



Release Notes

Version 4.2.2
April 2017

Contents

What's new in 4.2?	5
International interactive maps	5
German and Japanese localization	6
Nested Formulas	7
Formula support for chasm trap schemas	7
New keywords - Geospatial Search	8
Case configuration	9
Freeze column capability for data modeling page	9
Data Profile	9
Expanded Joins for Aggregated Worksheet	10
Inclusion Rule (Left Outer Join Inner Join Full Outer Join)	10
Non co-sharded joins	10
TQL Dependency Checker	11
Dictionary based query optimizations	11
Automatic periodic snapshot	11
Automatic periodic backup	12
Upgrade management	12
Support for large number of users on an appliance	12
ThoughtSpot Data Connect improvements	12
Data Connect data source re-configuration	12
Data Connect script-based processing	13
Other UX improvements	14
Enhanced color scheme for charts	14
Enhanced configurator	14
Pareto chart enhancement	14
Geo map data labels	14
Color scale on geo maps	14
Background filled world map	15
Faster loading of bubble and heatmap	15
Unsupported zip code area map	15
What's fixed in 4.2.2?	15
FIXED: Horizontal scrolling in a pinboard not working well	15
FIXED: Download pinboard not working	15
FIXED: Overlapping labels in pivot table on IE11	15
FIXED: "Show underlying data" doesn't show anything	15

FIXED: Cannot see users with access to pinboards	16
FIXED: Error shown instead of placeholder title for formula	16
FIXED: Context menu to drill down from a pinboard is not available	16
FIXED: Cannot disable NPS survey question	16
FIXED: Fiscal year-to-date formulas not working	16
FIXED: Filtering on null values causes no data to display	16
FIXED: Remove external join paths from worksheet scope	16
FIXED: Could not edit users in pinboard sharing	17
FIXED: Using IS NULL when defining a generic relationship not working	17
FIXED: Aggregated worksheet data values not showing in search suggestions	17
FIXED: Need ability to set Attribution Dimension on a column in a worksheet	17
FIXED: Non cosharded join caused system to hang	17
FIXED: Error when uploading modeling file	17
FIXED: Scheduled pinboards are not being sent out	18
FIXED: Add the ability to disable/enable localizations	18
FIXED: Lock contention during long running queries	18
FIXED: Add the ability to protect a snapshot from deletion	18
FIXED: Corrected pinboard impressions and ad hoc search metrics	18
FIXED: System tables getting indexed for search incorrectly	18
FIXED: Active Directory sync created duplicate users	19
FIXED: SAML auto redirect causes page not found error	19
FIXED: Security: Strong SSL security in nginx	19
FIXED: Data Connect: Provide a way to specify proxy details when enabling Data Connect	19
FIXED: Data Connect: Not able to edit data sources when target tables include generic relationships	19
FIXED: Data Connect: Informatica agent infraagent does not come up upon restart	20
FIXED: Data Connect: Loads in progress could not be restarted or aborted	20
FIXED: Data Connect: Tasks failing due to timeout	20
FIXED: Data Connect: Notification of failed job didn't go out if Informatica job was deleted	20
FIXED: Extended Enterprise Edition: Add custom branding to the browser tab favicon and text	20
FIXED: Extended Enterprise Edition: Remove Copy Link function from embedded visualization	21
What's fixed in 4.2.1?	21
Record client's identifier in statsDB	21
Login screen can be redirected to SAML location	21
tscli to increase number of points on a geo chart	21
FIXED: ThoughtSpot doesn't work on Firefox and Safari browsers	22
FIXED: Error - "object could not be opened due to errors in some of its dependencies"	22

FIXED: Periodic backup failure due to nodes not mounting NAS	22
FIXED: Schema Viewer doesn't work with localization	22
FIXED: SAML auto redirect logout returns empty page	22
FIXED: /session/user/createusers doesn't create users in bulk	22
FIXED: Data Connect - Cannot completely remove data sources	23
FIXED: Data Connect - Editing a data source changes metadata that breaks underlying objects	23
What's fixed in 4.2?	23
FIXED: Aggregation on top of an aggregated formula error	23
FIXED: Can't add new filter if there's a blank filter	23
FIXED: Pinboard slideshow resets after resetting drill down	23
FIXED: Case sensitivity is not preserved when importing URLs	23
FIXED: Group_function ratios are not summing to 100%	24
FIXED: Share dialog causes line wrapping for users or groups with very long names	24
FIXED: Unable to sync users/groups from AD using LDAP sync script	24
FIXED: Various data label issues with column and line charts	24
FIXED: Show underlying data broken for no_bucket/ detailed timestamp field	24
FIXED: Rename Column in TQL makes it unalterable	24
FIXED: Firefox/IE: Pinboard - Scrollbars inside the pinned visualizations are unresponsive to dragging	25
FIXED: Unable to determine which table is which	25
FIXED: Unable to open downloaded Excel	25
FIXED: Sage suggestions and formula definition are not rendered on IE11	25
FIXED: Download as CSV throws exception	25
FIXED: Row data still being shown in filters on worksheet with chasm trap involved	26
FIXED: Remove deprecated pinboards/worksheets	26
FIXED: Chart legend labels are cut off	26
FIXED: CSV/Excel downloads now indicate when results are incomplete	26
FIXED: SAML login link not appearing	26
FIXED: Asynchronous rendering of axis and data labels	26
FIXED: Format pattern disabled for non-additive attributes	26
FIXED: Filtering on day_of_week function fails	27
FIXED: Unable to modify fields via model file	27
FIXED: Incorrect counting of nodes	27
FIXED: Synonyms are not suggested in search bar	27
FIXED: Cluster doesn't start automatically after one node cluster booted	27
FIXED: Pinboard PDF download sometimes contains individual visualizations that have not finished rendering	27

FIXED: Data Connect - Tables imported from data connect don't have column level security option	28
FIXED: Data Connect - Getting error when setting up Data Connect	28
FIXED: Data Connect - Suggestions for column values are not appearing in the search bar	28
FIXED: Data Connect - Deleting a user removes jobs created by the user	28
FIXED: Data Connect - Falcon column changes don't propagate to Callosum	28
FIXED: Data Connect - (\$LastRunTime) - Set to GMT only	28
FIXED: Data Connect - Data load failed when initializing writer due to SSL cert	29
FIXED: Data Connect - data load failing for date type column with JDBC	29
Known issues	29
Applying filter on {blank} value results in no data	29
Proper error message for tsload errors	29
tscli monitoring command should check validity of email address	29
KPI line extends outside of the chart canvas	29
Resizing a column automatically sorts it	30
Show underlying data is broken on aggregated worksheets	30
Explanation for non co-sharded join prevention	30
Exclude is broken on measures	30
Single disk failures causes cluster update failures	30
Downloaded charts are poor quality on IE11	30
Supported upgrade paths	30

What's new in 4.2?

