different pages. Manual logout directs you to SAML sign out, and session expiry produces a session expired page.</p> <p>To redirect the user to a different URL on session expiry, an administrator can set the following parameter: <code>webapp.session.redirectUriOnExpiry</code>. This parameter applies to the following SSO environments:</p> <ul style="list-style-type: none"> <li>• <i>Configure SSO for SAML</i></li> <li>• <i>Configure SSO for Azure AD</i></li> </ul> |

### New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-52221 | <p>You cannot update your AWS configuration for per-user or per-workspace mode via UI.</p> <div> <p><b>Workaround:</b> You can switch to using AWS system mode with a single, system wide configuration, or you can use the APIs to make changes. See <i>API Workflow - Manage AWS Configurations</i>.</p> </div> |
| TD-      | Cannot select and apply custom data types through column Type menu.   |

|          |   |
|----------|---|
| 49559    | <p><b>Workaround:</b> You can change the type of the column as a recipe step. Use the Change column type transformation. From the New type drop-down, select <code>Custom</code>. Then, enter the name of the type for the Custom type value.</p> |
| TD-47784 | When creating custom datasets using SQL from Teradata sources, the <code>ORDER BY</code> clause in standard SQL does not work.  |
| TD-47473 | <p>Uploaded files (CSV, XLS, PDF) that contain a space in the filename fail to be converted.</p> <p><b>Workaround:</b> Remove the space in the filename and upload again.</p>   |

# Release Notes 6.8

## Contents:

- *Release 6.8.2*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 6.8.1*
    - *What's New*
    - *Changes to System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 6.8*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
- 

## Release 6.8.2

April 27, 2020

### What's New

#### Import:

- Enhanced full-screen interface for importing using custom SQL. See *Create Dataset with SQL*.

### Changes in System Behavior

None.

### Key Bug Fixes

| Ticket   | Description  |
|----------|--|
| TD-48245 | <p>By default, under SSO manual logout and session expiration logout redirect to different pages. Manual logout directs you to SAML sign out, and session expiry produces a session expired page.</p> <p>To redirect the user to a different URL on session expiry, an administrator can set the following parameter: <code>webapp.session.redirectUriOnExpiry</code>. This parameter applies to the following SSO environments:</p> <ul style="list-style-type: none"><li>• <i>Configure SSO for SAML</i></li><li>• <i>Configure SSO for Azure AD</i></li></ul> |

### New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-48630 | <p>Connection files used by the data service are not persisted in a Dockerized environment.</p> <div></div> |

|          |   |
|----------|---|
|          | <p><b>Workaround:</b> In the Admin Settings page, set <code>data-service.vendor</code> to a location that is persisted. Example value:</p> <pre>(Path-to-persistent-directory)/conf/data-service/application.properties</pre>   |
| TD-47696 | <p>Platform appears to fail to restart properly through Admin Settings page due to longer restarts of individual services. Symptoms:</p> <ul style="list-style-type: none"> <li>• Changes to settings may appear to have not been applied.</li> <li>• Admin Settings page appears to be stuck restarting.</li> </ul> <p><b>Workaround:</b> Restart can take up to several minutes. If the restart does not appear to complete, try reloading the page. If that doesn't work, restarting from the command line is more reliable. See <i>Start and Stop the Platform</i>.</p> |

## Release 6.8.1

February 7, 2020

This release enables some new features and makes some relational connections generally available.

### What's New

#### Install:

- Support for CDH 6.3. See *Supported Deployment Scenarios for Cloudera*.

**NOTE:** Support for CDH 6.0 has been deprecated. See *End of Life and Deprecated Features*.

#### Import:

- Upload tabular data from PDF documents.

**NOTE:** This feature is in Beta release.

**NOTE:** This feature must be enabled.

See *Import PDF Data*.

- Read support for ORC tables managed through Hive. See *Configure for Hive*.

#### LDAP:

- Support for initial binding to active directory using the user's account. See *Configure SSO for AD-LDAP*.

#### Cluster Clean:

- Cluster Clean standardization feature is now available in all product editions. See *Overview of Cluster Clean*.

#### Documentation:



- API: Improved documentation for the asset transfer endpoint. See *Changes to the APIs*.

## Changes to System Behavior

### Wrangler Enterprise desktop application:

**NOTE:** In a future release, the Wrangler Enterprise desktop application will be deprecated. Please switch to a supported version of Google Chrome or Mozilla Firefox. Support for Edge Chromium is expected in a future release. See *Browser Requirements*.

### General availability:

- The following relational connections are now generally available:
  - DB2 (import only)
  - Salesforce (import only)
  - Tableau Server (publish only)
 For more information, see *Connection Types*.

## Key Bug Fixes

| Ticket   | Description  |
|----------|--|
| TD-45492 | Publishing to Databricks Tables fails on ADLS Gen1 in user mode. |

## New Known Issues

| Ticket   | Description  |
|----------|--|
| TD-47263 | <p>Importing an exported flow that references a Google Sheets or Excel source breaks connection to input source.</p> <div> <p><b>Workaround:</b> If the importing user has access to the source, the user can re-import the dataset and then swap the source for the broken recipe.</p> </div> |

## Release 6.8

December 6, 2019

Welcome to Release 6.8 of Trifacta®. This release introduces several key features around operationalizing the platform across the enterprise. Enterprise stakeholders can now receive email notifications when recurring jobs have succeeded or failed, updating data consumers outside of the platform. This release also introduces a generalized webhook interface, which facilitates push notifications to applications such as Slack when jobs have completed. When jobs fail, users can download a much richer support bundle containing configuration files, script files, and a specified set of log files.

Macros have been expanded to now be export- and import-ready across environments. In support of this feature, the Wrangle Exchange is now available through the Trifacta Community, where you can download macros created by others and import them for your own use. Like macros, you can now export and import flows across product editions and release (Release 6.8 or later only).

In the application, you can now use shortcut keys to navigate around the workspace and the Transformer page. And support for the Firefox browser has arrived. Read on for more goodness added with this release.

## What's New

### Install:

- Support for ADLS Gen2 blob storage. See *ADLS Gen2 Access*.

### Workspace:

- Individual users can now enable or disable keyboard shortcuts in the workspace or Transformer page. See *User Profile Page*.
- Configure locale settings at the workspace or user level. See *Locale Settings*.
- You can optionally duplicate the datasets from a source flow when you create a copy of it. See *Flow View Page*.
- Create a copy of your imported dataset. See *Library Page*.

### Browser:

- Support for Firefox browser.

**NOTE:** This feature is in Beta release.

For supported versions, see *Browser Requirements*.

### Project Management:

- Support for export and import of macros. See *Macros Page*.
  - For more information on macros, see *Overview of Macros*.
- Download and use macros available through the Wrangle Exchange. See <https://www.trifacta.com/blog/crowdsourcing-macros-trifacta-wrangle-exchange/>.

### Operationalization:

- Create webhook notifications for third-party platforms based on results of your job executions. See *Create Flow Webhook Task*.
- Enable and configure email notifications based on the success or failure of job executions.

**NOTE:** This feature requires access to an SMTP server. See *Enable SMTP Email Server Integration*.

- For more information on enabling, see *Workspace Settings Page*.
  - Individual users can opt out of receiving email messages or can configure use of a different email address. See *Email Notifications Page*.
- For more information on enabling emails for individual flows, see *Manage Flow Notifications Dialog*.

### Supportability:

- Download logs bundle on job success or failure now contains extensive configuration information to assist in debugging. For more information, see *Configure Support Bundling*.

### Connectivity:

- Support for integration with EMR 5.8 - 5.27. For more information, see *Configure for EMR*.
- Connect to SFTP servers to read data and write datasets. See *SFTP Connections*.

- Create connections to Databricks Tables.

**NOTE:** This connection is supported only when the Trifacta platform is connected to an Azure Databricks cluster.

For more information, see *Databricks Tables Connections*.

- Support for using non-default database for your Snowflake stage.
  - Support for ingest from read-only Snowflake databases.
  - See *Enable Snowflake Connections*.

#### Import:

- As of Release 6.8, you can import an exported flow into any edition or release after the build number of the export. See *Import Flow*.
- Improved monitoring of long-loading relational sources. See *Import Data Page*.

**NOTE:** This feature must be enabled. See *Configure JDBC Ingestion*.

#### Transformer Page:

- Select columns, functions applied to your source, and constants to replace your current dataset. See *Select*.
- Improved Date/Time format selection. See *Choose Datetime Format Dialog*.

**Tip:** Datetime formats in card suggestions now factor in the user's locale settings for greater relevance.

- Improved matching logic and performance when matching columns through RapidTarget.
  - Align column based on the data contained in them, in addition to column name.
  - This feature is enabled by default. For more information, see *Overview of RapidTarget*.
- Improvements to the Search panel enable faster discovery of transformations, functions, and other objects. See *Search Panel*.

#### Job execution:

- By default, the Trifacta application permits up to four jobs from the same flow to be executed at the same time. If needed, you can configure the application to execute jobs from the same flow one at a time. See *Configure Application Limits*.
- If you enabled visual profiling for your job, you can download a JSON version of the visual profile. See *Job Details Page*.
- Support for instance pooling in Azure Databricks. See *Configure for Azure Databricks*.

#### Language:

- New trigonometry and statistical functions. See *Changes to the Language*.

#### API:

- Apply overrides at time of job execution via API.
- Define import mapping rules for your deployments that use relational sources or publish to relational targets.

- Export and import macro definitions.
- See *Changes to the APIs*.

## Changes in System Behavior

### Browser Support Policy:

- For supported browsers, at the time of release, the latest stable version and the two previous stable versions are supported.

**NOTE:** Stable browser versions released after a given release of Trifacta will **NOT** be supported for any prior version of Trifacta. A best effort will be made to support newer versions released during the support lifecycle of the release.

For more information, see *Browser Requirements*.

### Install:

**NOTE:** In the next release of Trifacta after Release 6.8, support for installation on CentOS/RHEL 6.x and Ubuntu 14.04 will be deprecated. You should upgrade the Trifacta node to a supported version of CentOS /RHEL 7.x or Ubuntu 16.04. Before performing the upgrade, please perform a full backup of the Trifacta platform and its databases. See *Backup and Recovery*.

- Support for Spark 2.1 has been deprecated. Please upgrade to a supported version of Spark.
  - Support for EMR 5.6 and eMR 5.7 has also been deprecated. Please upgrade to a supported version of EMR.
  - For more information, see *Product Support Matrix*.
- To simplify the installation distribution, the Hadoop dependencies for the recommended version only are included in the software download. For the dependencies for other supported Hadoop distributions, you must download them from the Trifacta FTP site and install them on the Trifacta node. See *Install Hadoop Dependencies*.
- Trifacta node has been upgraded to use Python 3. This instance of Python has no dependencies on any Python version external to the Trifacta node.

### Import/Export:

- Flows can now be exported and imported across products and versions of products. See *Changes to the Object Model*.

### CLI and v3 endpoints (Release 6.4):

**NOTE:** Do not attempt to connect to the Trifacta platform using any version of the CLI or the v3 endpoints. They are no longer supported and unlikely to work.

### In Release 6.4:

- The Command Line Interface (CLI) was deprecated. Customers must use the v4 endpoints for the APIs instead.
- The v3 versions of the API endpoints were deprecated. Customers must use the v4 endpoints for the APIs instead.
- Developer content was provided to assist in migrating to the v4 endpoints.
- For more information on acquiring this content, please contact *Alteryx Support*.