International interactive maps

ThoughtSpot now supports international county-level maps for a select number of countries. Here are the five newly added countries, along with their administrative regions:

- United Kingdom
 - County and Unitary Authority
 - Local Authority District
- Germany
 - State
 - District
- Sweden
 - County
 - Municipality
- South Africa
 - State
 - District
- France
 - Region
 - Department

You can use the **Specify Geographic Configuration** dialog while modeling your data, to select the specific GeoType setting.

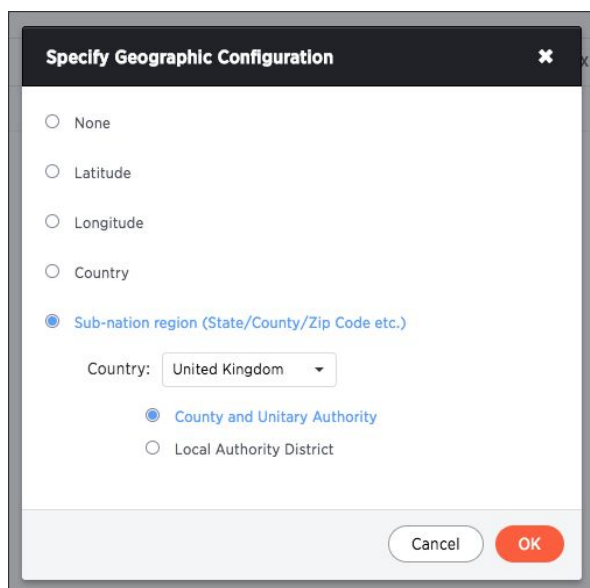


Figure 1: Specify Geographic Configuration for international sub-nation regions

The names of international country regions are supported in English, as well as in their respective foreign languages.

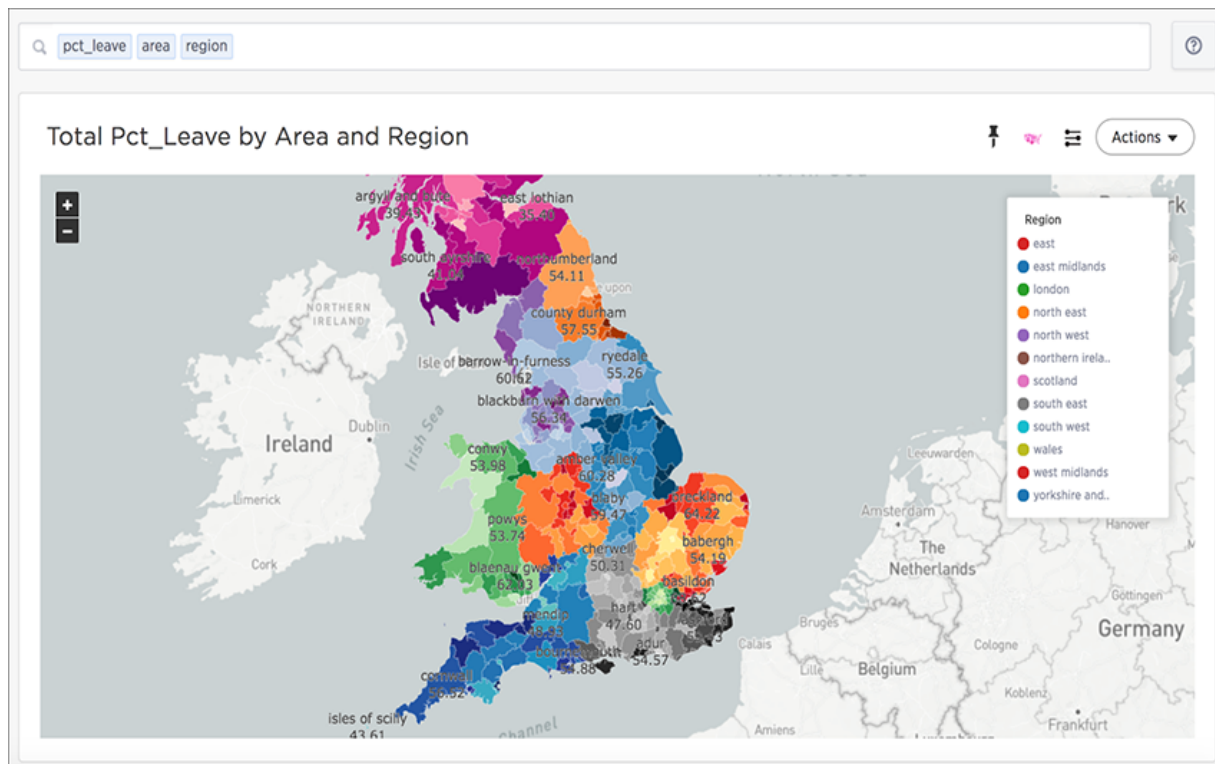


Figure 2: Example United Kingdom geo area map

German and Japanese localization

The ThoughtSpot user interface is now available in German and Japanese. To enable localizations, please contact ThoughtSpot Support. The language used to display menu items in ThoughtSpot is based on a setting in the user profile.



Figure 3: ThoughtSpot in Japanese

Keywords, operators, and error messages are included in the translated material. However, formulas and documentation are not translated. Also, all metadata will retain their user inputted format.

Nested Formulas

Nested formulas, or formula on formula, allow you to reference a formula within another formula.



Figure 4: formula1 nested in formula2

This feature enables encapsulation/decomposition of formula logic and minimizes duplication by allowing a formula to reuse another formula. These formulas become columns, and are even suggested within the Formula Builder. Invalid formulas, such as cyclic or inconsistent aggregation formulas, are prevented from being created.

In addition, these formulas have dependency awareness, which ensure that you do not break any dependent formulas while editing a formula with nested formulas. The system does not let you delete a formula that is nested in other formulas.

There is no limit to the number of nestings you can create.

Formula support for chasm trap schemas

Previously, ThoughtSpot didn't allow you to create a formula that spanned across a chasm trap in a worksheet. However, now you can create a formula that involves aggregated measures coming from multiple fact tables of a chasm trap.