### Key Bug Fixes

| Ticket | Description |
|--------|-------------|
|        |             |

|          |   |
|----------|---|
| TD-40348 | When loading a recipe in an imported flow that references an imported Excel dataset, Transformer page displays <b>Input validation failed: (Cannot read property 'filter' of undefined)</b> error, and the screen is blank. |
| TD-42080 | Cannot run flow or deployment that contains more than 10 recipe jobs  |

## New Known Issues

| Ticket   | Description  |
|----------|--|
| TD-46123 | <p>Cannot modify the type of relational target for publishing action.</p> <div> <b>Workaround:</b> Create a new publishing action with the desired relational target. Remove the original one if necessary. See <i>Run Job Page</i>. </div>  |
| TD-45923 | Publishing a compressed Snappy file to SFTP fails.   |
| TD-45922 | You cannot publish TDE format to SFTP destinations.  |
| TD-45492 | Publishing to Databricks Tables fails on ADLS Gen1 in user mode.   |
| TD-45273 | <p>Artifact Storage Service fails to start on HDP 3.1.</p> <div> <b>Workaround:</b> The Artifact Storage Service can reference the HDP 2.6 Hadoop bundle JAR. </div> <p><b>Steps:</b></p> <ol style="list-style-type: none"> <li>1. You can apply this change through the <i>Admin Settings Page</i> (recommended) or <code>trifacta-conf.json</code>. For more information, see <i>Platform Configuration Methods</i>.</li> <li>2. Locate the following property: <div> <code>"artifact-storage-service.classpath"</code> </div> </li> <li>3. Replace this value: <div> <code>:%(topOfTree)s/%(hadoopBundleJar)s</code> </div> </li> <li>4. With the following: <div> <code>:%(topOfTree)s/conf/hadoop-site/:%(topOfTree)s/hadoop-deps/hdp-2.6/build/libs/hdp-2.6-bundle.jar</code> </div> </li> <li>5. Save changes and restart the platform.</li> </ol> |
| TD-45122 | <p>API: re-running job using only the <code>wrangleDataset</code> identifier fails even if the original job succeeds when <code>writeSettings</code> were specified.</p> <div> <b>Workaround:</b> Use a full <code>jobGroups</code> job specification each time that you run a job. </div> <p>See <a href="https://api.trifacta.com/ee/es.t/index.html#operation/runJobGroup">https://api.trifacta.com/ee/es.t/index.html#operation/runJobGroup</a></p>  |
| TD-44429 | Cannot publish outputs to relational targets, receiving <code>Encountered error while processing stream</code> .   |

**Workaround:** This issue may be caused by the `trifacta` service account not having write and execute permissions to the `/tmp` directory on the Trifacta node.

If so, you can do either of the following:

1. Enable write and execute permissions for the account on `/tmp`.
2. Create a new temporary account and provide the service account write and execute permissions to it. Then, add the following to `data-service.jvmOptions`:

```
-Dorg.xerial.snappy.tmpdir=/new/directory/with/writeexecuteaccess
```

TD-44427

Cannot publish dataset containing duplicate rows to Teradata. Error message:

```
Caused by: java.sql.SQLException: [Teradata Database] [TeraJDBC 15.10.00.14] [Error -2802]
[SQLState 23000] Duplicate row error in abc_trifacta.tmp_218768523.
at
```

**Workaround:** This is a known limitation on Teradata. For more information on this limitation, see *Enable Teradata Access*.

# Release Notes 6.4

## Contents:

- *Release 6.4.2*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 6.4.1*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
  - *Release 6.4*
    - *What's New*
    - *Changes in System Behavior*
    - *Key Bug Fixes*
    - *New Known Issues*
- 

## Release 6.4.2

November 15, 2019

This release is primarily a bug fix release with the following new features.

### What's New

#### API:

- Apply overrides at time of job execution via API.
- Define import mapping rules for your deployments that use relational sources or publish to relational targets.
- See *Changes to the APIs*.

#### Job execution:

- By default, the Trifacta application permits up to four jobs from the same flow to be executed at the same time. If needed, you can configure the application to execute jobs from the same flow one at a time. See *Configure Application Limits*.

### Changes in System Behavior

None.

### Key Bug Fixes

| Ticket   | Description  |
|----------|--|
| TD-44548 | RANGE function returns null values if more than 1000 values in output. |
| TD-44494 | Lists are not correctly updated in Deployment mode                     |

|          |  |
|----------|--|
| TD-44311 | Out of memory error when running a flow with many output objects   |
| TD-44188 | Performance is poor for SQL DW connection  |
| TD-43877 | Preview after a DATEFORMAT step does not agree with results or profile values  |
| TD-44035 | Spark job failure from Excel source  |
| TD-43849 | Export flows are broken when recipe includes Standardization or Transform by Example tasks.<br><br><b>NOTE:</b> This Advanced Feature is available in Trifacta under a separate, additional license. If it is not available under your current license, do not enable it for use. Please feel free to contact your representative. |

## New Known Issues

| Ticket   | Description  |
|----------|--|
| TD-46185 | Stepping backward to an early step in a recipe sometimes fails to properly update the state of the quality bar and histograms in the data grid.<br><br><b>Workaround:</b> This issue is caused by caching of snapshot profiles from the data grid. The workaround is to reload the page through the browser. |

## Release 6.4.1

August 30, 2019 This release includes bug fixes and introduces SSO connections for Azure relational sources.

## What's New

### Connectivity:

- You can now leverage your Azure AD SSO infrastructure to create SSO connections to Azure relational databases. For more information, see *Enable SSO for Azure Relational Connections*.

## Changes in System Behavior

### Configuration changes:

- The parameter to enable custom SQL query has been moved to the Workspace Settings page.
- The parameter to disable schematized output has been moved to the Workspace Settings page.
- For more information, see *Changes to Configuration*.

## Key Bug Fixes

| Ticket   | Description                               |
|----------|---|
| TD-39086 | Hive ingest job fails on Microsoft Azure. |

## New Known Issues

None.



## Release 6.4

August 1, 2019

This release of Trifacta® features broad improvements to the recipe development experience, including multi-step operations and improved copied and paste within the Recipe panel. As a result of the panel's redesign, you can now create user-defined macros, which are sets of sequenced and parameterized steps for easy reuse and adaptation for other recipes. When jobs are executed, detailed monitoring provides enhanced information on progress of the job through each phase of the process. You can also connect to a broader ecosystem of sources and targets, including enhancements to the integration with Tableau Server and AWS Glue. New for this release: read from your Snowflake sources. Read on for additional details on new features and enhancements.

### What's New

#### Transformer Page:

- The redesigned Recipe panel enables multi-step operations and more robust copy and paste actions. See *Recipe Panel*.
- Introducing user-defined macros, which enable saving and reusing sequences of steps. For more information, see *Overview of Macros*.
- Transform by example output values for a column of values. See *Transformation by Example Page*.
  - For an overview of this feature, see *Overview of TBE*.
- Browse current flow for datasets or recipes to join into the current recipe. See *Join Window*.
- Replace specific cell values. See *Replace Cell Values*.

#### Job Execution:

- Detailed job monitoring for ingest and publishing jobs. See *Overview of Job Monitoring*.
- Parameterize output paths and table and file names. See *Run Job Page*.

#### Install:

- Support for RHEL/CentOS 7.5 and 7.6 for the Trifacta node. See *System Requirements*.
- Support for deployment of Trifacta platform via Docker image. See *Install for Docker*.

#### Connectivity:

- Support for integration with Cloudera 6.2.x. See *System Requirements*.

**NOTE:** Support for integration with Cloudera 5.15.x and earlier has been deprecated. See *End of Life and Deprecated Features*.

**NOTE:** Support for integration with HDP 2.5.x and earlier has been deprecated. See *End of Life and Deprecated Features*.

- Support for Snowflake database connections.

**NOTE:** This feature is supported only when Trifacta is installed on customer-managed AWS infrastructure.

For more information, see *Snowflake Access*.

- Support for direct publishing to Tableau Server. For more information, see *Run Job Page*.
- Support for MySQL database timezones. See *Install Databases for MySQL*.

Enhanced support for AWS Glue integration:

- Metadata catalog browsing through the application.
- Per-user authentication to Glue. See *Configure AWS Per-User Auth for Temporary Credentials*.
- See *AWS Glue Access*.

**Import:**

- Add timestamp parameters to your custom SQL statements to enable data import relative to the job execution time. See *Create Dataset with SQL*.

**Authentication:**

- Leverage your enterprise's SAML identity provider to pass through a set of IAM roles that Trifacta users can select for access to AWS resources.

**NOTE:** This authentication method is supported only if SSO authentication has been enabled using the platform-native SAML authentication method. For more information, see *Configure SSO for SAML*.

For more information, see *Configure for AWS SAML Passthrough Authentication*.

- Support for AzureManaged Identities with Azure Databricks. See *Configure for Azure Databricks*.

**Admin:**

- Administrators can review, enable, disable, and delete schedules through the application. See *Schedules Page*.

**Sharing:**

- Share flows and connections with groups of users imported from your LDAP identity provider.

**NOTE:** This feature is in Beta release.

See *Configure Users and Groups*.

**Logging:**

- Tracing user information across services for logging purposes. See *Configure Logging for Services*.

**Language:**

- New functions. See *Changes to the Language*.
- Broader support for metadata references. For Excel files, `$filepath` references now return the location of the source Excel file. Sheet names are appended to the end of the reference. See *Source Metadata References*.

**APIs:**

- Admins can now generate password reset requests via API. See *Changes to the APIs*.

**Databases:**

- New databases:
  - Job Metadata Service database

## Changes in System Behavior

**NOTE:** The Trifacta software must now be installed on an edge node of the cluster. Existing customers who cannot migrate to an edge node will be supported. You will be required to update cluster files on the Trifacta node whenever they change, and cluster upgrades may be more complicated. You should migrate your installation to an edge node if possible. For more information, see *System Requirements*.

**NOTE:** The v3 APIs are no longer supported. Please migrate immediately to using the v4 APIs.

**NOTE:** The command line interface (CLI) is no longer available. Please migrate immediately to using the v4 APIs.

**NOTE:** The PNaCl browser client extension is no longer supported. Please verify that all users of Trifacta are using a supported version of Google Chrome, which automatically enables use of WebAssembly. For more information, see *Browser Requirements*.

**NOTE:** Support for Java 7 has been deprecated in the platform. Please upgrade to Java 8 on the Trifacta node and any connected cluster. Some versions of Cloudera may install Java 7 by default.

**NOTE:** The **Chat with us** feature is no longer available. For Trifacta customers, this feature had to be enabled in the product. For more information, see *Alteryx Support*.

**NOTE:** The desktop version of Trifacta Desktop Wrangler will cease operations on August 31, 2019. If you are still using the product at that time, your data will be lost. Please transition to using the free Cloud version of Trifacta® Wrangler. Automated migration is not available. To register for a free account, please visit <https://cloud.trifacta.com>.

## Workspace:

- Configuration for AWS authentication for platform users has been migrated to a new location. See *Configure Your Access to S3*.

## API:

- The endpoint used to assign an AWSConfig object to a user has been replaced.

**NOTE:** If you used the APIs to assign AWSConfig objects in a previous release, you must update your scripts to assign AWS configurations. For more information, see *Changes to the APIs*.

## Documentation:

- In prior releases, the documentation listed UTF32-BE and UTF32-LE as supported file formats. These formats are not supported. Documentation has been updated to correct this error. See *Supported File Encoding Types*.