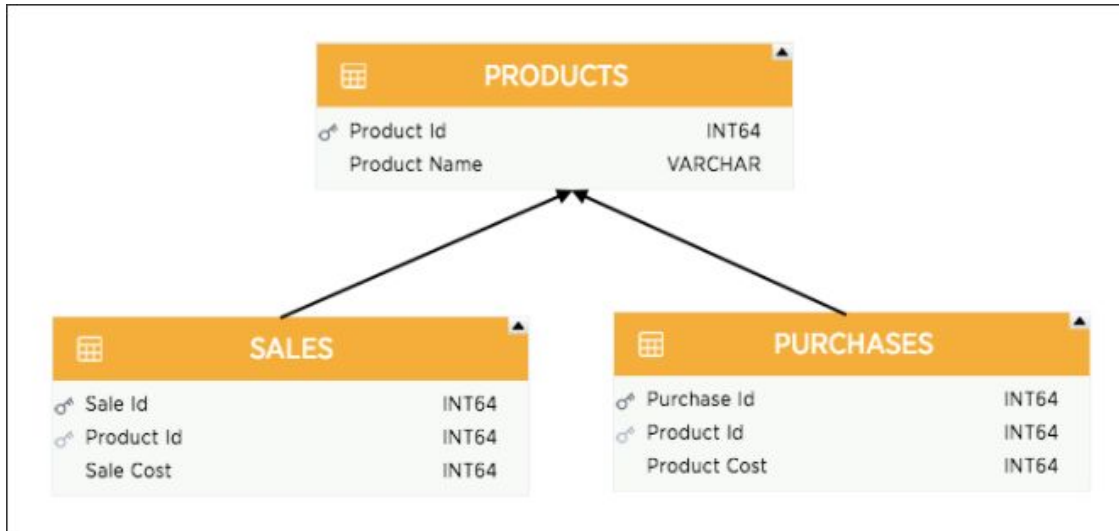


Figure 5: Formula: $\text{sum}(\text{Product Cost}) / \text{sum}(\text{Sales Cost})$

The Formula Builder now provides column suggestions across chasm traps.

New keywords - Geospatial Search

The new keywords **near**, **within**, and **farther than/from** allow you to search with spatial filters on data contained in a table or worksheet which contains geo type columns with latitude/longitude data. This lets you search based on spherical distance. The distance is calculated as a straight-line distance (not road distance) radius using miles, kilometers, or meters.

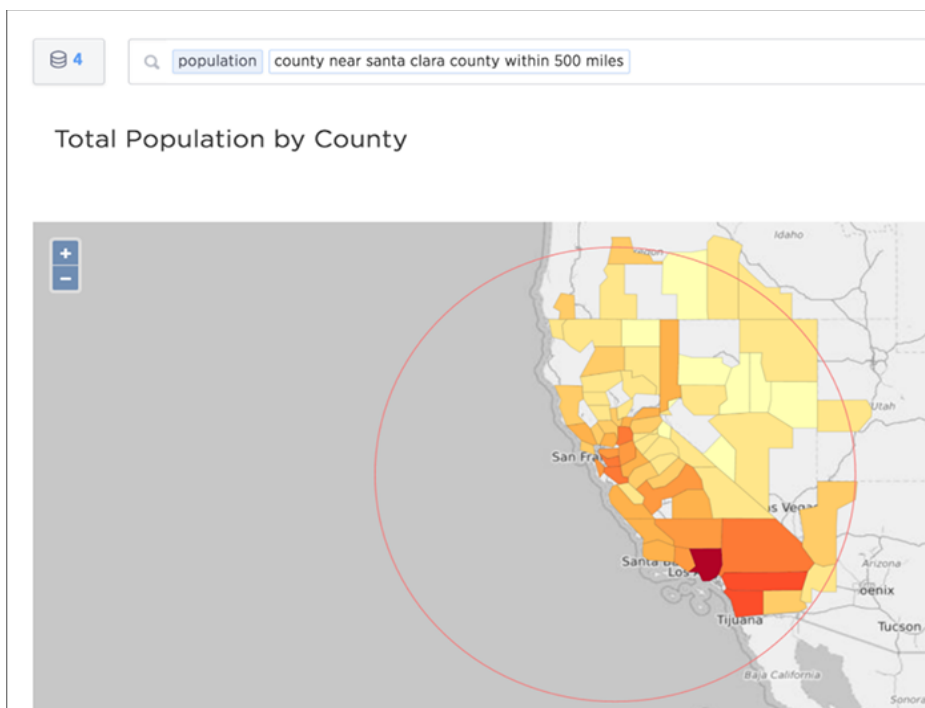


Figure 6: Geospatial search on a geo area map

The keywords follow the following formats:

- **near** (with default of 10 km)
- **within** X miles/kilometers/meters
- **farther than** X miles/kilometers/meters **from**

For example, your search might look like this:

revenue store name county **near** alameda **within 50 miles**

Or it might look like this:

Average hours worked branch **farther than 80 km from** scarborough

Case configuration

The case sensitivity for source data strings is now preserved in the display. So, the visual display of results is identical to the input case that is loaded.

It is important to note that string casings aren't applied globally, but by column. So datasets will have different string casings as long as they're in different columns. Tables that are already compacted will keep their lowercase format. In these cases, in order to get the specific string case that you want, you would have to truncate related tables and reload them.

To take advantage of case configuration, you need to have ThoughtSpot Support enable it on your cluster. In addition, title casing should be disabled for string casing to properly work.

Freeze column capability for data modeling page

As you edit the column properties of a table or worksheet, the Column Name column now freezes in place. This should make it much easier for you to know which column you're making changes for.

Data Profile

Basic column data profile information is available under **Profile** on the Data page.

	ORDER TOTAL PRICE	DISCOUNT	REVENUE	SUPPLY COST	TAX	COMMIT DATE
NULL VALUES	0	0	0	0	0	0
MIN	0	0	0	0	0	02/10/1992
MAX	43,226,660	10	10,304,950	125,159	8	10/24/1998
AVERAGE	19,100,242.55	5.05	3,608,079.64	89,749.2	4.04	N/A
SUM	95,558,513,471	25,254	18,051,222,4..	449,015,252	20,228	N/A

Figure 7: Column profile view

When data is first loaded into ThoughtSpot, it is often unfamiliar to users. This profile view should help you get a better sense of what’s there before asking questions on the data.

Expanded Joins for Aggregated Worksheet

You can now join aggregated worksheets to any base table or user imported data. In addition, aggregated worksheets can be included as tables in worksheets. This means aggregated worksheets now behave similar to tables, and can be used in the same way as a table, excluding any TQL manipulation.