## Key Bug Fixes

| Ticket   | Description   |
|----------|---|
| TD-41260 | Unable to append Trifacta Decimal type into table with Hive Float type. See <i>Hive Data Type Conversions</i> .   |
| TD-40424 | UTF-32BE and UTF-32LE are available as supported file encoding options. They do not work.<br><div><b>NOTE:</b> Although these options are available in the application, they have never been supported in the underlying platform. They have been removed from the interface.</div> |
| TD-40299 | Cloudera Navigator integration cannot locate the database name for JDBC sources on Hive.  |
| TD-40243 | API access tokens don't work with native SAML SSO authentication  |
| TD-39513 | Import of folder of Excel files as parameterized dataset only imports the first file, and sampling may fail.  |
| TD-39455 | HDI 3.6 is not compatible with Guava 26.  |
| TD-39092 | <code>\$filepath</code> and <code>\$sourcerownumber</code> references are not supported for Parquet file inputs.<br>For more information, see <i>Source Metadata References</i> .   |
| TD-31354 | When creating Tableau Server connections, the Test Connection button is missing. See <i>Tableau Server Connections</i> .  |
| TD-36145 | Spark running environment recognizes numeric values preceded by + as Integer or Decimal data type. Photon running environment does not and types these values as strings.   |

## New Known Issues

| Ticket   | Description   |
|----------|---|
| TD-42638 | Publishing and ingest jobs that are short in duration cannot be canceled.<br><div><b>Workaround:</b> Allow the job to complete. You can track the progress through these phases of the jobs through the application. See <i>Job Details Page</i>.</div> |
| TD-39052 | Changes to signout on reverse proxy method of SSO do not take effect after upgrade.   |

# End of Life and Deprecated Features

## Contents:

- *Release 9.2*
- *Release 9.1*
- *Release 8.11*
- *Release 8.7*
- *Release 8.2*
- *Release 8.1*
- *Release 7.9*
- *Release 7.8*
- *Release 7.7*
- *Release 7.6*
- *Release 7.1*
- *Release 7.0*
- *Release 6.10*
- *Release 6.8.1*
- *Release 6.8*
- *Release 6.4*
- *Release 6.0.2*
- *Release 6.0.1*
- *Release 6.0*
- *Release 5.1*
- *Release 5.0*
- *Release 4.2.1*
- *Release 4.2*
- *Release 4.1.1*
- *Release 4.1*
- *Release 4.0.1*
- *Release 4.0*
- *Release 3.2.1*
- *Release 3.2*
- *Release 3.1.2*
- *Release 3.1.1*
- *Release 3.1*
- *Release 3.0*
- *Release 2.7*

The following features have been deprecated or removed from Trifacta® in recent releases.

## Legend:

| Status      | Description  |
|-------------|--|
| deprecated  | Feature or capability is no longer actively supported. It may still work, but future fixes or enhancements are unlikely. |
| end of life | Feature or capability has been removed from the product.   |

## Release 9.2

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item             | Status  | Description  |
|------------------|---------|--|
| Hortonworks Data | Depreca | Integration with Hortonworks Data Platform (HDP) clusters has been deprecated. |

|                                    |          |   |
|------------------------------------|----------|---|
| Platform (HDP)                     | ted      | <p><b>NOTE:</b> As part of the above deprecation, the following Hortonworks-only integrations are no longer available for the Trifacta platform:</p> <ul style="list-style-type: none"> <li>• Hive Warehouse Connector</li> <li>• High Availability for Hive</li> <li>• Zookeeper Quorum</li> </ul> <p>The following Cloudera platforms and versions are still supported for Release 9.2:</p> <ul style="list-style-type: none"> <li>• CDH 6.2, CDH 6.3</li> <li>• CDP 7.1</li> </ul> <p>For more information, see <i>Product Support Matrix</i>.</p> |
| High Availability on AWS using EKS | Deprecat | <p>The method for deploying the Trifacta platform in a highly available environment on AWS using EKS (Kubernetes) clusters has been deprecated.</p> <p>The existing method of deploying using failover nodes is still supported for AWS. For more information, see <i>Install for High Availability</i>.</p>  |
| Custom Dictionary types            | Deprecat | <p>The ability to add custom data types using a dictionary file has been deprecated. Replacement options:</p> <ul style="list-style-type: none"> <li>• Create a custom data type using regular expressions method. See <i>Create Custom Data Types Using RegEx</i>.</li> <li>• Create a reference dataset and validate against using a left join. See <i>Validate Column Values against a Dataset</i>.</li> </ul>   |

## Release 9.1

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item   | Status   | Description  |
|--|----------|--|
| Databricks 8.3   | Deprecat | <p>Databricks no longer supports this version. Options:</p> <ul style="list-style-type: none"> <li>• Downgrade to Databricks 7.3 LTS</li> <li>• Upgrade to Databricks 10.x.</li> </ul> <p><b>NOTE:</b> If you have enabled user mode for ADLS Gen2, you must upgrade to Databricks 10.x. See <i>ADLS Gen2 Access</i>.</p> <p>Databricks maintains an 18-month support window. For more information, see <a href="https://docs.databricks.com/release-notes/runtime/releases.html">https://docs.databricks.com/release-notes/runtime/releases.html</a>.</p> |
| Databricks 6.x, except for Databricks 6.4 Extended Support | Deprecat | For Databricks 6.x, only Databricks 6.4 Extended Support is still supported. You should upgrade. For more information, see <i>Product Support Matrix</i> .   |
| Databricks 5.5 LTS   | Deprecat | Databricks 5.5 LTS is deprecated. Please upgrade to a supported version. For more information, see <i>Product Support Matrix</i> .   |

## Release 8.11

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item                   | Status | Description                               |
|------------------------|--------|---|
| Enable Onboarding Tour | EOL    | This feature is removed from the product. |