Inclusion Rule (Left Outer Join | Inner Join | Full Outer Join)

ThoughtSpot now supports full outer and right outer joins during worksheet creation, regardless of the sharding of underlying tables.

Choose the inclusion rule:

- Apply full outer join (Full outer join)
- Apply left outer join (Left outer join)
- Apply right outer join (Right outer join)
- Exclude empty rows (Inner join)

Figure 8: Inclusion rule choices

Non co-sharded joins

ThoughtSpot also supports joins between two sharded fact tables with different shard keys, otherwise known as non co-sharded tables. However, because data is being redistributed when

this type of join happens, search results will come back more slowly than with co-sharded tables. Therefore, we recommend that you use a common shard key for two fact tables.

You are not limited by the column connection or relationship type.

TQL Dependency Checker

TQL prevents disruptive changes by checking for dependencies when schema changes are made. In the case of a potential breakage, you are shown the dependent objects.

```
TQL [database=TPCH5K]> drop table SUPPLIER;
Logical tables:
  AggrWsWithPinnedMeasure
  SUPPLIER
  Formula Worksheet
  Formula Worksheet With Full Outer join
  Container Worksheet

Dependent question answer book:
  Brand Revenue Trend
  Average Revenue by Part
  Brand Revenue
  Region Revenue by Part

(4) logical tables, (0) pinboard answer books, (4) question answer books, (0) logical columns will be broken.
```

Figure 9: Dependency awareness in TQL

Dictionary based query optimizations

Dictionary compression, which replaces the original data with a map string, is already implemented in your ThoughtSpot application. This reduces memory usage, as the system only has to look at the column value associated with the original data, then look that up in the compressed dictionary to see what it equals to.

Now, we are leveraging dictionary compression to speed up queries, especially those involving strings. Query executions should only take place on the compressed dictionary domain, where only a small number of values have to be decompressed before returning an answer. This eliminates decompression overhead, and should improve queries that involve strings.

Automatic periodic snapshot

There is a new implementation of snapshots—manual or periodic. So you are shown two separate lists when you run the `tscli` command.

You can have up to 20 manual snapshots, after which, you have to clear one before you are able to create another. All existing snapshots from the previous version of ThoughtSpot will become manual snapshots when you upgrade.

Automatic snapshots follow time-based policies, e.g.:

- 3 * 1-hour snapshots
- 2 * 4-hour snapshots

These policies are configurable.

Automatic periodic backup

Time-based automatic backups are available as well, with an option of dataless, lightweight, or full backups. All modes of backups have local or NAS options.

There is also a richer garbage collection policy e.g.:

- 7 * 1-day backups
- 4 * 1-week backups

And like automatic snapshot policies, backup policies are also configurable.

Upgrade management

We have made upgrades more robust and controllable in the following ways:

- Preupdate checks for potential failures
- Ability to check the status of the update process
- Resumable upgrades from where it paused
- Bypass failures to continue upgrades

Support for large number of users on an appliance

ThoughtSpot can support up to 600 simultaneous users, and provide scale for more than 80,000 users in the system.

ThoughtSpot Data Connect improvements

There are a number of improvements around the self-serve data integration UI offered through ThoughtSpot.

Data Connect data source re-configuration

You can add or remove columns for existing tables. This means you can edit the data source definition to add or remove columns, and adapt existing jobs to change source schemas.

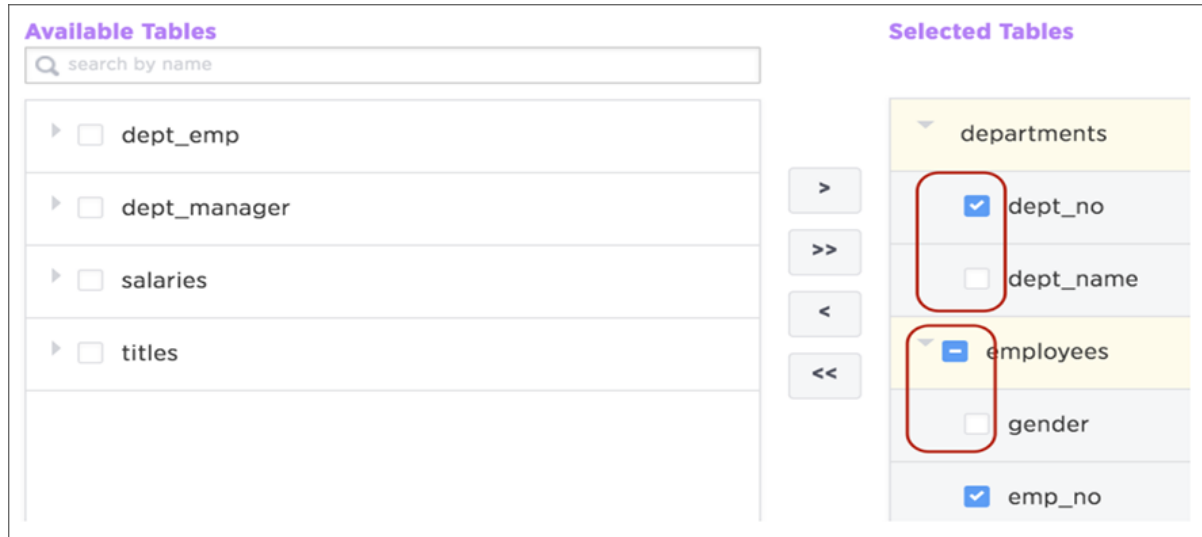


Figure 10: Re-configuring the data source

All relationships are captured with this feature. The relationships are automatically created for you when you add a primary key column.

It is recommended that you do not edit the schema when altering tables.

Data Connect script-based processing

You now have to option to process the data before and after any load by adding a pre- or post-script. A script is a set of TQL statements, and you can add a script of any TQL command. The main use case is to truncate existing tables.

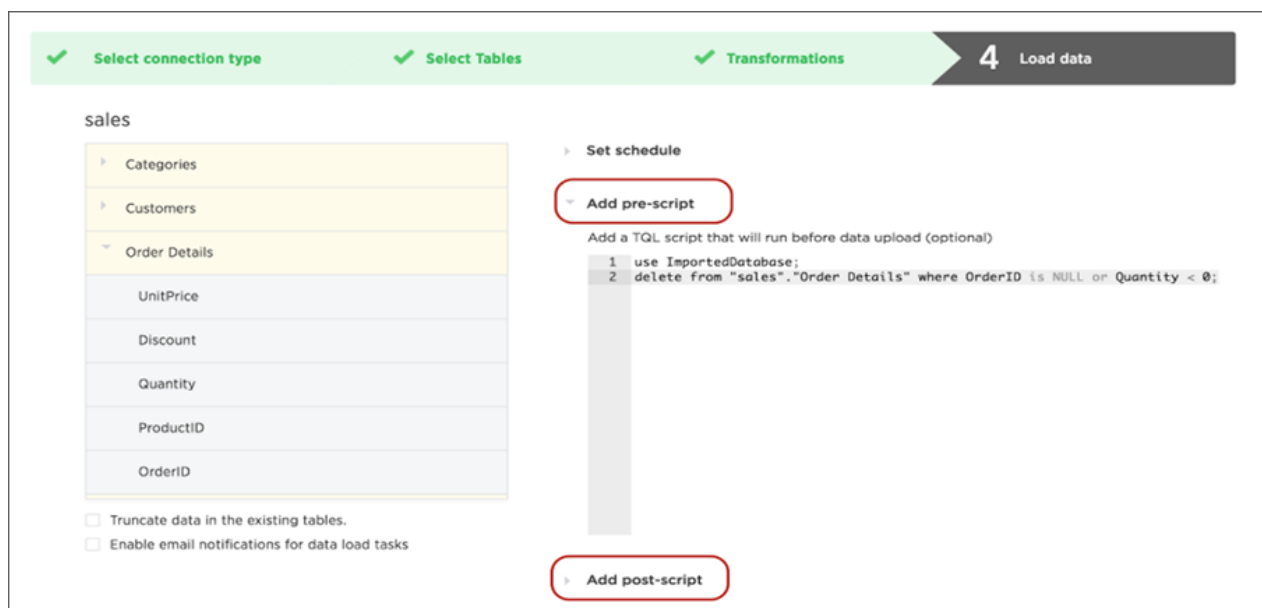


Figure 11: Options to add pre- or post-scripts

Other UX improvements

Other UX improvements include chart and geo map improvements.

Enhanced color scheme for charts

Colors are maintained across searches within a session. For example, when doing a search on revenue by state, each state will keep its color assigned to it even if you change the search or chart type.

Enhanced configurator

The chart configuration panel has been combined and moved to the right drawer on the answer page. This new configurator includes the option to manually set the y-axis scale.




Figure 12: The enhanced chart configurator

Pareto chart enhancement

The line portion of the pareto chart has been corrected to show data labels as percentages.

Geo map data labels

You can now see data labels on geo area and geo bubble maps by checking the **Show Data Labels** setting in the chart configurator.

Color scale on geo maps

You can now see color scale on the bottom right corner of the map for the geo area chart. The color scale will only show up when there is no legend column, or the legend column has only one value selected.

Background filled world map

When geo tiles are not enabled, we now show a filled world map in the background to give you a sense of the current location and zoom level on the map. The maps will show water as white and land as light gray.

Faster loading of bubble and heatmap

We now fetch only points, instead of the whole polygon boundaries, to draw bubble and heat maps, which makes them load faster.

Unsupported zip code area map

Geo area maps are no longer supported for zip codes. This is because zip codes are not comprised of areas, and it is therefore impossible to get the boundary data for zip codes. Earlier we supported it using a data file that only had boundaries for half of the zip codes.

What's fixed in 4.2.2?

FIXED: Horizontal scrolling in a pinboard not working well

ISSUE ID: SCAL-17688

Description: Trying to scroll horizontally (towards the right) on a pivot table within a pinboard, caused the columns to scroll as expected, but the very first column on the left started moving towards the right as well, taking the whole window along with it. This has been fixed.

FIXED: Download pinboard not working

ISSUE ID: SCAL-19433

Description: Download pinboard wasn't working in version 3.5.7. This has been fixed.

FIXED: Overlapping labels in pivot table on IE11

ISSUE ID: SCAL-19363

Description: In the Internet Explorer 11 browser the label for a pivot table column had overlapping letters. This has been fixed.

FIXED: "Show underlying data" doesn't show anything

ISSUE ID: SCAL-19381

Description: When using a Safari browser, choosing "Show underlying data" from a pinboard showed a message saying that some rows are being shown, when in fact no rows are shown. This has been fixed.

FIXED: Cannot see users with access to pinboards

ISSUE ID: SCAL-18751

Description: When editing multiple pinboards in bulk, to share with users and groups, everything works as expected. But afterwards, when editing permissions on the pinboards, only the groups with access, and not the users, are shown. However, the users still have permission to access the pinboards. This has been fixed, so that you now may view and edit the users who have access.

FIXED: Error shown instead of placeholder title for formula

ISSUE ID: SCAL-17764

Description: On the Internet Explorer 11 browser, trying to edit a formula showed an error message where the formula title should be. This has been fixed.

FIXED: Context menu to drill down from a pinboard is not available

ISSUE ID: SCAL-19251

Description: A user with **View** access to a pinboard and with access to the underlying data was unable to access the drill down menu on the pinboard when viewing, editing, or presenting. This has been fixed.

FIXED: Cannot disable NPS survey question

ISSUE ID: SCAL-19172

Description: After using the flag to disable the Net Promoter Score (NPS) question at the bottom of the browser window, the NPS question was still showing. This has been fixed.

FIXED: Fiscal year-to-date formulas not working

ISSUE ID: SCAL-18185

Description: Formulas that contained fiscal year-to-date calculations were not working previously. This has been fixed.

FIXED: Filtering on null values causes no data to display

ISSUE ID: SCAL-19348

Description: In the filter dialog, null values were displaying as {null} instead of {blank}. When selecting to filter on these values, the message “No data to display” was shown. This has been fixed.

FIXED: Remove external join paths from worksheet scope

ISSUE ID: SCAL-13553

Description: When creating a worksheet, you are prompted to select a join path if there are multiple possible join paths. In the past, all possible join paths were shown, including those that

involved tables outside of the worksheet scope. This made it hard to choose an appropriate join path. Now only join paths within the scope of the worksheet will be shown.

FIXED: Could not edit users in pinboard sharing

ISSUE ID: SCAL-18751

Description: When editing a group of shared pinboards, the groups with permissions on the pinboards could be changed, but the users with permissions didn't appear. This has been fixed.

FIXED: Using IS NULL when defining a generic relationship not working

ISSUE ID: SCAL-19089

Description: Using a generic relationship which included the IS NULL operator was causing the application server to fail. This has been fixed.

FIXED: Aggregated worksheet data values not showing in search suggestions

ISSUE ID: SCAL-17825

Description: For aggregated worksheets, data values were not being indexed for search suggestions. Now, if the value is an attribute with aggregation set to "none", the value from the base table will be used to generate the search suggestion for the aggregated worksheet.