## Release 8.7

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item           | Status     | Description  |
|----------------|------------|--|
| PostgreSQL 9.6 | deprecated | <p>Please upgrade.</p> <div> <b>NOTE:</b> If you are installing on Azure, you must upgrade to PostgreSQL 11. All other deployments must upgrade to PostgreSQL 12. </div> <p>For more information, see <i>Install Databases</i>.</p>  |
| HDP 3.0        | deprecated | <p>Please upgrade to using HDP 3.1.</p> <div> <b>NOTE:</b> In a future release, support for Hortonworks Data Platform (HDP) will be deprecated. Please migrate to using a different supported running environment. For more information, see <i>Product Support Matrix</i>. </div> |
| HDInsight      | deprecated | Support for HDInsight 3.5 and 3.6 has been deprecated. HDInsight is no longer supported. Please migrate to using a supported version of Azure Databricks. See <i>Product Support Matrix</i> .  |

## Release 8.2

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item               | Status      | Description  |
|--------------------|-------------|--|
| CDH 5.16           | deprecated  | Please upgrade to the latest supported version of CDH. See <i>Product Support Matrix</i> . |
| Alation            | end of life | This catalog integration is removed from the product.                                      |
| Waterline          | end of life | This catalog integration is removed from the product.                                      |
| Cloudera Navigator | end of life | This integration is removed from the product.  |

## Release 8.1

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item                  | Status               | Description  |
|-----------------------|----------------------|--|
| PostgreSQL 9.6        | deprecated (partial) | Support for PostgreSQL 9.6 is deprecated for customer-managed Hadoop-based deployments and AWS deployments. PostgreSQL 9.6 is supported only for Azure deployments. When Azure supports PostgreSQL 12 or later, support for PostgreSQL 9.6 will be deprecated in the subsequent release of Trifacta. |
| CDH 6.1               | deprecated           | Please upgrade to the latest supported version of CDH. See <i>Product Support Matrix</i> .   |
| HDP 2.6               | deprecated           | Please upgrade to the latest supported version of HDP. See <i>Product Support Matrix</i> .   |
| Alation integration   | deprecated           | This catalog integration is no longer supported.   |
| Waterline integration | deprecated           | This catalog integration is no longer supported.   |
| Navigator integration | deprecated           | This integration is no longer supported.   |

## Release 7.9

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item              | Status         | Description   |
|-------------------|----------------|---|
| Manage Users area | deprecat<br>ed | All of the functions of the Manage Users area of the Admin Settings page have been moved to the Workspace Users page. See <i>Users Page</i> . |

## Release 7.8

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item            | Status | Description   |
|-----------------|--------|---|
| TDE file format | EOL    | <p>Tableau has reached End of Life for the TDE file format.</p> <p>This file format is no longer supported in the product. All uses of this format have been automatically migrated to use HYPER, which is the replacement format now supported by Tableau.</p> |

## Release 7.7

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item                    | Status         | Description  |
|-------------------------|----------------|--|
| classic Flow View       | EOL            | The original (classic) version of Flow View is no longer available in the product. For more information on the new version, see <i>Flow View Page</i> .  |
| Parameter History Panel | deprecat<br>ed | As a part of collaborative suggestions enhancement, the support for Parameter History panel is deprecated from the software. For more information on collaborative suggestions feature, see <i>Overview of Predictive Transformation</i> . |

## Release 7.6

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item                                    | Status         | Description   |
|---|----------------|---|
| Wrangler Enterprise desktop application | end of<br>life | To access the Trifacta application, please use one of the supported browser versions instead. For more information, see <i>Browser Requirements</i> .   |
| CentOS/RHEL 7.1 and 7.2                 | deprecat<br>ed | Please upgrade to a supported distribution of either operating system. For more information, see <i>System Requirements</i> .   |
| Maintenance Release Updater script      | deprecat<br>ed | <p>This script enabled installation of maintenance updates to the Trifacta deployment. Maintenance updates include:</p> <ul style="list-style-type: none"><li>• Release X.Y.1 to Release X.Y.2 or later</li><li>• Hot Fixes</li></ul> |



## Release 7.1

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item  | Status         | Description  |
|---|----------------|--|
| Wrangler Enterprise desktop application                 | deprecat<br>ed | <p>The Wrangler Enterprise desktop application is no longer built and deployed with the software.</p> <div><b>NOTE:</b> You should switch to using a supported browser version at this time. In a future release, the Wrangler Enterprise desktop application will no longer work. For information on supported browser versions, see <i>Browser Requirements</i>.</div> <p>The Release 6.8 version of the Wrangler Enterprise desktop application is available on request. For more information, please contact <i>Alteryx Support</i>.</p>   |
| API reference document<br>ation with<br>product<br>docs | end of<br>life | <p>The API reference documentation is no longer published with the product documentation.</p> <p>API reference documentation can now be accessed directly from the product. For more information, see <i>API Reference</i>.</p>  |
| CentOS<br>/RHEL 6.x                                     | end of<br>life | Please upgrade to a supported version of CentOS/RHEL.  |
| Azure<br>Databricks<br>5.3 - 5.4                        | end of<br>life | <p>Please upgrade to Azure Databricks 5.5 LTS or Azure Databricks 6.x.</p> <div><b>NOTE:</b> Some versions of Azure Databricks 6.x have already reached End of Life. Customers should upgrade to a version of Azure Databricks that is supported by the vendor. These EOL versions are likely to be deprecated from supported by the Trifacta platform in a subsequent release. For more information on supported versions, see <a href="https://docs.databricks.com/release-notes/runtime/releases.html#supported-databricks-runtime-releases-and-support-schedule">https://docs.databricks.com/release-notes/runtime/releases.html#supported-databricks-runtime-releases-and-support-schedule</a>.</div> |

## Release 7.0

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item                     | Status         | Description  |
|--------------------------|----------------|--|
| Ubuntu 14.04<br>(Trusty) | deprecat<br>ed | Please upgrade the Trifacta node to use Ubuntu 18.04 (Bionic Beaver).  |
| Spark 2.2.x              | deprecat<br>ed | Please upgrade to Spark 2.3 or Spark 2.4, and upgrade your cluster distribution version if necessary. See <i>Configure for Spark</i> . |
| CDH 6.0                  | deprecat<br>ed | Please upgrade to CDH 6.3.   |
| EMR 5.8 - EMR<br>5.12    | deprecat<br>ed | Please upgrade to a supported version of EMR. See <i>Product Support Matrix</i> .  |