FIXED: Need ability to set Attribution Dimension on a column in a worksheet

ISSUE ID: SCAL-18478, SCAL-17009

Description: In the past, you could set the Attribution Dimension column setting on a base table, but not on a worksheet. Some use cases required this setting to be different for different worksheets built on the same base tables. Now you can make this setting at the worksheet level.

FIXED: Non cosharded join caused system to hang

ISSUE ID: SCAL-19118

Description: A worksheet that required joining tables that were not cosharded caused the system to hang. Checks have been put in place to prevent this in the future.

FIXED: Error when uploading modeling file

ISSUE ID: SCAL-17053

Description: After downloading and editing the modeling CSV file, uploading it produced an error. This has been fixed.

FIXED: Scheduled pinboards are not being sent out

ISSUE ID: SCAL-19339

Description: Scheduled pinboards email messages were not being sent out. This was due to a discrepancy between the capitalization in the recipients email addresses. This has been fixed.

FIXED: Add the ability to disable/enable localizations

ISSUE ID: SCAL-18697

Description: You can now enable and disable locale based translations of the ThoughtSpot application. These are disabled by default. If you want to enable or disable localizations, please contact ThoughtSpot Support for assistance.

FIXED: Lock contention during long running queries

ISSUE ID: SCAL-19534

Description: Long running searches were blocking answers from being returned for searches issued after the long running search. This has been fixed, so that those answers which can be returned will be, even if another search is taking a long time.

FIXED: Add the ability to protect a snapshot from deletion

ISSUE ID: SCAL-19056

Description: Added the ability to “pin” a snapshot so that it doesn't get automatically garbage collected. This is useful in cases where you want to save an automatic snapshot longer than it would normally be saved. For details on how to accomplish the pinning or unpinning of a snapshot, contact ThoughtSpot Support.

FIXED: Corrected pinboard impressions and ad hoc search metrics

ISSUE ID: SCAL-14765, SCAL-18432, SCAL-17841, SCAL-17970

Description: The pinboard impressions and ad hoc search metrics were computed in such a way that they undercounted impressions. This has been fixed.

FIXED: System tables getting indexed for search incorrectly

ISSUE ID: SCAL-19154

Description: Some new system tables were being indexed for search by default, when they should have been set to not be indexed. This has been fixed.

FIXED: Active Directory sync created duplicate users

ISSUE ID: SCAL-19122, SCAL-18842

Description: Active Directory sync was creating users in ThoughtSpot based on their memberships in AD groups. But the domain name was being created in lowercase (even when AD had the domain name in uppercase). This was causing duplicate users to be created. This has been fixed.

FIXED: SAML auto redirect causes page not found error

ISSUE ID: SCAL-18951

Description: The SAML auto redirect was causing a HTTP 404 Page Not Found error. This has been fixed.

FIXED: Security: Strong SSL security in nginx

ISSUE ID: SCAL-18386

Description: Implemented the latest security recommendations for nginx, including:

- Only enabling TLS 1.2; disabling TLS 1.1 and 1.0
- Disallowing unsafe ciphers
- Added a setting to enable choosing the server's cipher preference, rather than defaulting to the client's preference

FIXED: Data Connect: Provide a way to specify proxy details when enabling Data Connect

ISSUE ID: SCAL-17610

Description: When enabling the infaagent for Data Connect with a proxy, the setting was not persistent over a software update. You had to manually re-set the proxy for Data Connect. This has been fixed.

FIXED: Data Connect: Not able to edit data sources when target tables include generic relationships

ISSUE ID: SCAL-19392

Description: When the target tables in ThoughtSpot included generic relationships (those created using the syntax "CREATE RELATIONSHIP WITH"), the Data Connect data sources could not be edited. This has been fixed.

FIXED: Data Connect: Informatica agent infraagent does not come up upon restart

ISSUE ID: SCAL-18492

Description: When restarting ThoughtSpot, the infraagent did not restart, even though it was reported as up and running. This has been fixed.

FIXED: Data Connect: Loads in progress could not be restarted or aborted

ISSUE ID: SCAL-19146

Description: Some loading tasks that were in progress didn't have the ability to be stopped or restarted. This was due to the call to the Informatica API login failing or timing out. This has been fixed.

FIXED: Data Connect: Tasks failing due to timeout

ISSUE ID: SCAL-18420

Description: Timeout was causing Data Connect tasks to fail when fetching a table. This has been fixed.

FIXED: Data Connect: Notification of failed job didn't go out if Informatica job was deleted

ISSUE ID: SCAL-19410

Description: After deleting tasks on Informatica, email alerts were not sent out when the tasks didn't start due to these tasks not being found. This has been fixed.

FIXED: Extended Enterprise Edition: Add custom branding to the browser tab favicon and text

ISSUE ID: SCAL-15045

Description: In ThoughtSpot Enterprise Edition, you can now modify the browser tab favicon and text to reflect your company's brand. When you upload a new application logo, it will also be used as the favicon for the browser tab. You can also enter custom text that will appear as the page title in the browser tab. To enter custom text, go to the **Page Title** section at the bottom of the **Style Customization** page and enter the new text there.

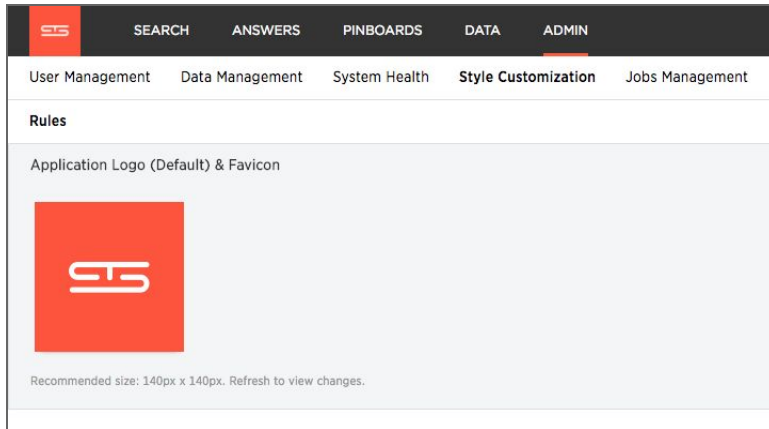


Figure 13: Customizing the favicon for the browser tab

FIXED: Extended Enterprise Edition: Remove **Copy Link** function from embedded visualization

ISSUE ID: SCAL-18852

Description: The **Copy Link** icon should have been hidden from embedded visualizations, because its function is to provide the link to use for embedding. It has been removed.

What's fixed in 4.2.1?

Record client's identifier in statsDB

ISSUE ID: SCAL-17698

Description: You can now find out the user id and session id for API calls. This helps you track stats like who the daily active users are.

Login screen can be redirected to SAML location

ISSUE ID: SCAL-18186

Description: You can have your users skip the ThoughtSpot login screen and be immediately redirected to your SAML location.

tscli to increase number of points on a geo chart

ISSUE ID: SCAL-18509

Description: Before you could only use a URL flag to increase the data points shown on a geo chart. Now there is a tscli command to increase the count to over 5,000. Contact ThoughtSpot Support if you would like to adjust this limit.

FIXED: ThoughtSpot doesn't work on Firefox and Safari browsers

ISSUE ID: SCAL-18132, 18405, 18551

Description: The pinboard page didn't load when using ThoughtSpot on Firefox, among other issues. ThoughtSpot now works on Firefox and Safari version 10.

FIXED: Error - "object could not be opened due to errors in some of its dependencies"

ISSUE ID: SCAL-18141

Description: Opening a worksheet with no missing dependencies would return an error. This would also affect row level security.

FIXED: Periodic backup failure due to nodes not mounting NAS

ISSUE ID: SCAL-18634

Description: A periodic backup failure would occur when there were four nodes that have not mounted NAS due to permissions on the NAS server side. In the future, NAS should be mounted correctly for all nodes using the `tscli nas mount` command before implementing periodic backup.

FIXED: Schema Viewer doesn't work with localization

ISSUE ID: SCAL-18750

Description: Clicking **View Schema** in the Data view returned a blank screen when using a localized version of ThoughtSpot. This has been fixed.

FIXED: SAML auto redirect logout returns empty page

ISSUE ID: SCAL-18951

Description: The automatic SAML logout redirect was broken, and would return an empty page.

FIXED: /session/user/createusers doesn't create users in bulk

ISSUE ID: SCAL-18997

Description: The `/session/user/createusers` endpoint has changed and would not create users in bulk. This has been fixed.

FIXED: Data Connect - Cannot completely remove data sources

ISSUE ID: SCAL-18998

Description: Deleting a data source would not completely remove it from the UI. Instead, an error would occur and the data source would persist.

FIXED: Data Connect - Editing a data source changes metadata that breaks underlying objects

ISSUE ID: SCAL-19024

Description: Editing a data source that references an existing table would change the metadata. This would then break underlying objects.

What's fixed in 4.2?

FIXED: Aggregation on top of an aggregated formula error

ISSUE ID: SCAL-12193

Description: An error message was returned when you performed an aggregation, such as sum, on top of an aggregated formula. This has been fixed by preventing you from saving inconsistent aggregations.

FIXED: Can't add new filter if there's a blank filter

ISSUE ID: SCAL-13678

Description: You could not add any additional filters to a pinboard if you already had a blank filter associated with it.

FIXED: Pinboard slideshow resets after resetting drill down

ISSUE ID: SCAL-14324

Description: While presenting a pinboard, drilling down then resetting would take you back to the beginning of the slideshow. Now, the slideshow remains on the visualization being resetted.

FIXED: Case sensitivity is not preserved when importing URLs

ISSUE ID: SCAL-14588

Description: Case sensitivity was not preserved when importing URLs to the data view. This issue is important since some URLs require case sensitivity. This has been fixed by support for mixed case in output.

FIXED: Group_function ratios are not summing to 100%

ISSUE ID: SCAL-15246

Description: Group_function ratios should sum to 100%, but this was not the case due to a rounding issue.

FIXED: Share dialog causes line wrapping for users or groups with very long names

ISSUE ID: SCAL-15541

Description: The share dialog caused line wrapping when users or groups with very long names were added to the share queue. This caused confusion as to which users or groups had been selected.

FIXED: Unable to sync users/groups from AD using LDAP sync script

ISSUE ID: SCAL-15783

Description: The users that were locally created in ThoughtSpot were being synced with AD but were not being dropped into their corresponding AD groups. This has been fixed.

FIXED: Various data label issues with column and line charts

ISSUE ID: SCAL-15942

Description: There were various issues with data labels and the tooltip when negative values or shared y-axis was involved. This has been fixed.

FIXED: Show underlying data broken for no_bucket/detailed timestamp field

ISSUE ID: SCAL-15994

Description: This error only occurred when one of the fields was a timestamp only field, and the bucket for it was set to no_bucket/detailed. When this field was removed, the underlying data worked.

FIXED: Rename Column in TQL makes it unalterable

ISSUE ID: SCAL-16002

Description: When you changed the name of a column on a base table using the TQL shell you could no longer use that column in other alter table statements. This issue has been fixed.

FIXED: Firefox/IE: Pinboard - Scrollbars inside the pinned visualizations are unresponsive to dragging

ISSUE ID: SCAL-16096

Description: The scrollbars on the visualizations in pinboards would not properly respond to dragging. Instead, the visualization would follow the cursor around. This stuck-mouse bug has been fixed.

FIXED: Unable to determine which table is which

ISSUE ID: SCAL-16566

Description: If you had multiple tables with the same or similar name in different databases you might have had a difficult time telling them apart. For example:

- vwBridge_DimSecuredDWUser_DimAccount
- vwBridge_DimSecuredDWUser_DimAuditForm
- vwBridge_DimSecuredDWUser_DimAuditType

This issue has been fixed by allowing the display to show the database/schema along with the tables.

FIXED: Unable to open downloaded Excel

ISSUE ID: SCAL-16854, 17388

Description: Downloaded excel files with long names would not open. This was because the exported filename exceeded the limit set by Microsoft Excel. This has been fixed.

FIXED: Sage suggestions and formula definition are not rendered on IE11

ISSUE ID: SCAL-16931, 16932

Description: While using Internet Explorer 11, editing a token in the middle of a query provided an option to scroll down for other suggestions, which is not the correct behavior. Also, editing a formula in an unsaved answer would render the formula editor with random text and a perpetual loading indicator.

FIXED: Download as CSV throws exception

ISSUE ID: SCAL-16936

Description: This issue was related to how answers for visualizations were generated when a pinboard filter was applied. This has been fixed.

FIXED: Row data still being shown in filters on worksheet with chasm trap involved

ISSUE ID: SCAL-17006

Description: Data that shouldn't be seen was shown in filters on worksheets with join paths on tables involved with chasm traps. This regression bug has been fixed.