## Release 6.10

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item                   | Status         | Description   |
|------------------------|----------------|---|
| v3 APIs                | end of<br>life | The v3 version of the APIs is no longer available. You must migrate to using the v4 APIs.   |
| Send a Copy<br>feature | end of<br>life | The Send a Copy feature is no longer available in the product. As a workaround, you can make a copy of the flow and share it. See <i>Flow View Page</i> . |

## Release 6.8.1

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item         | Status     | Description  |
|--------------|------------|--|
| Cloudera 6.0 | deprecated | Please upgrade to Cloudera 6.3. See <i>Supported Deployment Scenarios for Cloudera</i> . |

## Release 6.8

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item              | Status     | Description   |
|-------------------|------------|---|
| Spark 2.1.x       | deprecated | Please upgrade to a supported version of Spark. See <i>Configure for Spark</i> .                                      |
| EMR 5.6 - EMR 5.7 | deprecated | Please upgrade to EMR 5.8 or later, depending on the version of Spark you are using. See <i>Configure for Spark</i> . |

## Release 6.4

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item                       | Status      | Description   |
|----------------------------|-------------|---|
| v3 APIs                    | deprecated  | The v3 version of the APIs is no longer supported. Please migrate to using the v4 APIs.   |
| PNaCl                      | deprecated  | The PNaCl browser client extensions has been deprecated. Please verify that all users of Trifacta are using a supported version of Google Chrome, in which the WebAssembly client extension is automatically enabled. |
| Command Line Interface     | end of life | The CLI is no longer available. Please migrate to using the v4 APIs.  |
| Java 7                     | deprecated  | Support for Java 7 in the platform has been deprecated. Please upgrade the Trifacta node and any connected cluster to use Java 8.   |
| Chat with us               | end of life | The Chat with Us feature has been removed from the product. The underlying integration did not work.<br>If you need to reach out, please <i>Alteryx Support</i> .   |
| Cloudera 5.14.x and 5.15.x | deprecated  | Please upgrade to Cloudera 6.2.x.   |
| HDP 2.5.x                  | deprecated  | Please upgrade to HDP 3.1.x.  |

## Release 6.0.2

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

None.

## Release 6.0.1

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

## Release 6.0

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item     | Status     | Description                |
|----------|------------|----------------------------|
| CDH 5.13 | deprecated | Please upgrade to CDH 6.0. |
| CDH 5.12 | deprecated | Please upgrade to CDH 6.0. |

## Release 5.1

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item                    | Status     | Description   |
|-------------------------|------------|---|
| Diagnostics Page        | EOL        | The diagnostics server and related Diagnostics Page have been removed from the product. This feature has been superseded by Tricheck, which is available to administrators through the application. |
| PostgreSQL 9.3          | deprecated | PostgreSQL 9.3 is scheduled for deprecation in September 2018.<br>Please upgrade to PostgreSQL 9.6, which is the only supported version for this release.   |
| CDH 5.12                | deprecated | Please upgrade to CDH 5.15.   |
| wasb : storage protocol | deprecated | All interactions with WASB now occur over wasbs : secure storage protocol.  |

## Release 5.0

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item      | Status      | Description   |
|-----------|-------------|---|
| aggregate | end of life | The aggregate transform has been removed from the platform. Its functionality has been replaced by configuration in the pivot transform. See <i>Pivot Transform</i> . |
| CDH 5.11  | deprecated  | Please upgrade to CDH 5.14.   |
| EMR 5.7   | deprecated  | If you are integrating with an EMR cluster, you must upgrade to an EMR 5.11 cluster.  |

## Release 4.2.1

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item     | Status     | Description                 |
|----------|------------|-----------------------------|
| CDH 5.10 | deprecated | Please upgrade to CDH 5.13. |

## Release 4.2

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item             | Status      | Description   |
|------------------|-------------|---|
| wrangled dataset | end of life | Wrangled datasets are no longer available in the platform.<br><br>Instead, you create the actions that applied to wrangled datasets from recipes or two new objects: references and outputs. For more information, see <i>Changes to the Object Model</i> . |
| /docs URL        | end of life | This URL from the platform is no longer available.  |
| s3n protocol     | end of life | To connect to S3 sources, use of <code>s3n</code> is no longer supported in the platform. You must use <code>s3a</code> . See <i>S3 Access</i> .  |

## Release 4.1.1

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item                              | Status         | Description   |
|-----------------------------------|----------------|---|
| CDH 5.9                           | deprecat<br>ed | Please upgrade to CDH 5.12.   |
| Single-file CLI publishing option | deprecat<br>ed | Support for publishing to a single file as part of a CLI <code>run_job</code> action has been deprecated. Please use the external file method of specifying publishing targets. |

## Release 4.1

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item                         | Status         | Description  |
|------------------------------|----------------|--|
| MapR                         | deprecat<br>ed | Support for integration with MapR Hadoop clusters has been deprecated.<br><br>Please install a supported version of Cloudera or Hortonworks. See <i>Install Reference</i> .  |
| CentOS 6.2.x<br>CentOS 6.3.x | deprecat<br>ed | Please upgrade to the latest CentOS 6.x release.   |
| CDH 5.8                      | deprecat<br>ed | Please upgrade to CDH 5.11.  |
| HDP 2.4                      | deprecat<br>ed | Please upgrade to HDP 2.6.   |
| Pig running environment      | end of life    | The Hadoop Pig running environment is no longer available.<br><br>For running jobs on Hadoop, you must use the Spark running environment, which requires no additional configuration. See <i>Configure Spark Running Environment</i> . |
| Python UDFs                  | end of life    | You cannot use Python UDFs in the platform, due to the removal of the Hadoop Pig running environment.<br><br>All UDFs must be migrated to or authored in Java. See <i>User-Defined Functions</i> .                                     |
| Transform Editor             | end of life    | The Transform Editor for entering raw text Wrangle steps has been removed. Please use the Transform Builder for creating transformation steps.   |
| Standardize page             | end of life    | Feature flag-only feature will be replaced by a better standardization capability in a future release.   |
| standardize transform        | end of life    | Feature flag-only feature will be replaced by a better standardization capability in a future release.   |