FIXED: Remove deprecated pinboards/worksheets

ISSUE ID: SCAL-17063

Description: We have removed deprecated pinboards and worksheets that are no longer being used on the Resource Usage page.

FIXED: Chart legend labels are cut off

ISSUE ID: SCAL-17098

Description: Chart legend labels were cut off even when there was plenty of room.

FIXED: CSV/Excel downloads now indicate when results are incomplete

ISSUE ID: SCAL-17126

Description: CSV and Excel files now indicate at the top of the page the number of rows that have been successfully downloaded.

FIXED: SAML login link not appearing

ISSUE ID: SCAL-17154

Description: The SAML login link would not appear on the login page. This blocked SAML use and portal integration. The temporary workaround was to use the following link:

[http\(s\)://<hostname>/callosum/v1/saml/login](http(s)://<hostname>/callosum/v1/saml/login)

FIXED: Asynchronous rendering of axis and data labels

ISSUE ID: SCAL-17287

Description: There were significant delays in load time for pinboards due to synchronous rendering of all labels. Deciding to load labels asynchronously has reduced the load time.

FIXED: Format pattern disabled for non-additive attributes

ISSUE ID: SCAL-17440

Description: Format pattern option in data modelling was disabled for non-additive numeric attributes. This was an issue with the UI.

FIXED: Filtering on day_of_week function fails

ISSUE ID: SCAL-17468

Description: Filtering on day_of_week function from the search bar failed and returned a red bar error.

FIXED: Unable to modify fields via model file

ISSUE ID: SCAL-17566

Description:

Problem: Field descriptions for tables and worksheets were broken. This was due to the metadata cache not being cleared upon the updates.

FIXED: Incorrect counting of nodes

ISSUE ID: SCAL-17749

Description: This was due to a bug in `hadoop_tool.py`, whereby a new node would get counted when we tried to count the number of old nodes. This only happened when we tried to re-add an old node.

FIXED: Synonyms are not suggested in search bar

ISSUE ID: SCAL-17968

Description: Synonyms were not being suggested in the search bar. This was because data value suggestions were over prioritized over synonym suggestions.

FIXED: Cluster doesn't start automatically after one node cluster booted

ISSUE ID: SCAL-17976, 17998

Description: It seemed that one node was waiting on some kind of process when the process had not started and it was not listening on a specific port. The workaround was to manually stop and start the cluster.

FIXED: Pinboard PDF download sometimes contains individual visualizations that have not finished rendering

ISSUE ID: SCAL-18006

Description: Some visualizations would not be completed rendered in a PDF download.

FIXED: Data Connect - Tables imported from data connect don't have column level security option

ISSUE ID: SCAL-15871

Description: Data Connect imported tables need to be just like regular tables and support column level security. Row level security can now be applied to tables imported via Data Connect, and the RLS rules are correctly honored.

FIXED: Data Connect - Getting error when setting up Data Connect

ISSUE ID: SCAL-16955

Description: We would get a "No Json could be decoded" error when setting up Data Connect. This issue came about because the infa cloud IP addresses were not whitelisted.

FIXED: Data Connect - Suggestions for column values are not appearing in the search bar

ISSUE ID: SCAL-17382

Description: Suggestions for column values were not appearing in the search bar. This was because special characters prevented filtered results.

FIXED: Data Connect - Deleting a user removes jobs created by the user

ISSUE ID: SCAL-17445

Description: Deleting a user deleted all Data Connect jobs created by that user. Now, when a user is deleted, the unclaimed jobs are reassigned to tsadmin.

FIXED: Data Connect - Falcon column changes don't propagate to Callosum

ISSUE ID: SCAL-17706

Description: You can now rename columns and tables after a data load, using the DDL edit step or the ThoughtSpot web application.

FIXED: Data Connect - (\$LastRunTime) - Set to GMT only

ISSUE ID: SCAL-17726

Description: The \$LastRunTime was permanently set to GMT timezone and couldn't be reconfigured. You are now able to add an expression on top of \$LastRunTime/\$LastRunDate for a filter condition.

FIXED: Data Connect - Data load failed when initializing writer due to SSL cert

ISSUE ID: SCAL-17981

Description: An error would appear stating that the data load failed when initializing writer due to SSL cert. This has been fixed.

FIXED: Data Connect - data load failing for date type column with JDBC

ISSUE ID: SCAL-18001

Description: With Data Connect, data load for a table was failing for date type column with JDBC connector. This error occurred while writing data to the target native system.

Known issues

Applying filter on {blank} value results in no data

ISSUE ID: SCAL-7731

Description: Applying a filter on {blank} values results in an answer with no data.

Proper error message for tsload errors

ISSUE ID: SCAL-13149

Description: Error messages when loading data using tsload are incomplete.

tscli monitoring command should check validity of email address

ISSUE ID: SCAL-14471

Description: `tscli monitoring set-config --email` does not check the validity of email addresses. Instead, the command takes whatever value was entered.

KPI line extends outside of the chart canvas

ISSUE ID: SCAL-14788

Description: The range of the y-axis is not taking the KPI value into consideration, so the line exists beyond the scope of the chart canvas.

Resizing a column automatically sorts it

ISSUE ID: SCAL-15101

Description: When resizing the first column of a table, the action automatically sorts the column.

Show underlying data is broken on aggregated worksheets

ISSUE ID: SCAL-15298

Description: Selecting show underlying data on aggregated worksheets sometimes returns unexpected results.

Explanation for non co-sharded join prevention

ISSUE ID: SCAL-16006

Description: When you try to add a column that causes a non co-sharded join, you get a generic error message. The message should explain how the error came about.

Exclude is broken on measures

ISSUE ID: SCAL-16961

Description: Attempting to use Exclude on measures provides an irrelevant option. Exclude for measures will be replaced with “Show MEASURE >= NUMBER” and “Show MEASURE <= NUMBER”.

Single disk failures causes cluster update failures

ISSUE ID: SCAL-17340

Description: Single disk failure in a cluster derails the entire upgrade workflow, requiring manual intervention followed by a retry of the upgrade. Rebooting helps with the upgrade.

Downloaded charts are poor quality on IE11

ISSUE ID: SCAL-17361

Description: Charts downloaded while on Internet Explorer 11 are fuzzy and difficult to read.

Supported upgrade paths

If you are running a release that begins with one of the following ThoughtSpot version numbers, you can upgrade to this release directly:

- 3.5.x, 4.0, 4.1, 4.2, 4.2.1

If you are running a different version, you need to do a multiple pass upgrade. First, upgrade to one of the above versions, and then to this version.