## Release 4.0.1

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item    | Status     | Description  |
|---------|------------|--|
| MapR    | deprecated | Support for integration with MapR Hadoop clusters has been deprecated.<br><br>Please install a supported version of Cloudera or Hortonworks. |
| CDH 5.7 | deprecated | Please upgrade to CDH 5.10.  |

## Release 4.0

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item   | Status         | Description   |
|--|----------------|---|
| MapR   | deprecat<br>ed | Support for integration with MapR Hadoop clusters has been deprecated.<br><br>Please install a supported version of Cloudera or Hortonworks.  |
| Hadoop Pig<br>running<br>environment                 | deprecat<br>ed | This running environment has been superseded by the Scala-based running environment, which utilizes Spark's in-memory features for faster processing. The Scala Spark version is the default running environment for Hadoop environments. See <i>Running Environment Options</i> .  |
| Python UDF   | deprecat<br>ed | Python UDFs are only supported with Hadoop Pig running environment.   |
| Javascript<br>running<br>environment<br>and profiler | end of<br>life | The original Javascript running environment and profiling engine have been superseded by the Photon running environment, which was introduced in Release 3.2.<br><br>The Photon running environment is enabled by default for front-end processing. See <i>Running Environment Options</i> .  |
| Hadoop Pig<br>profiler                               | end of<br>life | This profiling engine has been superseded by the Scala-based Spark profiler, which utilizes Spark's in-memory features for faster processing. The Scala version of Spark Profiler is the default profiling engine for Hadoop environments. See <i>Profiling Options</i> .   |
| Python-based<br>Spark profiler                       | end of<br>life | In Release 2.7, a Python-based Spark profiler was released. This version of the Spark profiler required that Spark had to be installed on each node of the cluster, reducing efficiency.<br><br>This profiling engine has been superseded by the Scala-based Spark profiler, which utilizes Spark's in-memory features for faster processing. The Scala version of Spark Profiler is the default profiling engine for Hadoop environments. See <i>Profiling Options</i> . |
| /docs URL  | deprecat<br>ed | This URL from the platform points to an in-app page of documentation. All content in this location has been replaced and superseded by product documentation content.   |
| HDP 2.3.2  | deprecat<br>ed | Please upgrade to HDP 2.4 or 2.5.   |

## Release 3.2.1

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item            | Status     | Description                |
|-----------------|------------|----------------------------|
| CDH 5.5/CDH 5.6 | deprecated | Please upgrade to CDH 5.8. |

## Release 3.2

**Release End of Support date:** <https://www.trifacta.com/supportpolicy/>

| Item                           | Status         | Description  |
|--------------------------------|----------------|--|
| Aggregate Tool                 | end of life    | The separate page for building aggregations has been replaced by selecting your aggregation parameters through the Transform Builder. See <i>Transform Builder</i> .   |
| Add Sample Rows to Transformer | end of life    | In the Job Results page, it was possible to add the displayed sample rows back to the Transformer page as a new sample. With the new object model, this is no longer possible.   |
| MapReduce settings             | end of life    | Removed from Admin Settings page, as they are no longer applicable. Support for MapReduce 1 was previously deprecated. See <i>Admin Settings Page</i> .  |
| multisplit transform           | end of life    | This transform has been replaced by a more flexible version of the <code>split</code> transform. See <i>Split Transform</i> .  |
| Batch Server service           | end of life    | Internal platform service has been replaced by Batch Job Runner service.   |
| Monitor service                | end of life    | Internal platform service has been replaced by Batch Job Runner service.   |
| Zookeeper                      | end of life    | The Trifacta platform no longer utilizes the Zookeeper service for managing Hadoop Pig jobs. It has been replaced by Batch Job Runner service. <div><b>NOTE:</b> You may notice references to Zookeeper in configuration blocks or in the interface. These references will be removed in a future release.</div> |
| Support for CDH 5.3/5.4        | deprecat<br>ed | Please upgrade to HDP 5.8.   |

## Release 3.1.2

None.

## Release 3.1.1

| Item                | Description                           |
|---------------------|---------------------------------------|
| Support for HDP 2.2 | Please upgrade to HDP 2.3 or HDP 2.4. |

## Release 3.1

| Item                  | Description                                 |
|-----------------------|---|
| Support for HDP 2.1.1 | Please upgrade to HDP 2.3.                  |
| Support for CDH 5.2   | Please upgrade to CDH 5.3, 5.4, 5.5 or 5.6. |
| Support for MapR 4.1  | Please upgrade to MapR 5.1.                 |

Replacements:

**NOTE:** These features may be deprecated in the future in favor of their replacements.

- New `if` function is designed to replace the ternary construct.
  - See *IF Function*.
  - See *Ternary Operators*.
- New Spark Profiler on Scala is the default Spark profiler.
  - The old version, which required Spark, is still available for specific scenarios; we recommend using the new Spark Profiler on Scala.
  - For more information, see *Overview of Visual Profiling*.

## Release 3.0

None.

## Release 2.7

| Item                    | Description  |
|-------------------------|--|
| multiextract transform  | Transform was removed in favor of new multisplit transform.                        |
| Support for MapReduce 1 | Map Reduce 1 has been superseded by Map Reduce 2, also known as YARN.              |
| WebHCat for publishing  | For publishing to HCatalog, WebHCat has been replaced by publication through Hive. |

For more information, please see *Alteryx Support*.



Copyright © 2022 - Trifacta, Inc.  
All rights reserved.